

# Phenomenal Intentionality and a Voluntaristic Theory of Truth and Reference

Angela Mendelovici

Draft of August 15, 2020

## **Abstract**

According to the *phenomenal intentionality theory* (PIT), there is a kind of intentionality that is nothing over and above phenomenal consciousness (*phenomenal intentionality*) and all other kinds of intentionality derive from it. While PIT has been gaining in popularity in recent years, the view faces important challenges in accounting for how intentional states can be true or false or referring or non-referring. The aim of this paper is to precisify this challenge for PIT, outline some possible ways of meeting it, and provide a sketch of my favored approach, which is to supplement PIT with what I will call a “voluntaristic” theory of truth and reference, one on which the referents, truth-makers, or conditions of truth and reference of our contents are specified by us.

Keywords: phenomenal intentionality, consciousness, intentionality, truth, reference, non-relational views of intentionality, acquaintance, adverbialism

# 1 Introduction

According to the *phenomenal intentionality theory* (PIT), there is a kind of intentionality that is nothing over and above phenomenal consciousness (*phenomenal intentionality*) and all other kinds of intentionality derive from it. For example, according to PIT, a visual experience of a red square might represent the content <red square> simply in virtue of having a reddish-squarish phenomenal character. This experience of a red square is example of phenomenal intentionality. PIT also allows for non-phenomenal intentionality, so long as it is derived from phenomenal intentionality. For example, a standing belief about a red square might represent the content <red square> thanks to its relation to the phenomenal contents of experiences of red squares.<sup>1</sup>

This paper explores a challenge facing PIT that has been largely overlooked: that of accounting for truth and reference. On some conceptions of intentionality, intentional states are by definition identical to or importantly related to conditions of truth and reference. But even on conceptions of intentionality that do not tie it to truth and reference,<sup>2</sup> it is plausible that intentional states can at least sometimes be evaluated for truth or reference, and it is plausible that this is somehow connected to their being intentional. The problem is that it is hard to see how phenomenal states—the key ingredient required for intentionality, according to PIT—can secure any connection to the world beyond themselves of the sort required for truth and reference. Phenomenal states are just feelings, one might argue, and mere feelings cannot be true or false or refer to anything beyond themselves.

---

<sup>1</sup>See, e.g., Strawson 1994, Siewert 1998, Horgan and Tienson 2002, Pitt 2004, Loar 2003, Kriegel 2011, Pautz 2013, Farkas 2008, Bourget 2010a, Montague 2016, Woodward 2016, Mendelovici 2018a, Mendelovici and Bourget 2020 for versions of the phenomenal intentionality theory. Brentano 1874 is a notable historical precedent. See Kriegel 2013 and Bourget and Mendelovici 2016 for an overview.

<sup>2</sup>See, e.g., Kriegel 2011, ch. 1 and Mendelovici 2018a, ch. 1.

The aim of this paper is to precisify this challenge for PIT, outline some possible ways of meeting it, and provide a sketch of my favored approach, which is to supplement PIT with what I will call a “voluntaristic” theory of truth and reference, one on which the referents, truth-makers, or conditions of truth and reference of our contents are specified by us. I proceed as follows: §2 provides some background, §3 outlines the challenge for PIT from truth and reference, §4 considers some ways of responding to the challenge that I will ultimately reject, §5 provides a sketch of the voluntaristic theory of truth and reference, and §6 outlines a specific version of the voluntaristic theory.

## 2 Background

In order to get clear on PIT’s claims, we must first get clear on its explanatory target: intentionality. Intentionality is sometimes characterized as the “aboutness” or “directedness” of mental states. This characterization arguably succeeds gesturing towards the phenomenon of interest, but it is too vague and metaphorical to fix clearly on our target. A better way of fixing reference on our target is ostensively. Intentionality is arguably a phenomenon that we, at least sometimes and at least to some degree, notice introspectively in ourselves in certain paradigm cases. For example, a visual experience of a blue cup seems to be “of”, “about”, or “directed at” a blue cup, an auditory experience of sounds and words seems to “present” sounds and words, and a thought that grass is green seems to “say” that grass is green. We can thus take *intentionality* to be that feature of these and other paradigm cases that we are tempted to describe using representational terms like “of”, “present”, “say”, “about”, and “directed at”. We can say that what an intentional state is “of,” “about,” or *represents* is

