

# THE JOURNAL OF PHILOSOPHY

VOLUME CXVI, NUMBER 4  
APRIL 2019

*page*

181 *Acts and Alternative Analyses*

Arvid Båve

206 *Predicativism about Classes*

Kentaro Fujimoto

BOOK REVIEWS

230 Sarah Moss: *Probabilistic Knowledge*

Daniel Greco

236 NEW BOOKS: TRANSLATIONS

***Published by The Journal of Philosophy, Inc.***





---

---

# THE JOURNAL OF PHILOSOPHY

VOLUME CXVI, NO. 4, APRIL 2019

---

---

## ACTS AND ALTERNATIVE ANALYSES\*

A good case has recently been made for identifying propositions with a kind of structured act types. Unfortunately, as we will see in this paper, this “act-type theory of propositions” entails an unacceptable kind—and amount—of ambiguity in our language. Before outlining this problem, however, let me briefly introduce the theory by relating four central points about it.

Firstly, act-type theories divide into *Russellian* and *Fregean* variants. The former identifies the proposition that Socrates is wise with the act type of *predicating* the property of wisdom of Socrates, whereas Fregean variants would rather identify it with some act type directed toward the *senses* of ‘wise’ and ‘Socrates’, for instance, the act type of *saturating* the sense of ‘wise’ with that of ‘Socrates’.<sup>1</sup> (I will discuss only act-type

\* For valuable comments and help, I would like to thank Matti Eklund, Johan Gustafsson, Peter Pagin, and the audiences at the 2018 Workshop on the philosophy of Edward Zalta, Warsaw, the Kolloquium Fuhrmann, Philosophy Department at the Goethe University in Frankfurt, the higher seminar in Theoretical Philosophy at Uppsala University, and the CLLAM seminar at the Philosophy Department, Stockholm University.

<sup>1</sup> The Russellian variant is developed at length by Scott Soames, *What Is Meaning?* (Princeton, NJ: Princeton University Press, 2010) and *Rethinking Language, Mind, and Meaning* (Princeton, NJ: Princeton University Press, 2015). The Fregean variant has not quite been identified—let alone defended—in the literature, although Wayne Davis comes close in his *Meaning, Expression, and Thought* (Cambridge, UK: Cambridge University Press, 2003). In Peter Hanks, “Structured Propositions as Types,” *Mind*, cxx, 477 (January 2011): 11–52, and *Propositional Content* (Oxford: Oxford University Press, 2015), an ecumenical theory is defended, on which the proposition that Socrates is wise consists of three “component” act types: referring to Socrates, expressing wisdom, and predicating wisdom of Socrates. See also Jeffrey King, Scott Soames, and Jeff Speaks, eds., *New Thinking about Propositions* (Oxford: Oxford University Press, 2013).

Related but distinct contemporary views are found in Michael Jubien, “Propositions and the Objects of Thought,” *Philosophical Studies*, civ, 1 (May 2001): 47–62; Friederike Moltmann, “Propositions, Attitudinal Objects, and the Distinction between Actions and Products,” *Canadian Journal of Philosophy*, xliii, 5–6 (2013): 679–701; Friederike Moltmann, *Abstract Objects and the Semantics of Natural Language* (Oxford: Oxford University Press, 2013), chapter 4; and Sean Crawford, “Propositional or Non-propositional Attitudes?,” *Philosophical Studies*, clxviii, 1 (March 2014): 179–210.



## PREDICATIVISM ABOUT CLASSES\*

What are classes? More precisely, what are the objects of the second sort of second-order set theory? The objects of the first sort are, of course, sets, but classes have sets as their members and behave like sets. Allegedly, the subject matter of set theory comprises all the collections there are. If classes are collections of any kind, why can't we just count them among the subject matter of set theory? A plethora of paradoxes, however, teach us that many classes are "proper" and cannot be sets. So, what are classes after all?

## I. FOUR COMPETING VIEWS ABOUT CLASSES

In the present paper, we will particularly consider the following four prominent views about classes.

