

Draft of May 2025—please do not cite without permission.

Dennett as a (Weak) Mental Fictionalist

Ted Parent (Nazarbayev University)

nontology@gmail.com

1. Introduction

One of Dennett's major contributions to philosophy is his (1975, 1978, etc.) interpretivist view of belief (and other propositional attitudes). It occupies a middle ground between the so-called “hyper realism” of Fodor (1975, 2005) and the eliminativism of Churchland (1981), Churchland (1983), and Stich (1983). Roughly, Dennett's interpretivism claims that the best interpretation of the intentional facts *determines* those facts—and this might suggest a kind of *mental fictionalism*, where a certain sort of narrative determines all there is to intentionality. Yet Dennett often describes himself as a “mild realist” instead. But that may be misleading. Indeed, in a (2022) paper, he instead confesses that the answer to the question “Am I a Fictionalist?” is *yes and no*.

Going further, this paper argues that he should have just answered *yes*. This is due to a powerful objection against Dennett levied by Boghossian (2010) and Kriegel (2010). After presenting the objection, I shall offer a reply where Dennett's view is framed as a version of prefix-semantical mental fictionalism, albeit an ontologically neutral version rather than eliminativist one (cf. Wallace 2007/2022; Parent 2013, forthcoming; Parent et al., §2, in press). This is tricky, however, because we must also respect Dennett's realism concerning behavioral patterns, besides his claims about the “objectivity” of interpretation. But ultimately, the various pieces can be assembled together using a weak fictionalist framework, one which also thwarts the Boghossian-Kriegel objection, thereby suggesting a pattern that may have been obscured all along.

2. Real Exegesis

Dennett's view can be seen as having two parts—an account of an “intentional system” and an account of an intentional system having a “belief that *p*.” Both parts are based, however, in the so-called intentional stance, a strategy for generating predictions about behavior:

The Intentional Stance

1. Attribute S the beliefs, desires, etc., that it would be rational for S to have in its environment.
2. Predict that S will act in accord with those propositional attitudes.

Applying the strategy presupposes a grasp of what it is for a propositional attitude to be “rational” for S. Intentional systems theory aims to articulate this, although Dennett admits the endeavor remains rather inchoate. Yet Dennett (1981) offers us the following (quoting p. 49):

- (1) A system's beliefs are those it *ought to have*, given its perceptual capacities, its epistemic needs, and its biography. Thus, in general, its beliefs are both true and relevant to its life, and when false beliefs are attributed, special stories must be told to explain how the error resulted from the presence of features in the environment that are deceptive relative to the perceptual capacities of the system.
- (2) A system's desires are those it *ought to have*, given its biological needs and the most practicable means of satisfying them. Thus intentional systems desire survival and procreation, and hence desire food, security, health, sex, wealth, power, influence, and so forth, and also whatever local arrangements tend (in their eyes-given their beliefs) to further these ends in appropriate measure. Again, “abnormal” desires are attributable if special stories can be told.
- (3) A system's behavior will consist of those acts that *it would be rational* for an agent with those beliefs and desires to perform.

This may provoke questions, ones which have been pressed by Stich (1981), Haugeland (1993), and others, but let us leave these aside.

The two parts of Dennett's interpretivism can now be put as follows:

- (IS) *x* is an intentional system iff *x* has a behavioral pattern that is reliably predictable by the intentional stance (and not equally predictable by a simpler stance).¹
- (B) An intentional system *S* believes that *p* iff the belief that *p* occurs in the best (i.e., most predictive) interpretation of *S* (generated using the intentional stance).

¹ Dennett describes the “physical stance” and the “design stance” as alternative stances. The parenthetical above is to rule out that, e.g., lecterns have intentionality, owing to their being reliably predictable by the intentional stance. (“Lecturns just want to stay put.”) See Dennett (1975, p. 23)

And so, Dennett writes “*all there is* to really and truly believing that *p*...is being an intentional system for which *p* occurs as a belief in the best (most predictive) interpretation” (1975, p. 29).

The quotation clarifies that (B) is not simply specifying an accidental feature of believing that *p*. Nevertheless, it could suggest that (B) identifies the *only* factor in having the belief—that *all there is* to it is being optimally interpretable as having the belief. This would put Dennett at odds with Fodor’s Language of Thought hypothesis (LOT), where a belief is (at least) a matter of having some particular brain state. However, Dennett clearly does not want to rule out LOT; indeed, he claims that “provided one allows great latitude for attenuation of the basic, bold claim, I think some version of [LOT] will prove correct” (1975, p. 34).

So the “all there is” quotation was probably meant to say just that optimal interpretability is essential to belief, even though the underpinning may involve more. Nevertheless, consider a true belief-ascription such as ‘Barack believes that Malala is a hero.’ What makes this belief-ascription true if not some representational token in Barack’s head? N.B., the *content* of the token might not be fully “in the head” (Burge 1979). But it is tempting to think that a concrete, token brain-state in Barack must be part of what makes the attribution true.

Dennett (1991) would emphasize, however, Barack’s belief is akin to a center of gravity, where the latter is a point (or point-trajectory through time) which helps predict the behavior of a physical object in a gravitational field. A center of gravity for Dennett is not a material, concrete part of the object—he instead describes it as an abstractum invented by physicists, precisely for its predictive utility. Analogously, a belief is a “calculation-bound entity” or “logical construct” for

predicting the “real patterns” in a system’s behavior. (1981, p. 53).² (Dennett 1991 is the key source on real patterns, although talk of such patterns already occurs centrally in Dennett 1975.)

But belief as an abstract object again suggests an *incompatibility* with LOT, albeit in a different way. The issue here is not whether optimal interpretability is “all there is.” It is rather that a belief is identified as an abstractum, whereas LOT says it is a concretum. Yet if a belief is abstract, then it *could not be* concrete.

However, I blame this on a kind of *de dicto/de re* equivocation in English. Consider talk about “Mona Lisa’s smile” (n.b., it doesn’t refer *in res* to a smile). More relevantly, consider that a meter is a unit of measurement, a calculation-bound entity if anything is—and as such it is abstract. But the length of my desk is also a meter. Yet how can a calculation-bound entity (abstract) be the same as a real-world spatial magnitude (concrete)? The question is based on a confusion—talk of “a meter” is contextually shiftily between the two. Similarly, for Dennett, “beliefs” are abstract entities leveraged by intentional systems theory. Yet these entities may correspond to concrete particulars in the brain, in line with LOT. (Ignore the difficulties in “correspondence”.) Dennett is of course neutral on the correspondence, but there is no incompatibility here.

