

KRZYSZTOF POSŁAJKO *

DEFENDING MINIMAL NON-REALISM

SUMMARY: Reply to commentators in *Semiotic Studies* special issue on *Unreal Beliefs*. I defend the position developed in this book, i.e., minimal non-realism, against various charges. First, I defend the way I reconstructed folk commitments to realism about beliefs. Second, I deal with the arbitrariness charge against my definition of belief realism. Third, I critically evaluate a counterproposal to my view, namely the idea that beliefs are socially constructed natural kinds. Then I deal with the challenges levelled against my arguments against realism about belief. Finally, I reconsider some of the consequences of minimal non-realism and consider potential alternatives to it.

KEY WORDS: beliefs, folk psychology, anti-realism, dispositionalism, natural kinds.

1. Introduction

I am extremely grateful to all the commentators at the special issue on *Unreal Beliefs* for their insightful reviews, as well as to the editors of the volume, namely Agata Machcewicz-Grad and Maciej Tarnowski. Most importantly, the presented comments show that the issue of the reality of beliefs (and, more generally, of the reality of folk-psychological states) is still of significant interest to philosophers of mind. Showing that this question is vital and there are ways to make progress regarding it was the primary aim of my project.

The reviews have forced me to rethink some of the commitments I made in the book. I also had to reconsider the strength of my position, especially relative

* Jagiellonian University in Kraków, Department of Philosophy. E-mail: krzysztof.poslajko@uj.edu.pl. ORCID: 0000-0002-7883-1651.

to its competitors. Perhaps more importantly, the reviews pressured me to reflect more on the motivations which underlie the positions I developed in *Unreal Beliefs* (Posłajko, 2024). All in all, I still think that minimal non-realism is a viable position regarding the metaphysical status of beliefs. However, what I have come to realize is that the attractiveness of the view I presented strongly hinges on several assumptions that might not be universally shared. Thus, its attraction might be limited to philosophers who share a certain, perhaps slightly idiosyncratic, vantage point. For this reason, the defense of the view that I offer in this paper should mostly be read as spelling out the basis on which the project is built.

Many worries regarding my project were raised by more than one commentator. For this reason, I decided to organize the response thematically. I will start with the problem that has been raised most frequently, namely the question of whether my rendering of folk commitments was adequate. Special attention will be given to the empirical versions of that worry provided by Kalis (2026) and Komorowska-Mach (2026). Second, I will deal with a worry that many commentators have raised, albeit in different forms, namely that my definition of realism was arbitrary. Then I will tackle the suggestion made primarily by Curry (2026) and Zawidzki (2026) that an attractive alternative to my position is the view that takes beliefs to be socially constructed natural kinds, and that my book does not contain a satisfactory argument against such a proposal. Then, I will deal with the charges that my arguments against realism are faulty, focusing on those provided by Kozak (2026) and Horowitz (2026). Next, I will look at the consequences of my position, and finally at some of the possible alternatives to it, especially those presented by Bantegnie (2026), Horowitz (2026), and Mölder (2026).

2. Folk Commitments

Several commentators (including Bantegnie, Zawidzki, Mölder, Horowitz, Kalis, and Komorowska-Mach) have complained that my rendering of folk commitments regarding the reality of beliefs is mistaken. In short, I claimed (in Chapter 7 of *Unreal Beliefs*) that the folk consider beliefs to be strongly real in the sense that folk psychology carries the commitment to treating beliefs as internal causes of behavior. This, in turn, is why I postulated a mismatch between folk-psychological commitments and the fact that beliefs are non-real.

The reviewers have nearly unanimously claimed that I overstated the strength of folk commitments, which should be seen as less realist. This sort of complaint is hardly surprising, as many of the commentators adhere to the broadly understood Rylean tradition,¹ according to which the actual commitments of the folk

¹ As Bantegnie (2026) notes, my reading of Ryle (2009) is relatively superficial and ignores more detailed interpretations of the actual works of Ryle. I must admit that I relied on the popular picture of Ryle: in my defence, I might only say that it is this popular reading of Ryle that influenced later discussions on dispositionalism.

are far less substantial than philosophers tend to interpret them as being. As I admitted in the book, the part about folk commitments was perhaps the most risky part of the project. The idea that the folk are committed to there being internal causally active states called beliefs seems open to empirical falsification. Still, I will try to defend this idea.

2.1. Folk Commitments and Anti-Realism

A common complaint (levelled by, among others, Mölder, Bantegnie, and Horowitz) was that my insistence on the substantial character of folk commitments is an unnecessary addition to my general project of minimal non-realism. The gist of the complaint is that my main thesis was metaphysical: that beliefs do not constitute natural properties in Lewis's (1983) sense. Given this conclusion, considerations about what people think about their reality should not bear too much relevance to the issue of belief-realism.

I agree that one could, in principle, make one claim without the other: Mölder (2026), following Demeter (2009), distinguishes between the issues of entity realism and semantic realism. The claim of entity realism concerns the metaphysical status of beliefs, whereas semantic realism concerns the status of folk-psychological discourse. Being an entity realist (or anti-realist) concerning beliefs does not say anything about the status of folk psychology as a discourse. There is also no need to postulate any mismatch between folk psychology and reality in order to embrace entity anti-realism. Paradoxically, as Mölder points out, an entity realist might also postulate a sort of mismatch between folk psychology and reality: if folk psychology turned out to be insubstantial in its commitments, and beliefs turned out to be metaphysically real, then we would also have a mismatch between folk psychology and reality. This kind of mismatch, however, would not warrant anti-realist conclusions of any sort.

In response to this worry, I must say that my insistence on the idea of a mismatch is rooted in my motivation to preserve Churchland's central intuition of classical eliminative materialism (1991), namely that there is some important error in the way we normally conceptualize our mental lives. This motivation, as it turns out, is somehow idiosyncratic: it is rejected not only by intentional realists, but also by most Ryleans, who consider the folk to be free of error in their conceptualization of the intentional sphere. For a typical Rylean, the errors of folk psychology lie solely in the domain of realist philosophical theories.

Horowitz (2026) provides a useful terminological proposal to make matters clearer here: he proposes labelling the claim that beliefs are not metaphysically real "belief non-realism"; he reserves the name "belief anti-realism" for the further claim of a mismatch between substantial folk psychology and the non-realist metaphysics of beliefs. Framed in these terms, it would seem that most contemporary Ryleans are staunch belief non-realists, while belief anti-realism remains as unpopular as it ever was.

In my view, even showing that belief anti-realism can be coherently phrased is a theoretical gain. If we agree that eliminative materialism is a non-starter, then there is a philosophical gain in spelling out what it would mean to discover that the folk are in error about folk psychology. Even if the claim regarding the existence of this error remains tentative, the very articulation of how to make sense of the possibility of such an error widens the spectrum of possible positions.