*Reductionism:* Set theory need not take classes at face value, and talk of classes in set theory can and should be reduced to talk of sets.

*Plural Interpretation:* Classes are sub-pluralities of the set-theoretic universe  $\mathbb{V}$ .

*Mereological Interpretation:* Classes are mereological parts of  $\mathbb{V}$ .

*Predicativism:* Classes are predicates of sets.<sup>1</sup>

The theme of this article is predicativism. Predicativism in its modern form is credited to Parsons.<sup>2</sup> Nowadays, however, predicativism

\*I would like to thank Volker Halbach, Daniel Isaacson, Penelope Maddy, and Timothy Williamson for insightful and helpful comments. I am indebted to Ali Enayat for invaluable communications and information about the theory of satisfaction classes, and to Neil Barton and Sy-David Friedman for answering my questions about their article. I am deeply grateful to the two anonymous referees for very extensive comments and helpful suggestions, which led to a considerable improvement of the article. Finally, I would like to express my special thanks to Leon Horsten and Philip Welch for immensely fruitful discussions on the topic of this article.

<sup>1</sup>We use the term 'predicative' in the (literal) sense of "having the quality of predicating something" or "forming or having the function of a predicate" (*Oxford English Dictionary*, s.v. "predicative," accessed February 22, 2019, <http://www.oed.com/view/Entry/149848>) here and thus call the view at issue "predicativism." The term is, however, often used in philosophy and logic with a technical (figurative) meaning, according to which something is called predicative if it is defined, expressed, or constructed without quantification over domains that include it. If one wants to reserve the term 'predicative' for the latter technical use, one could alternatively call the view "predicationalism" or "predicational interpretation of classes."

<sup>2</sup>Charles Parsons, "Sets and Classes," in *Mathematics in Philosophy: Selected Essays* (Ithaca: Cornell University Press, 1983), pp. 209–20.



So far, so familiar. One way of understanding Moss's project is to see it as drawing out the consequences for this whole package of adopting a probabilistic conception of belief, of the sort standard in Bayesian epistemology. This is not exactly how Moss argues; she offers independent arguments for moving to probabilistic conceptions of belief, assertion, and knowledge, and treats each of the parts of her package as mutually reinforcing in light of the views about the function and norms of assertion just discussed. Nevertheless, perhaps because of my own philosophical background, I find it easiest to understand Moss's project as driven by a view about the nature of belief.

Bayesian epistemologists often argue that if we try to describe someone's opinions by just listing which possibilities are compatible with what she believes, the description we give is too coarse grained. We believe some things more strongly than others. That it will rain tomorrow and that it will rain every day next week are both compatible with my beliefs (as I type, it is looking gloomy). But I am more confident of the former than the latter. While some philosophers have speculated that strength of belief might turn out to be a kind of epiphenomenon of which propositions we accept in a binary fashion, in the Bayesian tradition strength of belief is modeled directly. An agent's state of opinion is modeled by a domain of doxastic possibilities, an algebra of subsets of that domain, and a probability measure on that algebra. How strongly an agent believes a proposition is modeled by what probability the measure assigns to the set of possibilities corresponding to it.

Even if this picture is an improvement over a simple binary approach to belief, it can seem highly artificial due to the precise levels of confidence it posits. I take seriously the possibility that I will have grandchildren who will live past 200, but I would have a very hard time quantifying my uncertainty on this matter. We can be unsettled as to just which possibilities we should be taking seriously, and also as to how seriously we should take them. For this reason, many writers, Moss included, favor a generalization of the standard Bayesian picture, on which our state of opinion is modeled not by a single probability space, but instead by a set of probability spaces. In this review I will call it a "belief set," but remember that is not a set of possibilities but a set of probability spaces (each of which itself includes a set of possibilities). If I am unsettled as to whether to take a possibility seriously *at all*—if I am not sure whether something *might* be true—then it will be included in some of the probability spaces in my belief set but not others. If I am unsettled as to *how* seriously to take some possibility—if I am not sure how likely it is—then not all of the measures in my belief set will assign it the same probability.