So what Dennett gives us (in strict terms) are the *objects of a theory*—objects that may or may not correspond to states of the brain. Still, Dennett prefers not to say that beliefs are theoretical *posits*; he regards them as abstracta rather than as *illata* (pp. 53-57) The difference seems to be one of causal potency: *Illata* are seen as causally operative in the world, at least according to the theory, whereas abstracta are not. However, despite the causal inertness associated with the term ‘abstracta’, Dennett (to repeat) allows that beliefs could end up being causally efficacious (as per LOT).

² There is a danger in equivocating between belief tokens and types. Dennett’s view is presumably about tokens, given “beliefs” could end up being concrete tokens in the brain. (A type could never be a token.) Besides, even if Michelle has the same belief-type as Barack, Dennett would distinguish the tokens by their relevance to explaining different behavioral patterns for two people. (Token behavioral patterns? Behavioral pattern types?)

Yet if he is not ready to see beliefs as causally potent, what kind of objects are they? At times, he calls them convenient theoretical *fictions*. Such “putative...states” are “idealized fictions in an action-predicting, action explaining calculus. (1978, p. 30). This requires qualification, but if intentional states are *some* sort of fiction, we might make headway if Dennett tells us what a fiction is. Unfortunately, he declines the invitation. In response to Haugeland (1993), he writes:

I wouldn't want to trot out *my* ontology and then find I had to spend the rest of my life defending or revising *it*, instead of getting on with what are to me the genuinely puzzling issues - like the nature of consciousness, or selves, or free will. The ontological status of fictional characters, haircuts, holes, and the North Pole may be deep and fascinating problems in themselves to some philosophers, but not to me... *That* is the game I am opting out of (1993, p. 212).

Soon after, he adds:

This doesn't mean that I think science is conducted in ontology-neutral terms, or that the ontologies scientists tacitly adopt don't influence (even cripple) their scientific enterprises. Quite the contrary; I think ontological confusions are at the heart of the lack of progress in cognitive science. But I don't think the way to overcome the problem is by stopping the science until the ontology is clear. Here is where we really are in Neurath's boat, and must rebuild it while we keep it sailing. How, then, do we rebuild it, if not by first developing a “systematic” ontology? By noting pitfalls, looking at analogies, keeping examples close to our attention, etc. (1993, p. 213).

Now I don't believe Haugeland wants to stop the science. However, Dennett is right that bad philosophy (more than occasionally) results from a poor diet of examples. Notwithstanding, he admits that ontological confusions are inimical to progress. Thus, even Dennett should agree that ontological confusions ought to be attended to. Yet he opts out of the game. In so doing, he apparently leaves us with the claim: Beliefs are idealized fictions, *whatever those are*.

This is disappointing not just for those sympathetic with Haugeland—it should also disappoint those sympathetic with Dennett. Interpretivism is supposed to be one of Dennett's major contributions. Allegedly, it finds a middle ground between Fodorian hyper-realism and

Churchlandian eliminativism. But all this concerns *the ontology of the mental*. So if Dennett pooos ontology, that seems rather self-effacing and probably at cross-purposes.

As indicated, however, there is more to take into consideration. Dennettian beliefs are apparently idealized fictions of a sort, yet we must also take note of his *rejection* of the ‘fictionalist’ label. The key data-point here is from his (1987, pp. 72-73):

Some instrumentalists have endorsed *fictionalism*, the view that certain theoretical statements are *useful falsehoods*... [Yet] as I said...”people really do have beliefs and desires, on my version of folk psychology, just as they really have centers of gravity.” Do I then grant that attributions of belief and desire...can be *true*? Yes, but you will misunderstand me unless you grant that the following are also true:

- (1) The gravitational attraction between the earth and the moon is a force that acts between two points: the two bodies’ centers of gravity.
- (2) Hand calculators add, subtract, multiply, and divide...
[etc.]

It is arguable that each of these is a useful, oversimplifying falsehood; I would rather say that each is a truth one must understand *with a grain of salt*. I have no official, canonical translation of that familiar phrase, but I also do not see the need for one. I would rather make my view as clear and convincing as I can by explaining why I think all belief talk has the same status...as [the examples listed]. (pp. 72-73).

So, while belief-attributions are not useful falsehoods, Hutto (2013) argues (and Demeter 2022 concurs) that Dennett could still be seen as a *weak* sort of fictionalist. Such is a view where ‘fiction’ is not used in an anti-commisive sense, but rather in a more timid, *non-commisive* sense. It is a view where to call something a “fiction” is not to insist that it is unreal. It is instead just to express neutrality about its existence. It helps here to remember that elements of “fiction” are sometimes real; Kripke (1977/2013) gives the example of Napoleon in *War and Peace*. With that in mind, it is fair to say that Dennett sees beliefs as part of a “fiction,” though again, he allows that they might exist. This highly qualified fictionalism may seem awkward, yet it would explain Dennett’s (2022) ‘yes and no’ answer to ‘Am I a fictionalist’?³

³ See also Dennett’s discussion of “two kinds of fiction” in his (2018, p. 34). Consider in addition Dennett’s (1992) repeated assertion that centers of gravity are “fictions” when comparing them to the self. But centers of gravity, recall, are the model for Dennettian beliefs.

But we have not yet mentioned that Dennett deliberately speaks of *objective facts* regarding intentionality—there is an objective fact on whether an interpretation reliably predicts a system’s behavior using the belief that *p*. True, Dennett is known for saying that when there is a tie between interpretations, there is no fact of the matter as to which interpretation is correct. But for Dennett, such cases are the exceptions that prove the rule. In the non-exceptional cases, it remains objective whether the best (most predictive) interpretation attributes the belief that *p*. So in these cases, apparently, it is an *objective fact* that the intentional system has the belief. At any rate, that is what (B) would suggest: If the righthand side of (B) describes an objective matter, then lefthand side would too. So, despite any kinship with non-commissive fictionalism, Dennett routinely self-identifies as a *realist*, albeit a “mild” one.

Let us step back a bit. Dennett’s remarks on ontology indicate:

- (D1) A commitment to “beliefs” as calculation-bound abstracta,
- (D2) No commitment to concreta corresponding to the abstracta,
- (D3) A commitment to objectively existing behavioral patterns whose prediction is assisted by the abstracta,
- (D4) A commitment to an objective fact that some behavioral patterns are reliably predicted by the intentional stance (and not some other stance), per (IS), and
- (D5) A commitment to an objective fact on whether the best (most predictive) interpretation attributes the belief that *p*.