Bantegnie (2026) points out another problem with my mismatch thesis: my reliance on the Thomassonian framework of easy ontology (Thomasson, 2015) strongly suggests that, on my view, the application criteria of the term “belief” are superficial. If this is so, then it is not clear how my view differs from dispositionalism. In response to this, I wish to point out that on my view the concept of “belief” (along with, perhaps, other folk-psychological concepts) is a complex beast as it contains both “application criteria” and “metaphysical commitments” (Section 7.3 of *Unreal Beliefs*). The application criteria of “belief” are indeed superficial; if we focused only on them, then the dispositional analysis would be mostly on track. However, if I am right, the folk have a tendency to make certain inferences about the metaphysical underpinnings of what is attributed on a superficial basis. This means that once we decide, on the basis of superficial dispositional evidence, that *S* believes that *p*, we have an inferential disposition to attribute deep, internal, causally active state of believing-that-*p* to *S*. This mechanism is not unique to folk psychology: in general, any essentializing process in human thought has this feature (e.g., Neufeld, 2022).²

This distinction between application criteria and metaphysical commitments is also important for my project as it supports the idea that the error of folk psychology is partial. It is not that we are entirely in the wrong when we attribute beliefs: there is only an aspect of our thinking that turns out to be mistaken. This diagnosis of only a partial mistake, in turn, motivates a kind of revisionary-preservationist approach, rather than wholesale elimination.

2.2. Functions of “Because”

Some commentators (e.g., Bantegnie, 2026; Zawidzki, 2026) have also taken truck with my principal argument for the realist reading of folk commitments, namely my thesis that the fact that the folk use “because” phrases in belief explanations is evidence for the claim that the folk take beliefs to be causes. Both Bantegnie and Zawidzki note that “because” does not necessarily mean that a causal relationship is being postulated. For Bantegnie, “because” phrases can be given a dispositionalist reading, whereas Zawidzki notes that “because” phrases can be also used in post hoc rationalizations. This is a fine point, and I agree that it is possible to propose an alternative reading of “because” in this context.

² Obviously, for Neufeld and other writers, essentializing is a process that regards concepts seen as internal mental states. In my approach, this essentializing is rather a feature of discourse than thought.

Still, my defeasible hypothesis is that—at least on the level of what Dennett (1989) called folk-psychological ideology—beliefs are taken to be causes; folk-psychological explanations are taken to be causal explanations; and “because” signals the commitment to such causal explanations. Admittedly, there is neither a quick and easy argument for my claim that “because” in the context of belief-explanations functions to signal causal relationships, nor for the opposite claim of Bantegnie and Zawidzki.

The only thing that can be said is that the causal interpretation seems to be the simpler one as it does not postulate a radical discontinuity between uses of “because” in psychological and non-psychological explanations. The postulate of this discontinuity, in my view, stems from the wish to fit the practice of folk psychology to the assumed theoretical framework, rather than from the observed actual difference between folk-psychological practices and other explanatory activities.

Bantegnie argues that “because” is sometimes used to provide non-causal explanations in other contexts as well. This is certainly true, but the question remains whether such uses can serve as an illuminating model for understanding uses of “because” in folk-psychological contexts. He gives two examples of such non-causal uses of “because” outside of psychology: “The temperature in the room has dropped because the kinetic energy of the air molecules has lowered” and “7 is a prime number because it is divisible only by 1 and itself” (Bantegnie, 2026, p. 155).

In my view, these examples fail to provide good analogies to the folk-psychological case as they are examples of broadly understood “a priori” explanations. Folk-psychological explanations do not strike me as being similar to a priori ones as there is a fair share of contingency involved in the relationship between actions and the reasons we postulate to explain them. All in all, I think that the onus here is on the non-causalists to provide a good model for non-causal explanations, and to show that this model is applicable to the folk-psychological case. As long as we do not have such a model, a naïve causalist interpretation seems to be the reasonable one.

2.3. Empirical Worries

The most substantial criticisms of my idea that the folk are committed to belief realism came from Kalis and Komorowska-Mach, who, importantly, base their criticism of my view on the empirical data. Kalis (2026, p. 89) bases her criticism of my view on folk commitments on “anthropological work in various cultural and linguistic communities”. Her main conclusion from these anthropological works is that folk psychology is not a uniform practice across cultures. More importantly, many of these practices are not “causalist” in the sense that they do not seem to rely on attributions of internal causal states. Still, the folk are doing fine in such cultures: they are able to successfully navigate interpersonal interactions and predict each other’s actions.

In a sense, the idea that causalist folk psychology is not a universal phenomenon is something which speaks in favor of the minimal non-realist view. If it is true that the folk in different cultures are able to do perfectly fine without postulating causally active beliefs, then the attraction of the traditional success argument (e.g., Fodor, 1987) for belief-realism wanes as we must seek the sources of success in interpersonal interactions in something different than the postulation of internal causal beliefs.

Still, the important worry that Kalis raises for my view is the issue of whether our contemporary Western culture really is committed to causalism. I agree that it might be the case that we engage in practices which are also not based on causalist assumptions: after all, if there are no beliefs understood as internal causal states, then there has to be some route by which we come to adequately predict each other, and this route must work without postulating beliefs. But the worry of Kalis is whether we “Westerners” really do postulate beliefs as causes at all. Kalis suggests that even if we sometimes do, this has more to do with the influence of philosophy and psychology on Western folk. Such a hypothesis is obviously tempting since it amounts to saying that reification of folk-psychological states is not something that comes to us naturally, but rather is an effect of theoretical impact on the folk’s linguistic practice.

But this is not the only possible way of seeing things. Speculating a bit, I would venture an opposite hypothesis: that the meteoric rise of philosophy of mind and theoretical psychology in the “West” was a result, rather than the cause, of our Western commitment to the idea of the “internal reality” of hidden causal mental states. Realist philosophy of mind (and the later theoretical psychology) might be seen as an attempt to systematize and vindicate a very parochial cultural construct of Western culture, namely the idea that there is a realm of internal, hidden intentional states. This is obviously purely speculative, and it would not be easy to verify such a hypothesis. Still, I think that the problem of the sources of our apparent commitment to causalism, especially given the assumption that beliefs are not real, is an important one.

Komorowska-Mach (2026) offers a subtle, empirically informed (by psychological literature on mindreading) proposal regarding three levels of engagement in folk-psychological practice. She distinguishes (1) mindreading as a craft, (2) folk-psychological talk, and (3) philosophical reconstruction of folk psychology. In her view, an adequate account of mindreading as a craft does not require postulating causally active internal states (with which I concur). It is also obvious that the philosophical reconstruction often takes the folk to be committed to internal states.