Most of the philosophical literature on imprecise probabilities has focused on epistemology and decision theory. How should they be updated? How should they interact with utilities (potentially imprecise themselves) to inform behavior? A much less familiar set of questions concern their implications for semantics and philosophy of language. Once asked, a number of natural applications suggest themselves. If we can have opinions modeled by sets of probability spaces, then it is natural to expect that we should be able to express those opinions. So a promising research program is to look for examples of assertions that are hard to make sense of if we must understand them as conveying traditional contents, but easier if we can understand them as conveying a probabilistic content, that is, a content corresponding to a set of probability spaces.

Explicitly probabilistic assertions are perhaps the most obvious place to look. When I sincerely assert that it will probably rain tomorrow, and my conversational interlocutors accept my assertion, how should we model what happens? Moss's view suggests a natural answer. I believe that it will probably rain, which on Moss's view amounts to its being the case that all the probability spaces in my belief set assign the (traditionally understood) proposition that it will rain some highish probability, say,  $> 0.5$ . That is, there is a probabilistic content—the set of spaces that assign the proposition that it will rain probability  $> 0.5$ —that I believe, because every probability space in my belief set is contained in that content. I assert the very content I believe, and if it is accepted by my interlocutors, then the conversation will evolve accordingly. As with the traditional view, the effect of an accepted assertion is modeled by set intersection; the ultimate state of the common ground is obtained by intersecting the initial set of probability spaces in the common ground with the set corresponding to my assertion. So if my assertion is accepted, the common ground will evolve to contain only spaces according to which rain is probable. If, in addition to accepting my assertion for the purposes of conversation, my interlocutors believe what I say, then they too will be confident that it will rain.

Moss's view suggests a similarly straightforward treatment of epistemic modals. Suppose I say that it *might* rain, rather than that it will probably rain. I believe that it might rain, which is to say that all the probability spaces in my belief set include rain as a possibility. That is, there is a probabilistic content—the set of spaces that include the proposition that it will rain as a possibility—that I believe, because my belief set is a subset of that content. I assert the very content I believe, and if it is accepted by my interlocutors, then the conversation will evolve accordingly; the set of probability spaces corresponding to the



common ground will contain only spaces according to which rain is possible. And if my interlocutors believe what I say, then they too will not rule out the possibility of rain. While I have focused on the applications that are simplest to describe, Moss also provides sophisticated treatments of nested epistemic modals and the interaction of conditionals, modals, and explicitly probabilistic language.

While Moss provides specific arguments against extant attempts to model the phenomena she treats using traditional contents, in my view one does not really need to be convinced by her negative case to find the positive case compelling. I remain agnostic, leaning skeptical, as to whether one *could* provide a similarly powerful and unified treatment of the topics she addresses without appeal to probabilistic contents. But I am thoroughly convinced that, once one uses probabilistic contents to model states of opinion—which is, essentially, what many Bayesian epistemologists have been doing for decades—there is no good reason why one should not *also* expect those states of opinion to find expression in language. And so when we look and see a number of linguistic phenomena that can be straightforwardly accounted for in those terms, bending over backward to account for them using traditional contents can look unmotivated and *ad hoc*.

So far I have focused on the role of probabilistic contents in belief and assertion, but I have not yet said much about the titular topic of the book: probabilistic knowledge. In the remainder of this review I will try to give a sense of what Moss means by positing probabilistic knowledge, as well as some of its putative epistemological implications.