So far, so good—this seems like weak mental fictionalism (plus realism about abstracta, patterns, and interpretations). But now, given (D5) along with (B), Dennett’s view also implies:

- (D6) A commitment to an objective fact on whether a system believes that *p*.

Granted, it is objective whether the best theory *attributes* a belief that *p*. But this apparently indicates, per (B), that the intentional system *objectively has* the belief that *p*. In turn, that suggest a physically real, state of the system. How does this square with the neutrality at (D2)?

Here we do well to recall Dennett’s remark about truth with a “grain of salt.” He again avoids details, but the phrase is patently a hedge on committing to the literal truth of a belief-

ascription. As such, we may read the hedge as in line with concreta-neutrality. But then, why affirm the “truth” of a belief-ascription in any sense? Well, notice that in (B), an intentional property of the system depends on a property of something else, a theory. Thus, the intentional property is theory-dependent. That is so, even though it is not theory-dependent *that* the theory attributes the property.

So perhaps, grain-of-salt-truth reflects that “having a belief” is not a theory-independent affair, for it is determined by our most predictive theory. Or rather, since “determined” insinuates that this is the *only* thing that matters, we should say that grain-of-salt-truth reflects that having a belief is “metaphysically correlated” with the best theory’s attribution of the belief. If one tolerates the language of “grounding:” The belief is *not grounded* in our best theory’s attribution of the belief, yet the system has the belief all and only those possible worlds where the best theory attributes it.⁴

But is the “best theory” the most predictive among *actual* theories, or the most predictive theory among *all possible* theories? These would surely be different, and so, would determine different beliefs for the system (even in run-of-the-mill cases). Kriegel (2010), whose objection to Dennett we discuss next, reads Dennett as intending the latter. This would put Dennett’s interpretivism closer to Davidson’s (1973, 1974, 1983); they would both champion a kind of “ideal observer theory” of belief (Kriegel’s terminology). The ideal interpreter would be one who knows *all facts* under purely non-intentional descriptions, and uses the intentional strategy in the ideal way (Kriegel 2010, p. 114) This way of seeing things leads Kriegel to formulate (B) slightly differently:

(B*) S believes that *p* iff S has a behavioral pattern that would prompt an ideal interpreter to attribute the belief that *p* to S.

The ideal-observer reading would comport well with Dennett’s judgment that the folk use psychological concepts that are borderline incoherent (1981, p. 47). Intentional systems theory is

⁴ By the way, even Quine (1973, §3) accepted the metaphysically charged ‘in virtue of’ as indispensable to theory-building. (Such an expression is often seen as denoting the grounding-relation). That is so, even though grounders often pick Quine as their arch nemesis, as in Schaffer (2009).

accordingly seen as a successor to folk psychology.⁵ So, one may well expect Dennett to avoid actual-world interpreters, and we shall henceforth assume as much.

3. The Regress Problem

The objection raised by Kriegel (and independently by Boghossian 2010) takes the form of a regress. The key observation is that a belief attribution to Barack, on the part of an interpreter, is itself a belief. It the interpreter's belief about what Barack believes. It may be part of an *ideal* interpretation of Barack, but it is a belief nonetheless. And so, (B*) would imply that the interpreter's attribution requires an ideal interpreter to attribute the attribution.

It might be debated whether a belief-attribution is always a belief, but it is at least an intentional state of some sort—and for convenience, let us continue to say it is a belief. Now the question is: How does Dennett understand *that* belief? Is the ideal interpreter attributing the attribution to herself? This would amount to a vicious circle: Prior to having *any* belief, the ideal interpreter would need to have a self-attributing-belief—which is incoherent. So apparently, there needs to be a second ideal interpreter who interprets the first interpreter. Yet for parallel reasons, the second interpreter would require a third, and so on.

This regress is not an artifact of the ideal-observer reading of Dennett. In fact, Kriegel's point is even easier to make if we opted for the “best actual observer” reading of Dennett. For the ideal-observer view could say against Kriegel that the regress is not problematic, given that it occurs only in other possible worlds. After all, there is no ideal interpreter in the actual world; hence, the

⁵ Still, we saw that Dennett (2018) emphasizes the centrality of the intentional stance to our lives—and famously, Dennett (1975, p. 21) declares that the folk use the intentional stance successfully all the time. These tell against the ideal-interpreter interpretation. On the other hand, even if the folk use the intentional stance, it does not follow that *intentional systems theory* is used by the folk. The theory is based on the strategy, but it is more than that. But of course, the ideal-interpreter reading remains disputable; for a different view, see Curry (2021).

ideal interpreter would live in some non-actual world. And the ideal interpreter of that interpreter might be placed in yet a third world. So potentially, the infinite series of interpreters is spread across the infinity of possible worlds—and that is not obviously a problem. (This makes Dennett’s position much more “metaphysical” than suits his temperament but let that pass.)

Notwithstanding, Kriegel offers an additional argument, intended to convince that the counterfactual regress is bad news as well. In brief, the explanation of one instance of an intentional kind leans on a different instance of the kind (albeit an instance housed in another world). This means “there is no genuine account of intentionality...we do not yet grasp the very first intentional fact” (2010, p. 120). In more detail, if we wanted an explanation of how *any* instance of intentionality is possible, then the explanation is circular. On the other hand, if we wanted to explain only how intentionality is *actual*, we make things more mysterious by appealing to non-actual intentionality (and never mind what explains *that*.)

Granted, explanation has to stop somewhere. But the awkwardness is that an intentional state is explained by something of the very same kind, viz, an intentional state. True, the latter is different in one respect—it is non-actual. But the nonactual feature is, at best, explanatorily irrelevant. Progress may come by explaining the non-actual state, but the most we get here is a *third* non-actual intentional state. This is a bust.

4. Ideal Interpretation

However, there is a rather simple question to ask at this point. Why has the focus shifted from interpretations to *interpreters*? Indeed, (B) originally used the former term and not the latter. I will later hypothesize on the shift, but for now, my claim is that Dennett can thwart the regress if he reclaims the talk of interpretations rather than interpreters.

Thus, let IT_t^S be an Ideal Theory for an intentional system S, relative to a time t .⁶ At minimum, let IT_t^S be a deductively closed theory which includes:

- (i) The set of true, non-intentional descriptions of all facts obtaining at or before t , and
- (ii) A set of belief ascriptions for S at t , generated per the intentional stance in conjunction with (i) [plus other principles of folk psychology; see below]. The relevant set entails the weighted-most (atomic) truths about S's behavior after t , compared to any other set of belief ascriptions for S at t .