What is ingenious in Komorowska-Mach’s reconstruction is the fact that she postulates another layer between mindreading as a craft and philosophical reconstruction: the level of folk-psychological talk. Komorowska-Mach characterizes the roots of this level as follows: “We package a richer, holistic understanding of somebody’s perspective, their ‘informational states’, into the noun-and-clause structure our language provides” (Komorowska-Mach, 2026, p. 105).

I think that this postulate of distinguishing these three levels of engagement is on the right track, and how Komorowska-Mach depicts what happens on these levels is a much better description of the reality of mindreading than what I offered in the book. In general, most of her proposal about the way folk psychology operates can be seen as consistent with the metaphysical picture I proposed (which is what Komorowska-Mach seems ready to admit).

The point of contention is fairly subtle: Komorowska-Mach has serious reservations about the claim that the second linguistic level actually contains substantial commitments. As she writes:

At the second level, it may seem that the folk are realists about belief, but this is merely an artefact of the constraints of natural language, which is geared toward talking about macroscopic, discrete objects [...]. The burden of proof lies with anyone who claims that these linguistic chunks express genuine ontological commitments. Absent further evidence, I take it to be entirely possible that the folk are neither realists nor anti-realists about beliefs, at least until philosophers prompt them to take a stance. (Komorowska-Mach, 2026, p. 106)

In response to this, I claim that even though the folk language does not commit us to beliefs as *entities*, its superficial reading does lead to commitment to internal *states*. The idea of beliefs as entities might strike us as an unreasonable interpretation of the folk language as it suggests an unattractive picture of beliefs as sorts of solid objects. Still, we do seem to be committed to beliefs as standing states, i.e., we do seem to think that our beliefs are something that are somehow stored in our minds and wait to be activated in the right circumstances.

The sort of reading I am offering is based on the idea that the ontological commitment of folk psychology can, to an extent, be read off the superficial linguistic form of folk-psychological talk. This superficial linguistic form might be a result of the imperfect projection of our actual practice onto the linguistic structure. Still, I concur that this form serves as the basis for reasonings the folk are disposed to make: from the results of complex mindreading processes to the postulation of internal states.

To sum up, I think that the empirical arguments raised by Kalis and Komorowska-Mach force us to admit that the reality of folk psychology is far more subtle than implied by my rough-and-ready characterization of it. Still, I think that nothing they wrote disproves my tentative claim that there is some level of engagement in folk psychology that licenses a causalist interpretation of it. This sort of engagement might be parochial, and it might be only one, not very important, aspect of the richness of folk-psychological practice. Still, it seems to be there.

3. Arbitrary Definition of Realism

Another complaint that many of the commentators share is that the project of minimal non-realism is based on arbitrary definitions of what it means to be

a realist (this was pointed out by Bantegnie, Mölder, Zawidzki, Horowitz, Kozak, and Hensel).

In the book I was ready to admit that my definition of realism was to a great extent projective. The proposed definition was intended to do two things: first, to capture what it is that philosophers care about when they debate about the reality of beliefs; second, to be a useful guide for conducting the debate further. Although I admit that this is not the only possible way to define realism, I still think my definition can be seen as useful for the debate.

3.1. Bantegnie on Dispositionalism

Bantegnie's (2026) complaint is that my definition of realism makes dispositionalism an anti-realist position by definition. He thinks this is a misunderstanding, as for him there are two versions of dispositionalism that deserve to be called "realism". The first one postulates that beliefs are multi-track dispositions that cluster together, and the clustering of these sub-dispositions is guaranteed by law-like generalizations. The second version is stronger and asserts the same claims as the first one, but adds a claim that there is a categorical basis to these multi-track dispositions. Bantegnie's criticism is that I unfairly excluded these positions from being classified as realist as I equated realism about belief with representationalism, and with the postulate of internal cogs in the cognitive machine.

I admit that I made representationalism a paradigm case of realism about beliefs, but this does not mean that I intended to a priori exclude other ways of being a realist. Most importantly, I hope I have made clear that I take the issue of realism to be a matter of there being a spectrum of positions: ranging from extremely realist to extremely anti-realist, with many possible options in between. My omission of the possible realist versions of dispositionalism, although lamentable, can thus be easily fixed. A stronger version of realist dispositionalism as described by Bantegnie should be classified as an unequivocally realist position as it meets the desiderata of "objective similarities" (provided that such a version of dispositionalism could ground the idea of beliefs having objective contents) and "entering causal explanations", even if those similarities and explanations are somehow "less exact" than on the representationalist model. One might wonder, however, whether a version of dispositionalism that postulates categorical bases for belief-related dispositions still counts as dispositionalism *sensu stricto*, but this, again, would be a purely definitional debate.

A more curious case is the case of dispositionalism that does not postulate categorical bases for belief-related dispositions, but claims that a cluster of such dispositions is connected by law-like generalizations. In this case we have a theory that does not carry a commitment to anything "internal" in the explanation of beliefs, but there seems to be something realistic about it. Such a version of realist dispositionalism might then be a prime example of a position which fits perfectly in the middle of the realism vs anti-realism spectrum. Such a view would keep some of the realist commitments (beliefs are somehow explanatory if we

agree that dispositions can serve as explanations), and there is a level of objective similarity between so understood beliefs.

In this way, my definition of realism about beliefs does not a priori exclude the kind of view Bantegnie proposes, but rather shows how they are located in relation to other more and less realist proposals. In general, I think that Bantegnie's insistence that I conflated the reality of beliefs with their internal character is a bit of a mischaracterization of my view. A representationalist view that postulates internal beliefs is just a paradigm of the realist position because it clearly meets the criteria of objective similarity, causation, and relation to fundamental properties. But this does not mean that positions other than the standard representationalist one cannot meet these criteria. Thus, I feel guilty of a failure of imagination as I have not predicted the possibility of the realist version of dispositionalism. However, I still feel that my definitions stand. I will revisit the idea of realist dispositionalism in Section 7.1.

3.2. Zawidzki on Imprecision

Zawidzki (2026) complains that I remain elusive about the exact location of beliefs, as I see them, on the spectrum of realness. He contrasts beliefs with arbitrary modal collections on the one hand, and with social kinds such as money on the other; he himself likens beliefs to "natural social kinds", such as money. I will discuss Zawidzki's positive proposal in the next section, but here I want to focus on the question of where I postulate locating beliefs. On my take, beliefs are less real than money (assuming that there indeed are non-trivial generalizations about money), but they are not fully arbitrary collections of modal objects.

If I were to postulate an analogy from social ontology, then I would suggest that beliefs as a category are somewhat similar to the category of "heads of states".³ Heads of state clearly exist: there are numerous people, past, present, future, actual, and non-actual, to whom we might predicate the property of being a head of state. There are some features that all heads of state seem to share, including certain social statuses and expectations. There are good pragmatic reasons for singling out this category, e.g., reasons concerning how heads of state should be treated according to diplomatic protocol. Yet, heads of state do not constitute any "natural kind", even if we allow social kinds to be natural. Heads of states vary immensely, both in their internal properties and in the social statuses that they are given (presidents can be either mostly ceremonial positions or hold near-dictatorial powers, etc.). Most importantly, there seem to be no important, non-trivial generalizations about heads of state, and the prospects for the discovery of such generalizations remain dim.