its *content*.<sup>3</sup>

PIT aims to offer a theory of intentionality, a specification of the nature of this phenomenon that we notice in paradigm cases. The main ingredient in PIT's account is *phenomenal consciousness*, the felt, qualitative, experiential, or subjective aspect of mental life. This is the “what it's like” (Nagel 1974) of being in certain mental states, the feature of mental life for which there is a “hard problem” of consciousness (Chalmers 1995). According to PIT, there is a kind of intentionality—*phenomenal intentionality*—that is nothing over and above phenomenal consciousness. There are various ways in which phenomenal intentionality might be nothing over and above phenomenal consciousness: it might be identical to, realized by, or grounded in phenomenal consciousness. But on all versions of the view, having a phenomenal state (perhaps of a certain kind) is metaphysically sufficient for representing a content.<sup>4</sup> Let us call states of phenomenal intentionality *phenomenal intentional states* and the contents of such states *phenomenal contents*. We can say that phenomenal contents are contents that are *phenomenally represented*.

According to PIT, phenomenal intentionality is the most basic, central kind of intentionality. It is “original” intentionality, intentionality from which all other kinds of intentionality (if there are any) derive. In short, then, the phenomenal intentionality theory (PIT) is the view that there is a kind of intentionality that is nothing over and above phenomenal consciousness (phenomenal intentionality) and that all other kinds of intentionality derive from it. Equivalently, PIT is the view that all original intentionality is phenomenal intentionality.

Some phenomenal intentionalists maintain that *all* intentionality is phenomenal intentionality (see, e.g., Pitt 2004, Farkas 2008, and Mendelovici 2018a).

---

<sup>3</sup>This way of fixing reference on intentionality is defended in *[redacted]*.

<sup>4</sup>PIT is compatible with the claim that not all phenomenal states are metaphysically sufficient for intentional states. Such states would be “mere qualia”, phenomenal states that float free of representing a content.

However, many accept that there is derived intentionality, which is ultimately derived from phenomenal intentionality (see, e.g., Horgan and Tienson 2002, Bourget 2010a, Kriegel 2011, and Pautz 2013). For example, Horgan and Tienson (2002) suggest that wide contents are derived from phenomenal contents together with grounding presuppositions and facts about the world, Kriegel (2011) argues that nonconscious content is derived from the phenomenal content of an ideal rational agent, and Bourget (2010a) suggests that non-phenomenal intentional states derive their contents through various mechanisms, such as composition, functional relations, descriptive reference, and deference. Mendelovici 2018a argues that occurrent thoughts and standing states derivatively represent various contents in virtue of dispositions to phenomenally ascribe these contents to our states or ourselves, though she claims that this kind of derived representation does not, strictly speaking, qualify as a kind of intentionality. As this diversity of views makes clear, there are various ways in which PIT might try to account for intentionality beyond the conscious realm.

Another choice point for PIT concerns the internal structure of intentionality. On the *relation view*, intentional states are relations to distinctly existing contents, which might be, for example, sets of possible worlds, structured propositions consisting of concrete objects and abstract properties, or abstract property clusters.<sup>5</sup> According to the alternative *aspect view*, intentional contents are aspects of intentional states, which which might be properties of intentional states, parts of such properties, or intentional states themselves<sup>6,7</sup>

---

<sup>5</sup>See Pautz 2007, Bourget 2010b, 2019a for relational versions of PIT.

<sup>6</sup>See Pitt 2009, Kriegel 2011, Mendelovici 2018a, Gow MS for aspect-view-theoretic versions of PIT.

<sup>7</sup>See Kriegel 2013, Mendelovici and Bourget 2014, 2020, and Bourget and Mendelovici 2016 for an overview of PIT.