Now this will hardly settle a unique interpretation of S, even in run-of-the-mill cases. It is necessary, for example, to say that IT_t^S is a theory that contains (i) and (ii) which also maximizes theoretical virtues (simplicity, scope, conservativeness, etc.).

In addition, it would be a mistake to ignore Lewis's (1974) ingenious tactics, developed for the purpose of radical interpretation. One could see these as a supplemental to, even if partially coinciding with, the intentional strategy. But the purpose would be to make more robust the process by which belief-ascriptions in (ii) are selected. Radical interpretation, for Lewis, also begins with access to all facts under non-intentional description, or what he calls the physical facts P. In radically interpreting a subject, Karl, the three variables of interest are:

M = the meaning of Karl's utterances
 Ao = Karl's propositional attitudes as expressed in our language, and
 Ak = Karl's propositional attitudes as expressed in Karl's language (ibid.).⁷

With characteristic alacrity, Lewis declares: "the problem of radical interpretation [is] as follows. Given P... solve for the rest" (ibid.).

The access to *all* physical facts makes Lewis akin to the ideal observer Dennett, and indicates that the problem of "radical interpretation" is rather different than the problem from Quine (1960, ch. 2; 1968). Lewis confirms, saying: "I am not really asking how *we* could determine these

⁶ Per usual, a theory is a set of formulae, an abstractum with other linguistic abstracta as members. This is admittedly metaphysical, but Dennett already invokes abstracta, as do the sciences in general (a fact that Quine famously bemoans).

⁷ In a postscript, added in his (1983), Lewis clarifies that this glosses attitudes *de se*, which are not really attitudes to a proposition. But for our purposes, we may safely ignore this.

facts. Rather: how do *the facts* determine these facts? By what constraints, and to what extent, does the totality of physical facts about Karl determine what he believes, desires, and means?” (p. 111).

The brilliance of Lewis’s strategy is in aggregating together the relevant hypotheses about Karl with a general semantical-psychological theory, allowing all these elements to interact and modulate each other to reach coherentist equilibrium. (In fact, Lewis describes a coherentist, holistic method as only one of three methods for approaching the problem, but I shall leave the others aside.⁸)

The principles of the general semantical-psychological theory, as identified by Lewis, are summarized thusly:

Principle of Charity: S has the propositional attitudes it ought to have, judged by what propositional attitudes *we* have, or what we would have had in S’s situation (including S’s life-history, training, etc.). This includes errors that we judge to be understandable in S’s situation.⁹

Rationalization Principle: The beliefs/desires ascribed to S ought to give S good reasons for S’s behavior.¹⁰

Principle of Truthfulness: S should respect conventions of truthfulness and trust that make possible communication in S’s language; cf. Lewis (1969).

Principle of Generativity: Fill in M by means of recursive clauses from a (finite) number of base clauses. The latter identify the meanings of certain morphemes in Karl’s language, and from those, finitely many recursions determine the meanings of more complex expressions.

Manifestation Principle: Propositional attitudes in Ak should normally manifest in Karl’s linguistic dispositions.

Triangle Principle: Karl’s propositional attitudes are the same whether expressed in his language or in ours.

⁸ Lewis himself favors a more piecemeal method, where we first use P (especially re: Karl’s behavior) to fill in Ao, guided by Charity and Rationalization. After that, Ao aids filling in M, with Truthfulness imposing the strongest constraint, followed by Generativity. Finally, Ao and M lead us to Ak via the Triangle Principle. (He says that Manifestation gets satisfied *eo ipso*.) But it is unclear why Lewis prefers this method. His use of Triangle might seem holistic, but it helps solve for Ak only, not to modulate Ao or M. Granted, people sometimes interpret piecemeal, but recall that Lewis is not describing how actual people radically translate. Yet as Bar-On (1992) argues, a desideratum on explaining semantic facts is to explain how speakers can *come to know* these facts in the language-acquisition process. (Unlearnable semantic facts could not be sustained in a language.) The greater epistemic fidelity in the “holistic” account may thus be best for Lewis’s metaphysics as well.

⁹ Note the difference between Lewis and the biology-based attribution of rational belief/desire from Dennett. Stich (1981) departs from Dennett in the same way.

¹⁰ Lewis’ claim here seems too strong, for reasons related to Stich’s (1981) objections to Dennettian rationality (and despite Lewis’s kinship with Stich; cf. the previous note). The problem in brief is that the folk *routinely* deviate from ideal rationality. Lewis makes things worse by explicating S’s “good reasons” in decision-theoretic terms, despite his pleas for the “commonsense” status of decision theory. Even though some *parts* of decision theory may be commonsense, it does not follow that *decision theory* is commonsense.

Lewis surely did not see these as exhaustive of the best semantical-psychological theory. For example, the theory would probably claim less global principles like those from Churchland's (1979) person-theory or "P-theory. Most relevantly (quoting pp. 92-93):

- Persons who believe that P tend to assent to P when queried.
- Persons who believe that P, where P elementarily entails that Q, tend to believe that Q.
- Given normal attention and background conditions, persons tend to perceive the observable features (i.e. the normal colors, shapes, textures, smells, sounds, and configurations) of their immediate environment.
- Barring preferred strategies and/or incompatible wants, persons who want that P, and believe that Q would be sufficient to bring it about that P, tend to want that Q.
- If a person wants that P, and believes that Q would be sufficient for P, and is able to bring it about that Q, then, barring preferred strategies and/or incompatible wants, [s/]he will bring it about that Q.

No doubt, some of these overlap with both Dennett's and Lewis's own claims, but redundancy is ok if it ensures that nothing important is overlooked. Indeed, I would submit that each of Dennett, Lewis, and Churchland add something unique, despite any overlap.

Again, we may see general principles of semantics/psychology as part of the Lewisian and Dennettian strategies for selecting the belief-ascriptions required in (ii). I have detailed the ideal theory to minimize hand-waving, but some open-endedness is practically unavoidable. In principle, however, IT_t^S is envisioned as robust enough to determine a uniquely best interpretation of an intentional system, at least for run-of-the-mill cases (though let us ignore Quinean skeptical interpretations).¹¹ This is to legislate against interpretive "gluts" for S—we assume that (ordinarily) S is not attributed incompatible beliefs by two interpretive theories tied for first place.