³ This analogy is inspired by Schwitzgebel's analogy between consciousness and "kings" (Schwitzgebel, 2025). However, I do not feel committed to the whole package of Schwitzgebel's views about consciousness.

Beliefs, on my view, are similar to such “thin” social categories. There are good reasons to single out beliefs as a category, and there is certainly a certain status that can be attached to being a believer that p ; however, if I am right, there are no interesting generalizations about this kind. More on the distinction between beliefs and natural social kinds will be said in Section 4.

3.3 Horowitz on Functional Properties

Horowitz (2026) claims that my definition of realism excludes functional properties from being natural. According to him, my definition of realism would require beliefs to have internal essences to be considered natural kinds. Thus, as beliefs are almost universally accepted to be functional kinds, then my conclusion that they are not natural kinds follows trivially. I think this is a misunderstanding. My definition of realness/naturalness does not exclude functional properties in general. Still, it puts some demands on those functional properties which we would like to classify as natural properties. Such properties would have to be projectible: there have to be some non-trivial generalizations regarding such properties. This is because we can claim that we have genuine explanations only if there are such generalizations.

It would be spurious to claim that all functional properties are natural: there are some functional properties which constitute classes whose members do not share anything more than the functional property used to define the kind in question. Take the example of “valve”: being a valve is a functional property, but there are lots of types of valves, and nothing of substance can be said about this property other than what is packed in its definition. On the other hand, there are more interesting functional properties. Being a predator is also an admittedly functional property. However, objects belonging to the class of predators apparently share lots of different properties, not only that of being an animal that hunts other animals. If indeed predators constitute such a class that shares numerous surprising properties, then “predator” counts as a natural, albeit functional, kind.

The question is then whether the class of “believers that p ” is more akin to the class of valves or to the class of predators. This translates to the question of whether this kind is such that its members share a number of different surprising properties. My answer to this question is negative, but it is not a matter of pure stipulation: it could turn out to be the other way round, and the claim of there being no natural category of believers is based on the available evidence, not on definitional fiat.

3.4. Hensel on Realism and Natural Kinds

Hensel (2026) is skeptical about my proposal to define reality in terms of natural kinds. For him, the idea that we can achieve progress in adjudicating the question of realism about beliefs in this way is mistaken, and we should revert to the original formulation of Churchland (1981), for whom the question of

realism about beliefs was to be seen as the question of the truth of folk psychology as a theory.

Hensel identifies two main problems with the natural kinds approach. The first is what he calls “the problem of disappearing substance”; the second is “the problem of disappearing reality”. The first problem goes as follows: natural kind theories together with natural kind-term semantics allow us to drop any commitment that we seem to associate with a given kind *K*. So, if we discover that there are no things that have feature *F* and we previously thought that *F* is essential to *K*, we are in no position to declare that this kind does not exist. A natural kind-style realist about *K* might postulate that we keep the commitment to *K* but drop the commitment to the idea that *K* must be characterized by *F*. In this way, a realist about *K* can survive (almost) any refutations of there being *K*.

In the case of beliefs, we have a clear example of this strategy. Once upon a time a realist about beliefs used to be committed to the idea of beliefs being “sentences in mentalese”. But what happens once the idea of mentalese is rejected? Realists about beliefs (as a natural kind) simply drop their commitment to mentalese and say that natural kind “beliefs” exist. We were just wrong about beliefs’ essential properties: we mistakenly assumed that they must be sentences in mentalese. In such a way, it might turn out that any negative result regarding beliefs might be brushed off by belief-realists.

The second problem identified by Hensel, “the problem of disappearing reality”, is that once we drop the classic micro-essentialist view of natural kinds, then kinds—understood in a less robust manner—become conventional. Moreover, it is very hard to ascertain which kinds are projectible as we still do not know which generalizations are true: it always might turn out that a kind we deem not to be projectible will become the basis of some future inductive generalizations.

I acknowledge that the problems identified by Hensel are very much real, but I do not find them unsurmountable. First, to avoid the problem of realists dodging their commitments, we should focus on what we care about when we debate the reality of beliefs. As Hensel notes, I claim that there are two essential commitments of belief realism that are worthy of its name: that beliefs have propositional semantic properties, and that they are causal explanations of actions. Of course, a realist could come and say that they would drop this commitment and still defend the reality of what they would call “beliefs”. But the reality of such understood “beliefs” would be of nearly no interest to philosophers of mind. We care about the reality of beliefs because we care about there being states that can combine being vehicles of propositional content with being causes of behavior. If there are some other states that are somehow related to beliefs so understood, then fine—there are no semantic police to stop anyone from calling them “beliefs”. But if there are no states that are at the same time propositional and causal, then we can claim victory for belief anti-realism because there is nothing that has the properties we care about.

This also resolves the worry about conventional kinds. As I argued in my earlier work (Poslajko, 2022), there are many possible precisifications of the term

“belief”, but there does not seem to be one that is objectively best. Does this bar us from making progress in the realism debate? I do not think so. If there is evidence that at least one kind that embraces our core postulates is projectible, then realism is “winning”. But as long as there is no acceptable—from the point of view of what we care about—precisification of belief that meets the criterion of projectibility, the anti-realist views retain the upper hand.

This brings us to the last issue, namely that we do not know which kinds will turn out to be projectible. This is uncontroversial, but this does not change the fact that we can make claims given the available evidence. I admitted in the book (e.g., Poślajko, 2024, p. 112) that the issue of the reality of beliefs is an empirically open one. Still, my anti-realist preference is based on what I take to be the current state of knowledge: currently, we do not have positive proof of there being genuinely surprising generalizations about the category of beliefs; therefore, we can, tentatively, accept non-realism.

All in all, there are genuine problems that arise when we tie realism to natural kindness, but these problems are not insurmountable. Moreover, the alternative that Hensel proposes—to go back to folk psychology as a theory—seems even more problematic. Gone are the days when seeing folk-psychological explanations as species of theory-based explanations was universally accepted. Today, those theorists who are still likely to compare folk-psychological explanations to scientific ones are more likely to see folk psychology as a model (e.g., Godfrey-Smith, 2005), but even this kind of assimilation of psychological to scientific vocabulary is controversial.

4. Socially Constructed Natural Kinds

Both Curry (2026) and Zawidzki (2026) locate a lacuna in my arguments. They note that I envisage the possibility of beliefs being seen as socially constructed natural kinds, and I reject this possibility (Poślajko, 2024, Section 8.6), but according to them my arguments against this view are too weak.