### 3 The challenge from truth and reference

Intuitively, it seems that the contents our mental states represent can be true or false, accurate or inaccurate, satisfied or unsatisfied, veridical or nonveridical, or referring or non-referring. For simplicity, we can divide this family of features into two types: *truth* and *reference*. Let us say that a content is *true* when it “corresponds to” or “agrees with” the world, perhaps by being related to some fact or other entity that makes it true (its *truth-maker*). A content’s *truth conditions* are the conditions in which it would be true. For example, the content <grass is green> is true in conditions in which grass is in fact green. If such conditions obtain, then the fact that grass is green is the content’s truth-maker. Note that on this broad notion of truth, desires that are satisfied and experiences that are accurate or veridical qualify as instances of truth. We can think of truth as a kind of seal of approval, a positive epistemic status, that a propositional content gets to have when it is related to the world in the right way.

We can say that a content *refers* to something—its *referent*—when it “picks out” or “selects” that thing as its target. A content’s *conditions of reference* are the conditions in which it refers to particular items. We can think of conditions of reference as sets of ordered pairs of conditions and items such that if the conditions in a pair obtain, the content refers to the item in the pair. For example, the content <my dog> refers to Lily in conditions in which Lily is my only dog, Mango in conditions in which Mango is my only dog, Barkley in conditions in which Barkley is my only dog, etc. Since [redacted] is my only dog, <my dog> in fact refers to [redacted]. As we are using the terms here, reference is any semantic feature that is a matter of “picking out” or “selecting” something

in the world.<sup>8,9</sup>

PIT faces a challenge in accommodating truth and reference. According to PIT, all intentionality is nothing over and above phenomenal consciousness or derived from such intentionality. The problem is that it is not clear how phenomenal states can relate us to or be made true by things in the world. How can there be a fact of the matter as to which particular worldly thing a phenomenal state picks out or which state of the world it is supposed to go with?<sup>10</sup>

It is worth distinguishing between different kinds of truth and reference that we might want to accommodate. *Internal-world* truth and reference is a matter of referring to or according with our conscious states and their features—i.e., with the *internal world*. *External-world* truth and reference is a matter of referring to or according with things beyond our conscious states—i.e., with the *external world*. While PIT owes us a story of both kinds of truth and reference, it faces a special challenge in accommodating external-world truth and reference. This is because it is not implausible that we are in some sense *acquainted* with (at least some of) internal-world items in the more or less Russellian sense of “acquaintance” (Russell 1997): they are immediately given to us, “before our mind’s eye,” and non-inferentially knowable by us. This arguably allows us to

---

<sup>8</sup>Unlike in the case of intentionality, we arguably don’t have a clear, ostensive grip on truth and reference. We might represent contents that are true or refer, but, at least when their truth-makers and referents are not our own mental states, we have no access to their referents, truth-makers, and our relations to these referents and truth-makers except via our access to our contents. We only have direct access to the mind side of mind-world relations.

<sup>9</sup>Although I will speak of contents being true or false or referring, we can also say that mental states, mental representations, or other items that have contents are true or false or refer to something just in case their contents do.

<sup>10</sup>This kind of objection to PIT has been forcefully put by Walter Ott (2016), who argues that many versions of PIT do not provide a satisfactory account of intentionality because they do not account for what he calls “representation,” which requires a connection to the mind-independent world. As I’ve defined “intentionality,” intentionality does not *by definition* involve such a connection to the world, so PIT does not automatically fail if it does not provide an account of what Ott calls “representation.” However, insofar as there is good reason to think that intentional states secure such a connection to the world, we have an objection to PIT.

pick them out by acquaintance, that is, to pick them out by being acquainted with them, in much the same way that Russell claimed we can refer to things by acquaintance. Much more needs to be said about how we can pick things out by acquaintance, but it is not entirely mysterious that this might be possible. Once something is before your mind's eye, it is available to you to single it out as a target of your thoughts or other intentional states, perhaps by selectively attending to it, inwardly demonstrating it, or embedding it in a thought or other intentional state in the right way. For example, suppose you are acquainted with a pain P. On one picture, you might refer to it by entertaining a content <the pain P>, which includes P itself as a constituent.<sup>11</sup>