However, there may remain interpretive "gaps." The uniquely best theory might be incomplete; it may give no verdict on S's propositional attitudes in certain situations. Yet for my

¹¹ I must add that Dennett apparently underestimates Quine's argument, speaking as if it applies only to exceptional cases when different interpretations of S are highly relevant alternatives (see Dennett 1975, p. 40). In fact, Quine's argument applies *ubiquitously* and concerns alternative interpretations that are far from commonsense ("temporal stages of undetached rabbit parts"). I might add that Quine's skeptical interpretations can be argued to be just as "ideologically parsimonious" as commonsense interpretations; *vide* Goodman's (1955) grue paradox. These are not simple issues.

part, this is what we should expect. Some gaps seem real. Does Joe Biden *believe* Euler's identity if he only vaguely recalls hearing about such a thing? Is a reflex to shield your head from a projectile *unintended*? Does Miley really want to quit smoking but is unable to—or is she able to, but doesn't really want to? Granted, not everyone will find these plausible as indicating psychological gaps, but the point is that, if gaps are predicted by our theory, this need not be a fatal objection.

Appealing to our idealized theory, we may update the formulation of (B) thusly:

(B**) S believes that p iff, according to IT_t^S , S believes that p .

As with (B), it is assumed that the righthand side concerns objective facts, and so ditto with the lefthand side.

5. Prefix-Semantical Mental Fictionalism

What I wish to emphasize now is that (B**) has precisely the form of a “prefix semantics” from Lewis (1978). Recall that in Lewis, the goal is to explain the sense in which ‘Sherlock Holmes lives in London’ is a correct thing to say, even though it is not actually true. Lewis suggests we can do this provided that we analyze the sentence as containing an implicit “story prefix,” which is made explicit in the following:

(H) ‘Sherlock Holmes lives in London’ is true iff, according to the story, Sherlock Holmes lives in London.

The “story prefix” on the righthand side blocks any implication of Holmes being in (real world) London. Yet the facts of the story determine a real sense in which the Holmes-sentence *is* true.

Different stories will be relevant for different ficta, but (H) exhibits a template for making sense of such discourse. And the thing to notice is that (B**) fits that template.¹²

¹² In the postscript to Lewis (1974) in his (1983), Lewis himself notes the parallel between his semantics for fiction and his view of psychological states.

One well-known problem for prefix-semantical fictionalism, associated with Brock (1993) and Rosen (1993), is that (L) seems to yield an ontological commitment to Holmes after all. In brief, Lewis affirms the righthand side of (H)—but per (H), this is to affirm the lefthand side. And affirming the lefthand side amounts to saying *that Holmes lives in London*. However, there are now a half-dozen or so responses to the objection, and both Brock (2002, n. 9) and Rosen (1995) have conceded Noonan's (1994) response. My own preference is for a reply based in Nolan (1997), Woodward (2005), and Liggins (2008), where we stress that 'Homes lives in London' is merely *shorthand* for the story-prefixed version of the sentence, viz., 'According to the fiction, Holmes lives in London'. Caveat: The fictionalist sometimes opts for the face-value reading of 'Holmes lives in London', as when she says that the sentence is untrue. But in other contexts, she regards it as elliptical (roughly speaking) for the story-prefixed counterpart—and that *is* literally true.

But while this may be apt for Lewis and Holmes, it is less apt for Dennett and belief. One minor issue is, again, that 'fiction' for Dennett should be non-commisive; thus, to say that a belief-ascription is derivable from a "fiction" is not to indicate that it literally false otherwise. But the big problem is that Dennett wants belief-ascriptions to be *objectively true*, as emphasized at (D6). This may suggest that Dennett would reject the fictionalizing interpretation of a belief-ascription and simply affirm its objectivity. Yet that conflicts with his concreta-neutrality at (D2).

Claim: Dennett can have it all by going ahead with fictionalizing belief-ascriptions. For he can simultaneously recover the "objective truth" of the ascription in terms of the objective facts about what the fiction says. In one sense, it *is* objectively true that Sherlock Holmes lives in London, even though Holmes does not live in *real-world* London. But one fact about the honest-to-goodness real world is that Holmes lives in London *according to the Conan Doyle stories*—a point made by

Sainsbury (2010).¹³ Adapting this stratagem, Dennett thus can say: It is objectively true that, according to the psychological fiction IT_t^S , S believes that p . The prefix-semantical approach, then, would be “S believes that p ” in Dennett’s mouth is shorthand for this, at least in many contexts.

Dennett can also add that the belief exists objectively as an abstractum. That may be restate (D1), but the word ‘objective’ is not found there, and could be included to tighten things up. Would Dennett then be satisfied with the Sainsbury idea about “objective fiction,” supplemented by (D3) - (D5) and a possibly strengthened (D1)? Is that the right amount of objectivity? I cannot be sure, but it seems to capture the essentials of what he wants to say.

6. Deflationary Fictionalism

We have presented Dennett as endorsing weak, prefix-semantical mental fictionalism (*cum* realism about abstracta, patterns, and interpretations), with clarifications about the kind of “objectivity” intended in (D6). But further refinements are needed. The glitch is the talk of what a fiction *entails*—for “entailment” is a semantic notion, defined in terms of truth, hence, representation. (What a fiction “entails” is how it *represents* things) So the talk of entailment seems to assume the existence of representation. And it seems to assume its existence not merely “in the fiction,” but in the real world. Dennett is ultimately trying to describe the real world after all—even though the real world might contain fictions and complications arising therefrom. But if his description of reality includes describing “entailments,” then he would be ontologically committed to intentionality.

This is a version of the “cognitive collapse” objection for mental fictionalism more broadly. The fictionalist eschews ontological commitment to intentionality, but apparently requires

¹³ Hutto (2013) also invokes Sainsbury’s point to help interpret Dennett.

intentional notions to explicate ‘according to the fiction’. As an attempt to fend this off, mental fictionalists sometimes point to future cognitive science for a non-intentional account of fictionality, whatever it may be (e.g., Wallace 2016, 2022; Joyce 2013). But passing the buck seems unsatisfying (Parent forthcoming). A better approach is to replace the semantic terms in the fictionalist account with non-semantic counterparts, so to yield a kind of deflationary inferential role semantics as described by Field (1994a, b; cf. Field 1977). Hence, if we let ‘true*’ express the deflated notion of truth, a principle like (B**) should be transmuted thusly (where ‘ $\vdash$ ’ expresses the derivation-relation):

$$(B\uparrow) \text{ “S believes that } p\text{” is true* iff } IT_t^S \vdash \text{ ‘S believes that } p\text{’}.$$

This assumes that IT_t^S is given in a purely formal language, attended by suitable derivation rules (more details are in Parent et al., in press, §2.)