On their approach, socially constructed kinds can be natural because members of such kinds might share important properties, and this projectibility can be supported by social mechanisms. For Curry and Zawidzki, beliefs should be seen exactly as this kind of social natural kind. As Curry characterizes the issue: “The question, then, is whether folk-psychological processes do construct beliefs such that those beliefs go on to play a distinctive role in the causal nexus of the world (rendering them projectible)?” (Curry, 2026, p. 119).

Curry starts from the assumption that beliefs are for attributors, and that once an attributor attributes beliefs veridically, then they are able to make many successful predictions about the person to whom they attributed the belief. He contrasts this case with the case of race (an analogy I relied on heavily in my book; Poślajko, 2024, Section 7.4), because for Curry if all there is to race is the racialization process, and there are no true generalizations about racial groups, then race attributions are never veridical.

In my view, Curry's argument relies on the erroneous conflation of veridicality and projectibility. Curry seems to claim that if a given social predicate is not projectible, then it cannot be veridically attributed. In my opinion, this view is patently false. There are many social categories that are non-projectible but veridically applicable. Take the example I used above, namely that of the category "heads of state". There seem to be no interesting, surprising generalizations about heads of state. Still, it seems to me that obviously it is true of many people that they are heads of state.

Still, the question of whether beliefs-as-socially-constructed are projectible is a vexed one. It hinges on the question of whether beliefs-as-constructed-by-attribution really provide us with a basis for predicting behavior. Curry and Zawidzki certainly think that this is the case. As Zawidzki claims: "to believe that p is to adopt the very fine-grained social role of a believer-that- p , maintained in the same way as more canonical social roles, like motherhood: in virtue of deontic attitudes prevailing in communities" (Zawidzki, 2026, p. 24). The fact that believers occupy such roles, and these roles are stabilized by social mechanisms, is the basis on which we might make substantial predictions about the behaviours of believers.

I agree with Zawidzki that being a believer-that- p is a kind of social role, but I think that this role is relatively thin and not much can be inferred from it. This is because beliefs are hopelessly holistic. Let us take a look at an example of holism from Crispin Wright: "I may manifest my belief that the sea is cold by refusing to bathe (if I do not like cold water bathing) or plunging in joyously (if I do, or want to present an example of British fortitude, etc.)" (Wright, 1993, p. 22).

What such examples show is that we hold some normative expectations for believers-that- p , but I contend that we are surprisingly tolerant about deviations from "normally expected behaviour". We are well aware that people suffer from weakness of will, that they might have eccentric desires, that they might have non-standard background assumptions, etc. Thus, even in the simplest cases of beliefs, such as the belief that the sea is cold, there is a surprisingly rich body of behaviours that count as being coherent with holding this belief. This means that the role of a believer-that-the-sea-is-cold is very thin: it probably involves very little besides saying so or nodding approvingly in the relevant contexts.

Zawidzki seems to be aware of these holistic considerations, and he explicitly contrasts belief-explanations with causal-explanations (Zawidzki, 2026, p. 19). However, it is unclear to me how this non-causal-explanatory role that he envisages for belief-explanations is to be squared with the idea of there being a rich social role of a believer-that- p .

This skepticism about the richness of the social role of a believer-that- p corresponds nicely with the considerations presented in Section 2.3 above. If Komorowska-Mach and Kalis are right that the actual mechanisms of mindreading are independent of the (parochial) linguistic practice of belief-attributions, then it is no wonder that, appearances to the contrary, we gain little predictive power by making explicit belief-attributions.

All in all, Curry and Zawidzki are right that I have not provided conclusive proof that beliefs cannot be a natural social kind. On the other hand, they have not provided a decisive reason to think that beliefs are such a natural social kind. Rather, they have merely expressed their optimism about the to-be-discovered importance of beliefs-qua-social-roles for the process of social predictions. In my view, however, such optimism is not enough to warrant the claim that beliefs are natural social kinds.

5. Faulty Arguments Against Realism

Some commentators have expressed misgivings about my arguments against realism. Zawidzki (2026) claims that it would be very hard to explain the entrenchment of folk psychology if there were no genuine generalizations about belief. This seems to me to be a variation on the theme of success argument: why would we use belief-talk if the attribution of beliefs did not lead to epistemic gains? But, to repeat myself, if the underlying sub-personal processes of mindreading and mindshaping are distinct from the outward processes of linguistic belief attributions, then this argument does not lead us to any form of realism. Moreover, the very idea of entrenchment might be simply false: if Kalis is right about the scope of variance in cross-cultural folk psychology, then beliefs are not as entrenched as is assumed.

5.1. Kozak on Rationality

Kozak (2026) offers an interesting argument for belief realism, which he calls “the rationality argument”, which is different from the standard success argument. This argument is based on the general idea that norms of reasoning can alter our behavior, especially how we reason. The very aim of there being norms of reasoning, according to Kozak, is that they can guide our reasoning. So, norms of rationality can alter our behavior (broadly conceived). If this is so, there has to be some intermediary between norms of rationality and behavior: beliefs are precisely what can be said to be this bridge. How can a norm of reasoning, say *modus tollens*, influence my reasoning practice? According to Kozak, it is only because I believe that said norm is true and correct. However, if this picture is on the right track, then beliefs must be treated as causes. Only if my belief that a given norm of reasoning is treated as a cause can we then treat it as something that changes my reasoning activities. This argument is not meant to show us what kind of theory of beliefs is correct: it could turn out to be either representationalist or dispositionalist. What needs to be accepted, as Kozak argues, is the basic realist commitment that beliefs are causally efficacious.

This is an intriguing line of thought, and it brings attention to certain, admittedly controversial, features of my position. If we treat denial of causal efficacy as a central commitment of non-realism about belief (as I am inclined to do), then a non-realist such as myself must reject Kozak’s conclusion. And, as his

arguments seem sound, the only way to do it is by denying the premise that norms of rationality can causally influence behavior (this sort of response was suggested to Kozak by Maja Kittel).

Contrary to what Kozak suggests, this does not mean that norms of rationality are not *applicable* to behavior. It is vital in this context to distinguish between norms of rationality being *applicable* to humans, and norms of rationality being able to *causally influence* behavior. That norms of rationality apply to human behavior means that we can judge ourselves and others by whether our reasonings adhere to norms of rationality. As I see it, we can do this without postulating that people can intentionally change their behavior by adopting norms of rationality.

The subtle and crucial difference here is between norms of rationality being “in force” and norms of rationality being prescriptive. This distinction was crucial in the debate about normativity of meaning back in the day (e.g., Hattiangadi, 2006). Semantic norms are, uncontroversially, in force, in the sense that people can and do judge others by the standards of linguistic correctness. But to say that semantic norms are prescriptive is to say something more controversial—namely, that people ought to follow them. But if it were true that people *ought to* follow semantic norms, it would mean that they should *be able to* decide to follow them. Still, it is far from obvious that many ordinary people have ever made any conscious decision of the form “from now on I will follow rules of standard English”.