In contrast, we arguably are not acquainted with external-world items—they are not kicking around in our minds, immediately available to form the targets of further thoughts and concepts—and so we arguably cannot pick them out by acquaintance. What is required for external-world truth and reference is a link between mental states and the extra-mental world. The problem is that it's hard to see how there can be such a link between phenomenal or phenomenally-determined states and external-world items. For example, according to PIT, my concept of Justin Trudeau might have an associated phenomenal character, which might constitute an associated phenomenal content. But it is not clear how this phenomenal character and any phenomenal or derived content it gives rise to is supposed to pick out the actual, flesh-and-blood Justin Trudeau. Any connection between the Justin-Trudeau-ish phenomenal character and Justin Trudeau himself seems entirely arbitrary.

We can appreciate the challenge for PIT from external-world truth and

---

<sup>11</sup> Arguably, picking something out by acquaintance requires more than just acquaintance. There are presumably many items that we are acquainted with, including many of our own mental states, but we do not automatically refer or otherwise pick out all of them. For example, we might be acquainted with both a pain and an itch but only single out the pain. Intuitively, we have referred to the pain but not the itch.

reference by comparing the case of PIT to that of an alternative family of theories of intentionality, *tracking theories*, on which intentionality is a causal, informational, or other kind of tracking relation to things in the environment.<sup>12</sup> Tracking theories are often thought to offer what is sometimes called a *referential theory of content*, a theory of content on which at least some contents are identical to their referents. For example, according to a tracking theory, the concept CAT might bear a content-endowing tracking relation to the property of being a cat, which is both its content and that content’s referent. On such a picture, a connection to the external world is established at the level of content, not merely at the level of reference. If the story succeeds as a story about content, then there is no further question remains as to how it can accommodate (at least some instances of) external-world reference.<sup>13</sup> In contrast, it is not clear that PIT can offer such a referential theory of content (though see §4.2). It seems that PIT must forge the connection between mind and world from scratch at the level of truth and reference, and it is not clear how it might do so.

While much more remains to be said about internal-world truth and reference (and I will say a bit more in §5), it is clear enough for now that the external-world case poses a bigger challenge for PIT. In what follows, I will focus on this part of the challenge. I first consider three strategies for addressing it that I will reject. I then turn to my favored strategy, that of offering what I will call a “voluntaristic” theory of truth and reference.

---

<sup>12</sup>See Fodor 1987, Dretske 1995, and Neander 2017 for representative tracking theories.

<sup>13</sup>Unfortunately, I think the story does not succeed as a story of content (*[redacted]*).

## 4 Some failed attempts

### 4.1 Eliminativism about external-world truth and reference

One way of dealing with the challenge of accommodating external-world truth and reference is to simply deny that there is such thing as external-world truth and reference, at least in any full-blooded, substantive sense. According to *eliminativism* about external-world truth and reference, contents do not have truth-values, referents, or conditions of truth and reference. Eliminativism is nonetheless compatible with the claims that there are intentional states, that the external world exists, and that the external world causally interacts with us and our intentional states in various ways. The view might be combined with a deflationary story of truth and reference, one that takes the interesting work to be done in the area of truth and reference to be that of explaining how truth and reference predicates work in a language.<sup>14</sup>

Eliminativism is not entirely out of the question. One might try to support the view by suggesting that the very idea of external-world truth and reference is defective or confused. The only cases of truth and reference that we can directly observe are cases of internal-world truth and reference. These are cases in which we can observe not only the contents that are true or refer but also the relevant truth-makers and referents themselves and perhaps even the way in which we are related to them. In these cases, truth and reference involves singling out some thing we are acquainted with. External-world truth and reference, if it is exists, would have to be an external-world version of this same phenomenon. But it is unclear just what it would be for this same phenomenon to occur in the external-world case. What would it be for us to single out something beyond

---

<sup>14</sup>For an overview of deflationary theories of truth, see Stoljar 2008.

our conscious minds? The very idea, one might suggest, is defective or confused.