Bayesian epistemology is often contrasted with traditional epistemology. Part of the apparent divide is that the two approaches involve different ways of characterizing our states of opinion. The Bayesian epistemologist works with credences, while the traditional epistemologist works with beliefs and knowledge. Moss aims at a reconciliation, albeit one that may be more congenial to those that start out with Bayesian sympathies. For Moss, to have a credence *is* to have a belief—for example, an imprecise but highish credence that it will rain *just is* a belief that it will probably rain. What does it take for such beliefs to amount to knowledge? Following Williamson, Moss is skeptical that we can provide an analysis of what it takes for a belief (probabilistic or not) to amount to knowledge. But there are informative things we can say—for example, that lucky guesses do not amount to knowledge, that only truths can be known, and so on. This latter claim turns out to be trickier to make in the case of probabilistic beliefs than in traditional ones—ordinarily, we think of credences as more or less accurate, but not true or false. In chapter 6, Moss provides a deflationist-inspired account of truth for probabilistic contents that is



reminiscent of similar strategies that meta-ethical quasi-realists and expressivists have used to make sense of truth attributions for normative claims; very roughly, if  $p$  has a probabilistic content, then the claim that it is true that  $p$  will have the same probabilistic content that  $p$  does.

So we end up with the picture that some credences amount to knowledge. A highish credence that it will rain, if formed in the right way (for example, not by a lucky guess), and true (in at least the deflationist sense), can amount to knowledge that it will probably rain. This lets Moss appeal to ideas familiar from knowledge-first epistemology—for example, that actions are only appropriate when based on knowledge—and apply them to debates that seem more naturally cast in probabilistic terms. She defends L. A. Paul's claim that we cannot rationally decide to undergo certain high-stakes transformative experiences,<sup>4</sup> on the grounds that evidence about how others similar to us have reacted to such experiences does not give us probabilistic knowledge concerning how we will react. She argues that legal burdens of proof should be understood as requiring probabilistic knowledge. Racial, ethnic, and gender profiling are wrong in part because the probabilistic beliefs they produce do not amount to probabilistic knowledge. In general, the move is to take some situation in which our evidence seems to make some claim probable, but it does not seem right for that probability to guide action, and to argue that in such cases even high credences based on solid evidence fail to constitute (probabilistic) knowledge.

Before I close, I want to sound a note of skepticism regarding these applications. For reasons I have not gone into, Moss requires a hefty dose of contextualism to make her account work (see especially chapter 7). For example, she wants to secure the attractive result that it can be true to say, of a fair lottery, that each ticket is a probable loser, but also that the winning ticket is not a probable loser. And for this to work, those two claims require a context-sensitive parameter to be filled in differently. But this kind of flexibility, while attractive, creates difficulties for some of the normative implications Moss hopes to draw from her account. The same contextualism that makes it possible—in certain contexts—to truly say that each ticket is a probable loser also makes it possible to truly ascribe probabilistic knowledge—for example, concerning criminal guilt—on the basis of merely statistical evidence. Moss attempts to address this problem by positing moral requirements on how context-sensitive probabilistic language should

<sup>4</sup>L. A. Paul, *Transformative Experience* (Oxford: Oxford University Press, 2014).



be interpreted. While this seems to me the right general form for a solution to take, I was not satisfied by the particular solution offered. Moss proposes a "rule of consideration," according to which it is often morally required to keep in mind the possibility that an individual may be a counterexample to statistical generalizations she falls under. While undoubtedly true, this struck me as less explanatory than one might hope. Sometimes we think it *is not* morally required to do so—for example, when we are comfortable convicting somebody on "direct" but nevertheless fallible evidence—and the rule does not help us decide how to classify tricky cases. And even if it could be fleshed out, it probably would not help with the non-moral applications. For example, relative to some contexts, statistical information *will* let me know how I will probably experience parenthood. So if Moss wants to vindicate Paul's pessimism about the rationality of undergoing a transformative experience, she will need a different kind of explanation of what is wrong with those contexts, or why the knowledge ascriptions that can be truly uttered in them should not guide action.

I rate these as minor quibbles. *Probabilistic Knowledge* offers a persuasive case for a radical revision of some foundational ideas about the nature of content and communication. While the implications of this revision for downstream debates about statistical evidence strike me as unclear, any such implications would be icing on an already rich and rewarding cake.

DANIEL GRECO

Yale University