In the case of logical norms, the very question of whether it is at all possible to decide to follow the rules of standard or non-standard logic is a subject of heated debate under the banner of the “adoption problem” (e.g., Birman, 2024). The worry is, in short, that it is unclear whether one can really *adopt* a logic, because doing so seems to presuppose already using a certain logic. So, one would need logic to adopt logic.

The issue of adopting logic is a vexed one, and I do not aim to claim that it shows the impossibility of norms of rationality influencing behavior. I only mean to show that the answer to the question of whether norms of rationality actually do influence behavior is far less obvious than Kozak suggests. It is true that an anti-realist about belief must deny that norms of rationality can be adopted by an intentional act. Indeed the central tenet of belief anti-realism is that the link between the “realm of reasons” and the “realm of causes” must be severed. However, my contention is that such an anti-realist picture is not as hopeless as is suggested.

The view of how a belief non-realist views norms of rationality goes roughly like this: We pick up, by means of sub-personal, non-intentional processes, certain ways of reasoning. These ways of reasoning are then subject to social criticism, as others might be prone to treating our ways of reasoning as deficient. Then, another subtle sub-personal, non-intentional process of mindshaping might get us to reason better. In this way, the norms of rationality are still applicable to us, and we still might get to be better reasoners, but the process does not involve

intentional acts of changing norms of reasoning.⁴ I admit that denying that norms of rationality can be intentionally adopted might be a high price for belief non-realism. However, a non-realist about beliefs has a story about the role of norms of rationality, and this story might be closer to the truth than the one espoused by rationalistically inclined philosophers.

5.2. Horowitz on Fine-Grained Contents

Horowitz (2026) concurs that if we individuate beliefs by their folk contents, then beliefs should turn out to be non-natural and not causally active. However, he points out that one could also individuate the contents of beliefs in a more fine-grained psychofunctional manner. Horowitz gives Loar (1998) as an example of a theorist who distinguishes between contents of beliefs as we ordinarily attribute them and “psychological contents”, which are determined by the conceptual roles of the subject. Horowitz does not follow Loar, but he notices the possibility that if we individuated beliefs in a more-fine grained manner than we do in common practice, then such more finely individuated beliefs could play the explanatory role. So, if two subjects’ behavioral dispositions differ, then more finely individuated beliefs also differ. In this way, we could, potentially, save the idea of beliefs being explanatory because the differences in fine-grained beliefs are mirrored by differences in behavior.

My worry with this suggestion is that there is a whiff of circularity and arbitrariness about such a way of defining contents of beliefs. If we decide to reject folk attributions as a way to determine contents of beliefs, then there are very many ways in which we might determine contents in a more-fine grained manner. In fact, one could devise such a radically fine-grained way of content individuation such that no two people—nor even two distinct time-slices of the same person—could harbor beliefs with the same content. For example, the belief that Birmingham is north of Oxford that I held before I visited either of these cities might be said to have different psychofunctional content than when I had visited them because my, say, evidential situation became different. But now, as my memories of England fade, perhaps the fine-grained content starts to differ again. Or perhaps it has always been the same? How are we to tell?

Horowitz toys with the suggestion that we should differentiate the contents of beliefs on the basis of subjects’ behavioral dispositions, but this seems to be on the wrong track. If we define contents of beliefs by dispositions and then use this definition to show that beliefs are causally active, then this reasoning seems hopelessly circular. This is because we first use the causal profile to define beliefs, then we show that beliefs defined by reference to their causal profile indeed have a causal profile.

⁴ This picture is obviously inspired by the Kripkensteinan skeptical solution to the problem of rule-following (Kripke, 1982).

All in all, Horowitz does not back the standard version of realism about beliefs, and I find his remarks about the limits of the success argument most helpful—specifically, his remark that “the explanatory and predictive success of a theory only indicates the truth of those parts of it that are explanatory or active in the theory’s explanations and predictions” (Horowitz, 2026, p. 62) seems to be right on track. Where I disagree with Horowitz is the extent to which the success of folk psychology commits us to anything resembling beliefs as ordinarily understood: I will contrast my views with the positive proposal of Horowitz in Section 6.

6. Consequences of Minimal Non-Realism

6.1. Schwitzgebel on Robots

In his intriguing review, Schwitzgebel (2026), who expresses sympathy with many aspects of my view, focuses on an entirely different issue: the problem of the consequences of the general superficialist approach (to which camp my proposal clearly belongs) to the issue of the attribution of mental states to robots. His contention is that the superficialist should welcome such a possible extension of the intentional idiom: given that some future robots would become behaviorally complex enough, then our linguistic practice would be such that we would find it apt to describe them using the intentional idiom.

I think that this is an extremely important aspect of my view that was somehow underdeveloped in the book: the question of the consequences of adopting minimal non-realism about beliefs for the issue of the correct scope of belief-attributions. *Prima facie*, as Schwitzgebel suggests, the consequence is clear: if we consider beliefs to be not strongly real, then we should allow for a significant degree of liberalism in the scope of beliefs’ attribution. If there is no hidden essence to beliefs, then there seems to be no grounds to deny that artificial systems are believers. Or so it seems.

On my view, this conclusion seems too hasty, and the link between superficialism and extreme liberalism about the scope of belief attributions—although in a way natural—could and perhaps even should be severed. As I see it, there should be a way for a superficialist to say that although a certain system seems, at first glance, to be a belief-holding system, in fact it is not really one. It would be, as I reckon, an extremely uncomfortable position to hold if one could not ever claim that what we are dealing with is a system that is just a seeming believer. For example, Dennett, who for a long time was a champion of artificial systems being true believers, ultimately rejected the idea that contemporary LLMs are true believers, calling them “counterfeit people” (Dennett, 2023). So, even a staunch superficialist might get fed up and decide that what we have in front of us is not the real McCoy.

The central question that Schwitzgebel’s commentary forces us to consider is the following: On what basis can a superficialist deny that a system has genuine

beliefs and relegate this system to the sphere of mere copycats? This is admittedly a vexed question that would require much more consideration than can be given here, but let me offer some preliminary thoughts.

The opponents of superficialism have it easier when it comes to denying the status of a true believer: a proponent of identity theory can appeal to similarities (or lack thereof) in brain structure; a dualist can appeal to non-physical phenomenology; and a representationalist can appeal to considerations about the cogs in cognitive machinery. A superficialist must, however, restrict themselves to purely dispositional considerations to achieve this result. So, my preliminary proposal is that in denying that a system is a true believer, we must resort to considerations about the richness of its dispositional profile. There can be a system that has *some* behavioral dispositions that are more or less the same as those of a typical believer; but this, after reflection, could turn out to be not enough to treat this system as a true believer. To classify a genuine believer, it is necessary to have these dispositions embedded in a thick web of many other dispositions that are characteristic of other mental states that typically interact with beliefs.