The important consequence of deflationism is that the truth* of a sentence does not indicate some robust accurate-representation-relation between the sentence and the world. To label a sentence as true* is only to license certain disquotational inferences. However, some have forcefully argued that deflationism thereby makes it impossible to express ontological commitment (Boghossian 1990, p. 178; Dreier 2004). Uttering a declarative in a deflated language does not amount to *representing* the world in some specific way. It just amounts to tokening a syntactic string for use in a formalistic engine. This means that, ontologically speaking, the truth* of ‘Sherlock Holmes is a fictional character’ would register the same weight ontologically as the truth* of ‘The Earth exists’, namely, zero.

Yet all this may be quite apt for Dennettian belief-attributions, and such a thing has been defended previously by Kukla’s (2018). Re: the inability to express genuine ontological commitment, Kukla claims that this is unproblematic, for discussions about ontology are at bottom

concerned with practical questions. One immediately thinks of Carnap's framework-metaontology, but Kulka clarifies (p. 28):

This is not to reduce metaphysical questions to pragmatic epistemic questions; we can of course be *wrong* about what it is we are dealing with. It is rather to deny that there is any *separate question to be asked* about the literal reality of something beyond questions about whether it is there to be coped with.

So unlike Carnap, Kukla does not hold that existence-claims are true merely because of the linguistic conventions that define a framework. But for her, it remains that we cannot reasonably deny the existence of x in a context where x is at the center of our plans and projects.

Dennett's (2018) reply to Kukla applauds her emphasis on the intentional strategy's significance in the practical sphere. Notwithstanding, he shows resistance to the claim that talk about what exists aligns perfectly with talk about what concerns us. He reports (p. 34):

I am tempted by [Kukla's] view...but I do see a few problems...People *engage* with Santa Claus, shaping their children's early lives to include many close encounters (of sorts) with him, but it seems to me that no matter how richly Santa Claus might come to dominate our lives in the future, he could never cross the threshold and become as real as the table that my coffee is sitting on.¹⁴

I would add that the pragmatist amendment, whatever its merits, seems non-responsive to the problem which occasioned it. The problem was to explain how we can express ontological commitment *at all*, whether relative to current purposes, or not. Perhaps, however, Kukla's idea is that we never express "ontological commitment" as such; we only express practical commitment. But this would seem to revert to the purely Carnapian line where ontological disputes reduce to disputes about what language games to play, given our current interests.

Carnap is not to be sneered at—yet a purely Carnapian line would not be congenial to Dennett either. The fact is that intentional systems theory explicitly claims ontological commitments, viz., to abstracta, patterns, and interpretations. But again, the problem is that in a

¹⁴ Or a grain of sand, to take an even more inconsequential object.

deflationary language, ‘Abstracta lack spatiotemporal location’ would be *void* of ontological import. And unfortunately, intentional systems theory *must* be couched in a deflationary language, to avoid semantical notions like entailment. So the question remains: How is the theory ontologically committed to anything?

Prefix semantics to the rescue. It turns out that the semantics enables ontologically significant distinctions even within a deflationary language, although admittedly, *full-blown Realist* commitment remains inexpressible (cf. Parent, forthcoming). Yet there is an important distinction between deflated discourse that is hedged with a story-prefix, verses a deflated discourse without such prefixing. Thus, if Dennett utters ‘Abstracta lack spatiotemporal location’ in a deflated language, this will still fail to *represent* the bona fide existence of abstracta or their features. However, it would still be importantly different from an utterance of ‘It is objectively true that S believes that *p*.’ For if we take (B†) seriously, the latter is shorthand for: It is objectively true* that $IT_t^S \vdash$ ‘S believes that *p*’. That is quite thin—its truth* indicates only that a specific syntactic string is derivable from IT_t^S . Most tellingly, the truth* of the belief-attribution does not derive the flatfooted ‘There is at least one belief’. It at most derives ‘ $IT_t^S \vdash$ ‘There is at least one belief’.’ So from the perspective of the deflationary language, the utterance is thoroughly non-committal regarding intentionality and S, even though it may have looked otherwise. It ends up expressing nothing about intentionality *at all*, not even in a deflated way.

As a contrast, ‘Abstracta lack spatiotemporal location’ is not elliptical in any sense for a story-prefixed sentence. So it will not just have as (non-trivial) formal consequences that some linguistic strings are derivable from some theory. Rather, its formal consequences include *unhedged* sentences such as ‘Abstracta are atemporal’, and ultimately, ‘Abstracta exist’. Even from the perspective of the deflated language, ‘Abstracta exist’ has more ontological significance than

‘According to the fiction, abstracta exist’. Again, the inferential role of ‘Abstracta exist’ is not limited to various derivability claims. That is a kind of ontological “significance” even if it is not the capital ‘R’, Realist kind of significance. We could add that the truth* of ‘Abstracta exist’ has first-order consequences for the model of the deflationary language—yet we must hasten to add that such a model is not necessarily identical to the real world, and that the model itself is characterized in a deflationary metalanguage. Thereby, the truth* in the model of ‘Abstracta exist*’ only licenses disquoting inferences between the metalanguage and the object language. Nonetheless, deflationists speak the metalanguage without guile, and so this applies to utterances of ‘Abstracta exist*’, despite its have only a deflated semantics.

Objection: Since the utterance of “S believes that p ” is ontologically minimal even relative to a deflationary language, it seems unfaithful to Dennett’s talk of the “objectivity” of “true” belief-ascriptions. I confess I find the objectivity-rhetoric a bit of a red herring. Yet in line with the Sainsbury idea, it remains objective whether a particular sentence is derivable from IT_t^S . Granted, deflationary truth* makes this “objective truth” deviate from expectations. Effectively, it just licenses disquotational inferences between the story-prefixed sentence ‘ $IT_t^S \vdash$ ‘S believes that p ’ and an ascription of truth* to the sentence. But guileless utterance is again the norm,. It is not as if the deflationist would utter the story-prefixed sentence tongue-in-cheek, and ditto with the shorthand version “S believes that p .”

Nevertheless, granting the ontological interest in truth* versus truth* in fiction, this still leaves open what is *actually-factually-no-fingers-crossed TRUE*. Dennett indeed uses an inflationary language at times to express that the hard sciences give us Truths. And yet, his broader attitude toward manifest objects (e.g., middle-sized dry goods) is rather non-committal. He allows the

possibility that many object represented in ordinary discourse really exist, but remains fairly neutral.