Another reason to doubt the existence of external-world truth and reference is that we have no clear evidence for its existence. There is no phenomenological evidence for relations of external-world reference and truth-making: while representing a content can make a phenomenological difference for us, the content's being true or referring to something makes no further difference—there is no phenomenological mark of successful reference or truth. There is also no broadly psychological evidence for such relations: while representing a content can make a difference to our further mental states and behaviors, the content's being true or referring makes no further contribution to the content's psychological role.<sup>15</sup> We likewise have no phenomenological or psychological evidence for the existence of *conditions* of truth and reference: that our mental contents would be true or refer to some thing if certain worldly conditions were to obtain does not make a phenomenological or psychological difference for us. If all this is right, then we have no reason stemming from phenomenological considerations or considerations of psychological role to think that there is such thing as external-world truth and reference. Since it is hard to see what other empirical or theoretical considerations might support their existence, we might conclude that a belief in external-world truth and reference is unfounded.

If all this is right, it is difficult to rule out eliminativism, and there might even be some positive considerations in its favor. Of course, the end result is an unhappy view. According to eliminativism, we cannot refer to our friends and family, and our philosophical theories about the external world are neither true or false. But that a view has unwanted consequences does not provide reason to

---

<sup>15</sup>Of course, a content's being true or referring might involve the world being a certain way, which might make a causal difference to our subsequent phenomenology, mental states, and behaviors, but it is the world's being that way, not the relation of truth-making or reference that our content bears to the world, that makes this difference. Truth and reference remain phenomenologically and psychologically epiphenomenal.

think it is false.<sup>16</sup>

In the remainder of this paper, I will consider some ways of attempting to accommodate external-world truth and reference and eventually argue that there is such thing as external-world truth and reference. We will see that the notion of external-world truth and reference is not in fact defective or confused and that there are theoretical reasons to think that the conditions required for external-world truth and reference can indeed be fulfilled.

## 4.2 A relational view of intentionality

One might suggest that a relational version of PIT can readily account for truth and reference. Recall that according to the relation view, intentional states are relations to distinctly existing contents, which might be, for example, sets of possible worlds, structured propositions consisting of concrete objects and abstract properties, or abstract property clusters. According to the alternative aspect view, intentional contents are aspects of intentional states, which might be properties of intentional states, parts of such properties, or intentional states themselves.

Relational versions of PIT—versions of PIT that endorse the relation view—can readily take their distinctly existing contents to be referents, truth-makers, or conditions of truth and reference, perhaps even providing a referential theory of content of the sort described in §3.<sup>17</sup> If contents are identical to their referents or truth-makers, then intentionality itself relates us to referents and truth-makers. If contents are conditions of truth and reference, then contents themselves specify

---

<sup>16</sup>One might object that the view cannot even be truly thought, since that would require external-world truth. But an advocate of the view might be able to truly think it in a way that does not require external-world truth, perhaps by thinking that internal-world truth and reference is all the truth and reference there is. In any case, even if the view cannot be truly thought, the state of affairs of there not being external-world truth and reference might nonetheless obtain.

<sup>17</sup>The analog of a referential theory of content in the case of truth is an identity theory of truth, on which contents are identical to truth-makers.

which worldly entities must exist or obtain in order for them to be true or refer to something. Either way, a picture of truth and reference falls out of our theory of intentionality.<sup>18,19</sup>

Unfortunately, this overall picture faces difficulties both as an account of intentionality and as an account of truth and reference. As an account of intentionality, it faces well-known problems in accommodating all the contents we can manifestly represent. While it might seem natural to say that a perceptual experience of a cup involves a concrete cup and its concrete property instantiations as part of its content, the view does not readily extend to hallucinations, illusions, reliably misrepresenting experiences (Mendelovici 2013, 2016), and false thoughts. In some cases, the concrete world is simply not enough to capture all the contents we can manifestly represent. For example, suppose that the contents of color experiences are primitive colors that happen not to be instantiated. In order to accommodate the representation of such contents, the relation view would have to posit entities existing beyond the concrete world—e.g., Platonic properties—that play the role of contents. This is problematic not only because it leads to an inflated ontology but also because the ontological extremes the relation view leads to suggests that the entire approach is wrongheaded.<sup>20</sup>

A second problem with the relation view is that it is not clear how bearing a relation to something can make that thing entertained or otherwise represented by us. Our contents are *psychologically involved* in that they play various psychological roles in the cognitive economy, such as those of being related to

---

<sup>18</sup>If the relationalist takes us to be acquainted with some of the distinctly existing items playing the roles of contents, she can further say that we can pick them out by acquaintance in much the same way that I suggested earlier we can pick out our conscious states and their properties by acquaintance. Pautz's (2007) and Bourget's (2010b, 2019a, 2019b) views can be understood in this way.