Consider a robot that has some belief-related dispositions: it can answer questions and even sometimes argue in favor of a given answer. It can also guide its behaviours using its “beliefs” to achieve some pre-determined aims. But this robot has no dispositions that would justify ascribing any independent motivations or plans to it (so, it is basically a mix of a chatbot and a Roomba). So, on the basis of its web of dispositions being too thin, a superficialist could tentatively justify their doubt that this robot is a true believer.

This is what differentiates my view from Schwitzgebel’s proposal. He also treats the thickness of the web of dispositions as a criterion for belief ascriptions, but his understanding of this web is far less substantial than mine. For Schwitzgebel, for a system to meet the superficial criteria for a believer, it is enough to have a coherent set of verbal dispositions and the ability to guide its own behavior (so, for him a, LLM-powered Roomba might “deserve” to be called a believer if there is enough consistency in its verbal outputs). On my view, this might not be enough. For a system to qualify as a believer, verbal dispositions and behavior guidance should be connected to a richer set of motivational and plan-related dispositions. The robot must be interpretable as being an independent planner, capable of, as it were, being its own center of actions and having independent motivations. In this way, as I see it, a dispositionalist might set a higher bar for believers than Schwitzgebel envisaged.

Additionally, as I indicated in Section 9.3 of my book, there is also an important pragmatic aspect to the question of the scope of belief attributions. This means that we should ask not only the question of whether we *can* extend the use of belief-attributions to robots, but also whether we *should*. Given that beliefs do not have hidden essence, there is a substantial wiggle room to make the concept of belief answer our pragmatic and moral needs. So, the question we might wish to consider before we agree to call robots believers is whether this kind of change in use would actually help us with our social and moral aims.

As I said, these are just preliminary thoughts, and I hope to be able to say more about restrictions of the scope of belief attributions in the future.

6.2. Mölder on the Aim of Attributions

Mölder (2026, p. 10) raises an interesting and vexed question: What would be the aim of belief attributions if we revised the concept of beliefs along the lines that I proposed in the book? On my proposal, in order to retain the concept of belief, it would have to undergo a significant revision: we would have to drop the realist commitment to the idea of beliefs being explanatory of actions in a causal sense. Mölder then asks: What use would belief-attributions be in such a scenario if we are now using them to describe causal relationships?

This is indeed a very good question, and I can offer only a gesture at an answer here, with full knowledge that it demands a more thoughtful elaboration. The proposal would go along these lines: once we get rid of the causal-explanatory function of belief attributions, we would start to use these attributions as parts of certain stories that could render people's behaviours understandable. In this way, belief attribution practices would take an explicitly justificatory aim. Instead of asking the "Why did they do it?" question with an intent to look for causes, we would start to ask questions such as "In light of what reasons are their actions justified?"

Obviously, such a justificatory aim is already present in our current belief-attribution practice, but it is intermixed with the causal aim. So, the endgame of the proposed revision would not be to propose an entirely new function of folk psychology, but rather to purify it from the unwanted causal element. This might be seen as a highly speculative proposal, but let us remember that—according to non-causal theories of folk psychology—this is what our practice already looks like. So, my proposal is that even if folk psychology does not look like the non-causalists envisage it, we should try to make it look like that. Obviously, more should be said about the specifics of such a proposal, but this must be left for a later occasion.

7. Alternatives to Minimal Non-Realism

My proposal of minimal non-realism was intended to be an alternative to both the strong realist view on the one hand, and to error theories on the other. As it turns out, there are many other theories of belief that try to achieve the same aim, and several commentators pointed to them in their reviews. Here I will discuss some of them.

7.1. Mölder on Interpretivism

Interpretivism, fathered by Dennett (1989), and more currently championed by Mölder (2010) and Curry (2020), is a theory that places itself in roughly the

same place on the spectrum of realism about belief as my proposal. My earlier discussion of interpretivism (Posłajko, 2020) focused on the issue of judgement-dependence and metaphysical constitution, and Mölder in his commentary tries to answer precisely these challenges.

It must be admitted that Mölder's version of interpretivism, namely the canonical ascription approach, is a theory with which my view shares a great deal of affinity, and it served as a vital point of inspiration to my project. However, I still think Mölder fails to answer the most important challenge I raised to his view: to specify, in a non-trivial manner, in what sense are attitudes constituted. Mölder agrees with me that interpretivism should not overgeneralize so as to include cases such as boulders. The existence of boulders, even though they are pleonastic entities, does not involve any process of interpretation. However, on Mölder's version of the interpretivist view, the existence of beliefs is closely tied to interpretation. As Mölder puts it: "the relevant distinction lies between cases where meeting the application conditions of a term is straightforwardly empirical (as in the case of boulders) and cases where meeting them is sensitive to interpretation, as with mental states" (Mölder, 2026, p. 142–143). Furthermore,

mental state ascription is not simply a matter of applying fixed criteria in a straightforward empirical manner. Making an ascription requires constructing a coherent arrangement of information from various ascription sources. This arrangement is unique in each case of applying a mental term, such as "belief". (Mölder, 2026, p. 143)

This proposal, in my view, fails to answer the challenge fully. I am still not sure whether the difference between boulders and beliefs is, on Mölder's account, epistemological or metaphysical. Does the difference lie solely in the fact that it is epistemically harder to ascribe beliefs than to ascertain that a given rock is a boulder? If so, then this is something that nearly any middle-ground position in the metaphysics of beliefs would agree on. Thus, the epistemological reading does not distinguish between interpretivism and dispositionalism.

The metaphysical reading would provide the relevant distinction, but I have an impression that Mölder shies away from it. He invokes a "lightweight" version of constitution, but I fail to see how this version would provide a relevant difference. All in all, I think that much more would need to be done to provide a satisfying notion of constitution of beliefs, and this is, in my view, necessary to render interpretivism a satisfying alternative.

7.2. Bantegnie on Realist Dispositionalism

As was noted in Section 3.1 above, Bantegnie toys with the idea of what he calls the "realist dispositionalist" view (although he falls short of endorsing it). He considers two versions of it. The first, weaker one, postulates that beliefs are identical with clusters of dispositions, and that elements of these clusters are tied

by law-like generalizations. The stronger version postulates that these clusters share some categorical basis.

This is a *prima facie* attractive proposal and worthy of at least brief consideration, even if Bantegnie only presents it as a possible view. It is, to an extent, a simpler view than mine, as in my view belief-dispositions are somehow haphazard. So why should we reject realist dispositionalism? The answer is pretty simple—because such a view would require positive proof of there being generalizations that are responsible for clustering of belief-dispositions. My insistence on the idea that our belief-attribution practices pick up sets of dispositions that are not strongly related to each other is based on the epistemic principle that we need to have positive evidence to the effect that a given category is projectible in order to postulate it being so. Such available evidence is scant in the case of beliefs.