Dennett (2017) writes:

[T]he items in the official ontology of the scientific image really exist, but solid objects, colors, ... words, and so on, don't really exist. They are useful illusions, perhaps, like the user-illusion of the desktop icons. The patterns of colored pixels on the computer screen are real, but they portray entities that are as fictional as Bugs Bunny and Mickey Mouse. Similarly... the manifest image has some reality as a collection of images...but it is a mistake to think of the "things" we interact with...as reality.

That's...a *version* of what I have said about the manifest image...: a user-illusion brilliantly designed by evolution to fit the needs of its users. My version differs only in being willing and eager to endorse these ontologies as ways of carving up *reality*, not *mere* fictions but different versions of what actually exists: real patterns. (p. 222)

It is unclear, however, how a solid object might both exist as a pattern in mind-independent reality *and* a pattern that is part of a user-illusion (though 'illusion' is surely meant in the non-commisive sense). The problem is that there are two patterns, one subjective and one objective. Yet there is only one object, and one thing cannot be identical to two. So he most likely means: The user has a *representation* of solid objects, and those intentional objects might *correspond* to objects in the Real World, i.e., the representations might be actually-factually True.

Such representationalism is suggested even more strongly in the claim that objects in scientific ontology really exist. It insinuates that science tells us things *correspond to* reality in virtue of their non-deflationary semantic powers. But that would be a commitment to semantically charged representations, rather than representations as per a deflationary approach. The very project of trying to describe what is True would betray Dennett's neutrality on whether representational content is real.

This is a variant of the argument from Boghossian (1990). Boghossian was targeting eliminativists, taking them to task for rejecting representation and yet committing to Truths in science. But Boghossian's point also works against Dennettian neutrality. My own view is to take Boghossian's critique seriously, and to opt for a principled *quietism* on what is Really Real.

Quietism is admittedly not the same as agnosticism, although agnosticism is obviously implied. Yet quietism includes a norm against trying to dispel agnosticism.

To be clear, quietism is not a rejection of LOT if it is formulated to fit the deflationary temperament. The deflationaist could hypothesize that concrete symbolic tokens are somehow functionally realized in the brain, and even accept this hypothesis as theoretically superior (whatever that involves). But the hypothesis would be formulated in an ontologically deflated language, as would every other hypothesis.

The deflationist does not forego ontology entirely; we have seen how to identify deflated discourse with some ontological interest. Of course, there is the metaphysical urge to plumb the depths of Reality further—but the fictionalist-cum-deflationist-cum-quietist suppresses that (cf. Parent 2015 and forthcoming). On pain of inconsistency, there is no commitment to what is True, for there cannot be a commitment to Representation in a non-deflated sense. Such downgrading of ontology is quite consonant with the attitude from Dennett we have seen. So though it requires distance from scientific Realism, quietism might ultimately fit better his metaphysical antipathy.

7. The Nonexistence of the Regress

A deflationary Dennett avoids Kriegel's regress, for if intentional systems theory occurs in a deflationary language, then a belief-attribution does not require some prior instance of intentionality. The origin of the regress, I suspect, lies in the thought that an interpretation is an *activity* rather than a mere sentence. I can grant that such an activity exists, but once we have an interpreter, we are naturally inclined to apply the intentional stance to *her*. We thus regard the activity under a further assumption, namely, that it is the product of an *intentional agent*, one who is a source of “underived”

intentionality, who supplies her interpretation with content. Of course, my suggestion for Dennett would be to reject that view of interpretive activity.

Still, our well-entrenched training in the intentional strategy makes this difficult. And it is not just a psychological obstacle, but a philosophical one. If a person's interpretation ultimately lacks content, then in Wittgenstein's phrase, the interpretation seems to "hang in the air" along with what it interprets. (Nothing would be accomplished.) Yet the interpreter as a real sanctum of intentional content is exactly what generates the regress. The way to avoid the regress, therefore, would be to embrace the possibility of contentless interpretation.

I stress the *possibility* of contentless interpretation, for the present fictionalism is again neutral on whether intentionality exists in a concrete, theory-independent way. But if the regress is to be avoided, it behooves us to embrace this as a real possibility. And the fictionalist-deflationist approach not only supplies details here; it also helps bring together Dennett's multidimensional perspective. The sole departure is the quietist attitude toward science—but as argued, his neutrality is unstable without this. Yet deflationary prefix-semantical fictionalism, along with quietism, still has much to offer him: Theoretical clarity, regress avoidance, and compatibility with his broader metaphysical temperament.¹⁵

Bibliography

- Bar-On, D. (1992). Semantic verificationism, linguistic behaviorism, and translation. *Philosophical Studies* 66: 235-259.
- Boghossian, P. (1990). The status of content. *Philosophical Review* 99: 157-184.
- _____. (2010). The transparency of mental content revisited. *Philosophical Studies* 155: 457-465.
- Brock, S. (2002). Fictionalism about fictional characters. *Nous* 36: 1-21.
- Burge, T. (1979). Individualism and the mental. *Midwest Studies in Philosophy* 4: 73-122.
- Churchland, P.M. (1979). *Scientific Realism and the Plasticity of Mind*. New York: Cambridge University Press.

¹⁵ For valuable feedback on this material, I thank Devin Curry, William Lycan, audiences at the University of Tartu and at the 2024 online meeting of the Australasian Association of Philosophy.