7.3. Horowitz on Syntax Without Content

Horowitz's view, which he also develops in his recent book (2024), is that in order to account for the success of folk psychology we should postulate internal psychofunctional states, but only postulate that these states have syntactic structures without propositional content. This is also an interesting position: it rejects the standard realist view because standard realism is tied to the idea of psychological reality of content. On the other hand, Horowitz's view is more realist than mine as it postulates that there is at least some correspondence between what we talk about in folk discourse and the internal reality.

This proposal would require much more engagement than can be provided here, but let me offer some preliminary remarks. What needs to be admitted is that a position like mine does carry a commitment to there being internal psychofunctional states which are the actual causes of human behavior. There is also no easy way to a priori exclude the possibility that such internal psychofunctional states would have some sort of syntactic structure. The question of there being such an internal syntax is an empirical one, and philosophers cannot answer it from an armchair.

What I take to be highly doubtful, however, is the idea that these possibly discovered psychofunctional states with syntax would have syntax that could reasonably mirror the syntax of beliefs as we ordinarily attribute them.⁵ In short, postulating such syntactic mirroring would require believing in some sort of pre-established harmony, for which an extraordinary amount of evidence would be needed. And, as I have argued above, what we care about in debating the reality of beliefs is whether there is really something that is similar to beliefs as we ordinarily attribute them. One cannot exclude that there is some sort of internal

⁵ I am unclear, however, to what extent Horowitz actually endorses the idea of such a mirroring. Section 5.3 of (Horowitz, 2024) does seem to imply certain skepticism about it, but more careful analysis would be needed.

syntax, but this hypothesis of internal syntax in itself does not save belief-realism, even in its weaker form.

Again, my argument relies on the idea that the lack of positive evidence for the realist hypothesis lends credence to anti-realism. However, I think that is exactly right: it is the belief-realist who incurs substantial metaphysical debts, and they need to provide us with positive, even if tentative, reasons to accept this view. And as the evidence is lacking, anti-realism should be the default position.

REFERENCES

- Bantegnie, B. (2026). Dispositionalism and Realism. *Studia Semiotyczne*, 39(1), 147–157.
- Birman, R. (2024). The Adoption Problem and the Epistemology of Logic. *Mind*, 133(529), 37–60.
- Churchland, P. M. (1981). Eliminative Materialism and Propositional Attitudes. *The Journal of Philosophy*, 78(2), 67–90.
- Curry, D. S. (2020). Interpretivism and Norms. *Philosophical Studies*, 177(4), 905–930.
- Curry, D. S. (2026) Beliefs Are Real(ly for Social Animals). *Studia Semiotyczne*, 39(1), 113–130.
- Demeter, T. (2009). Two Kinds of Mental Realism. *Journal for General Philosophy of Science*, 40(1), 59–71.
- Dennett, D. C. (1989). *The Intentional Stance*. Cambridge, MA: MIT Press.
- Dennett, D. C. (1998). Two Contrasts: Folk Craft versus Folk Science and Belief versus Opinion. In D. Dennett (Ed.), *Brainchildren* (pp. 135–148). Cambridge, MA: MIT Press.
- Dennett, D. C. (2023). The Problem With Counterfeit People. *The Atlantic*. <https://www.theatlantic.com/technology/archive/2023/05/problem-counterfeit-people/674075/>
- Fodor, J. A. (1987). *Psychosemantics: The Problem of Meaning in the Philosophy of Mind* (Vol. 2). Cambridge, MA: MIT press.
- Godfrey-Smith, P. (2005). Folk Psychology as a Model. *Philosopher's Imprint*, 5(6).
- Hattiangadi, A. (2006). Is Meaning Normative? *Mind & Language*, 21(2), 220–240.
- Hensel, W. (2026). Troubles With the Natural-Kind Approach to the Debate Over the Reality of Belief. *Studia Semiotyczne*, 39(1), 29–47.
- Horowitz, A. (2024). *Intentionality Deconstructed: An Anti-Realist Theory*. Oxford University Press.
- Horowitz, A. (2026) Poślajko on Belief Similarities and Belief Causation. *Studia Semiotyczne*, 39(1), 49–66.
- Kalis, A. (2026). The Folk Are Doing Fine (But the Philosophers Are in Trouble). *Studia Semiotyczne*, 39(1), 89–100.

- Komorowska-Mach, J. (2026). Rethinking Folk Realism About Beliefs. *Studia Semiotyczne*, 39(1), 101–111.
- Kozak, P. (2026). Why Should We Be Realists About Beliefs? The Rational Indispensability Argument for Realism. *Studia Semiotyczne*, 39(1), 67–88.
- Kripke, S. A. (1982). *Wittgenstein on Rules and Private Language: An Elementary Exposition*. Harvard University Press.
- Lewis, D. (1983). New Work for a Theory of Universals. *Australasian Journal of Philosophy*, 61(4), 343–377.
- Loar, B. (1988). Social Content and Psychological Content. In R. H. Grimm, D. D. Merrill (Eds.), *Contents of Thoughts* (pp. 99–110). Tucson: University of Arizona Press.
- Mölder, B. (2010). *Mind Ascribed. An Elaboration and Defence of Interpretivism*. Amsterdam: John Benjamins.
- Mölder, B. (2026). Belief: Minimal Non-Realism and the Constitutive Role of Interpretation. *Studia Semiotyczne*, 39(1), 131–146.
- Neufeld, E. (2022). Psychological Essentialism and the Structure of Concepts. *Philosophy Compass*, 17(5). <https://doi.org/10.1111/phc3.12823>
- Posłajko, K. (2020). Can Deflationism Save Interpretivism? *Philosophia*, 48(2), 709–725.
- Posłajko, K. (2022). The Lycan-Stich Argument and the Plasticity of “Belief”. *Erkenntnis*, 87(3), 1257–1273.
- Posłajko, K. (2024). *Unreal Beliefs: An Anti-Realist Approach in the Metaphysics of Mind*. Bloomsbury Publishing.
- Ryle, G. (2009). *The Concept of Mind*. New York: Routledge.
- Schwitzgebel, E. (2025, March 06). Kings, Wizards, and Illusionism about Consciousness. *The Splintered Mind*. <https://eschwitz.substack.com/p/kings-wizards-and-illusionism-about>
- Schwitzgebel, E. (2026). Superficialism About Belief, and How We Will Decide that Robots “Believe”. *Studia Semiotyczne*, 39(1), 159–171.
- Thomasson, A. L. (2015). *Ontology Made Easy*. Oxford: OUP.
- Wright, C. (1993). *Realism, Meaning, and Truth*. Blackwell.
- Zawadzki, T. (2026). Locating Belief on the Spectrum of Existence. *Studia Semiotyczne*, 39(1), 15–27.