- _____. (1980). Eliminativism and the propositional attitudes. *Journal of Philosophy* 78: 67–90.
- Churchland, P.S. (1986). *Neurophilosophy: Toward a Unified Science of the Mind-Brain*. Cambridge, MA: MIT Press.
- Curry, D. (2021). Interpretivism without judgment-dependence. *Philosophia* 49: 611–615.
- Davidson, D. (1973). Radical interpretation. *Dialectica*, 27: 314–328.
- _____. (1974). Belief and the basis of meaning. *Synthese* 27: 309–323.
- _____. (1983). A coherence theory of truth and knowledge. In D. Henrich (ed.), *Kant oder Hegel?* (pp. 423–433). Stuttgart: Klett-Cotta.
- Demeter, T. (2022). A mental fictionalism worthy of the name. In T. Demeter, T. Parent, & A. Toon (eds.), *Mental Fictionalism: Philosophical Explorations* (pp. 86–103). New York: Routledge.
- Dennett, D. (1975). True believers: The intentional strategy and why it works. In A.F. Heath (ed.), *Scientific Explanation: Papers based on the Herbert Spencer Lectures given at the University of Oxford* (pp. 53–75). Oxford: Clarendon Press. Pagination is from the reprint in Dennett (1987).
- _____. (1978). *Brainstorms: Philosophical Essays on Mind and Psychology*. Cambridge, MA: MIT Press.
- _____. (1981). Three kinds of intentional psychology. In R.A. Healy (ed.), *Reduction, Time, and Reality* (pp. 37–61). Cambridge: Cambridge University Press. Pagination is from the reprint in Dennett (1987).
- _____. (1987). *The Intentional Stance*. Cambridge, MA: MIT Press.
- _____. (1991). Real patterns. *Journal of Philosophy* 88: 21–51..
- _____. (1992). The self as a center of narrative gravity. In F. Kessel, P. Cole, & D. Johnson, (eds.), *Self and Consciousness: Multiple Perspectives* (pp. 103–115). Hillsdale, NJ: Erlbaum.
- _____. (1993). Back from the drawing board. In B. Dahlbom (ed.), *Dennett and his Critics: Demystifying the Mind* (pp. 203–235). Malden, MA: Blackwell.
- _____. (2017). *From Bacteria to Bach and Back: The Evolution of Minds*. New York: W.W. Norton.
- _____. (2018). Reflections on Rebecca Kukla. In B. Huebner (ed.), *The Philosophy of Daniel Dennett* (pp. 32–35). New York: Oxford University Press.
- _____. (2022). Am I a fictionalist? In T. Demeter, T. Parent, & A. Toon (eds.), *Mental Fictionalism: Philosophical Explorations* (pp. 352–361). New York: Routledge.
- Dreier, J. (2004). Meta-ethics and the problem of creeping minimalism. *Philosophical Perspectives* 18: 23–44.
- Field, H. (1977). Logic, meaning, and conceptual role. *Journal of Philosophy* 74: 379–409.
- _____. (1994a). Deflationist views of meaning and content. *Mind* 103: 249–285. Reprinted with a new postscript in his (2001).
- _____. (1994b). Disquotational truth and factually defective discourse. *Philosophical Review* 103: 405–452. Reprinted in his (2001).
- _____. (2001). *Truth and the Absence of Fact*. Oxford: Oxford University Press.
- Fodor, J. (1975). *The Language of Thought*. New York: Thomas Y. Crowell Co.
- _____. (2008). *LOT 2: The Language of Thought Revisited*. Oxford: Clarendon Press.
- Goodman, N. (1955). *Fact, Fiction, and Forecast*. Cambridge, MA: Harvard University Press.
- Haugeland, J. (1993). Pattern and being. In B. Dahlbom (ed.), *Dennett and his Critics: Demystifying the Mind* (pp. 53–69). Malden, MA: Blackwell.
- Hutto, D. (2013). Fictionalism about folk psychology. *The Monist* 96: 582–604.

- Joyce, R. (2013). Psychological fictionalism and the threat of fictionalist suicide. *The Monist* 96: 517-538.
- Kriegel, U. (2010). Interpretation: Its scope and limits. In A. Hazlett (ed.), *New Waves in Metaphysics* (pp. 111-135). London: Palgrave-McMillan.
- Kripke, S. (1973/2013). *Reference and Existence: The John Locke Lectures*. Oxford: Oxford University Press.
- Kukla, R. (2018). Embodied stances: Realism with literalism. In B. Huebner (ed.), *The Philosophy of Daniel Dennett* (pp. 3-31). New York: Oxford University Press
- Lewis, D. (1969). *Convention: A Philosophical Study*. Oxford: Blackwell.
- _____. (1974). Radical interpretation. *Synthese* 23: 331-344. Pagination is from the reprint with additional postscripts in Lewis (1983).
- _____. (1978). Truth in fiction. *American Philosophical Quarterly* 15: 37-46.
- _____. (1983). *Philosophical Papers Vol. I*. New York: Oxford University Press.
- Liggins, D (2008). Modal fictionalism and possible-worlds discourse. *Philosophical Studies* 138: 151-160.
- Nolan, D. (1997). Three problems for ‘strong’ modal fictionalism. *Philosophical Studies* 87: 259-275.
- Noonan, H. (1994). In defence of the letter of fictionalism. *Analysis* 54: 133-139.
- Parent, T. (2013). In the mental fiction, mental fictionalism is fictitious. *The Monist* 96: 605-621. Revised and expanded as Chapter 9 of Parent (2017).
- _____. (2015). Rule following and metaontology. *Journal of Philosophy* 112: 247-265.
- _____. (2017). *Self-Reflection for the Opaque Mind: An Essay in Neo-Sellarsian Philosophy*. New York: Routledge.
- _____. (forthcoming). Ontology after folk psychology; or, why eliminativists should be mental fictionalists. *Analytic Philosophy*.
- Parent, T., Toon, A. & Demeter, T. (in press). *Mental Fictionalism: Elements in Philosophy of Mind*. Cambridge: Cambridge University Press.
- Quine, W.V.O. (1960). *Word and Object*. Cambridge, MA: MIT Press.
- _____. (1968). Ontological relativity. *Journal of Philosophy* 65: 185-212.
- _____. (1973). *The Roots of Reference*. La Salle, IL: Open Court.
- Rosen, G. (1993). A problem for fictionalism about possible worlds. *Analysis* 53: 71-81.
- _____. (1995). Modal fictionalism fixed. *Analysis* 55: 67-73.
- Sainsbury, R.M (2010). *Fiction and Fictionalism*. London: Routledge.
- Schaffer, J. (2009). On what grounds what. In D. Chalmers, D. Manley, & R. Wasserman. *Metametaphysics: New Essays on the Foundations of Ontology* (pp. 347-383). Oxford: Clarendon Press.
- Stich, S. (1981). Dennett on intentional systems. *Philosophical Topics* 12: 39-62.
- _____. (1983). *From Folk Psychology to Cognitive Science: The Case Against Belief*. Cambridge, MA: MIT Press.
- Wallace, M. (2016). Saving mental fictionalism from cognitive collapse. *Res Philosophica* 93: 405-424.
- _____. (2022). Mental fictionalism. In T. Demeter, T. Parent, & A. Toon (eds.), *Mental Fictionalism: Philosophical Explorations* (pp. 27-51). New York: Routledge. An unpublished version with minor differences was first written in 2007.
- Woodbridge, J. (2005). Truth as a pretense. In M. Kalderon (ed.), *Fictionalism in Metaphysics* (pp. 134-177). Oxford: Oxford University Press.