<sup>19</sup>See especially Bourget 2019b for discussion of truth and reference on a relational version of PIT.

<sup>20</sup>Crane 2001, Thau 2002, Kriegel 2007, Schellenberg 2011, and Mendelovici 2018a, §9.2.1 raise different versions of this kind of worry.

further thoughts and behaviors, being the targets of higher-order thoughts and introspective states, and contributing to our overall phenomenology, including the phenomenology of “grasping” certain contents. It is unclear how entities existing distinctly from our intentional states—particularly entities existing beyond the concrete world—can play any such roles.<sup>21</sup>

Even if we waive these worries with the relation view, it is not clear that a relational version of PIT can readily and unproblematically accommodate truth and reference. As we saw, the relation view cannot take all contents to be referents or truth-makers, since not all contents refer or are true—consider nonveridical hallucinations, illusions, reliable misrepresentations, and false thoughts. So we can rule out the picture of truth and reference on which all contents are identified with truth-makers and referents.

Perhaps, though, the relation view can still say that at least some contents are *conditions* of truth and reference. For reasons already considered, these conditions of truth and reference will have to at least sometimes be made up out of non-concrete items, like Platonic properties. So truth and reference will have to involve relations between these non-concrete items and the things in the world that are our referents and truth-makers. The problem is that it is not clear what kinds of facts could make it the case that there is a right or wrong way for these items to relate to the rest of the world in the ways required for truth and reference. One might suggest that these items are Platonic properties or Platonic states of affairs and that the relations in question are simply the relations of instantiation or obtaining, respectively. The problem with this is that, while these relations might sound familiar, it is not so clear what these instantiation and obtaining relations would amount to, particularly on the Platonic picture of properties and states of affairs. We would still need to explain why certain

---

<sup>21</sup>Kriegel 2011, BonJour 1998, Papineau 2014, and Mendelovici 2018a, §9.2.2 raise this kind of objection. Pautz 2007, 2010 considers and rejects such worries.

Platonic entities “go together” with certain worldly items and not others. While there may be ways of fleshing out the picture, one might worry that the overall view ends up being quite mysterious.<sup>22</sup>

If all this is right, then there are reasons to reject a relational version of PIT. Even if we set them aside, there are reasons to think that simply adopting a relational version of PIT does not provide a satisfactory way of dealing with the challenge from truth and reference. In what follows, I will explore aspect-theoretic responses to the challenge from truth and reference. However, the view I will eventually recommend can be modified to fit the relation view.

### **4.3 Supplementing PIT with an independent theory of truth and reference**

Another way in which we might try to accommodate external-world truth and reference is by simply combining PIT with an independent theory of truth and reference, such as a causal, similarity, or primitivist theory. For example, we might say that an intentional state’s content refers to or is made true by whatever causes it (perhaps by a particular kind of causal relation), or that an intentional state’s content refers to or is made true by whatever it is similar to (perhaps by a particular kind of similarity relation). On such a view, truth and reference are independent of intentionality. Intentionality is a matter of phenomenal consciousness, while truth and reference are a matter of causation, similarity, or something else.

One worry with this kind of view is that it seems to be committed to brute and arbitrary semantic facts. For example, suppose we say that it is a particular

---

<sup>22</sup>This kind of argument is made in Mendelovici 2018a, §9.3.4, where I argue that the resulting relational view of truth and reference is no more attractive than an aspect theory that goes primitivist about truth and reference (which is pretty unattractive—see the next subsection for an argument against such a primitivist view).

causal relation, R, that links contents to their referents. The fact that it is R—and not, say, some other causal relation or an entirely different kind of relation—that is the reference relation is a brute and arbitrary fact about reference. There is nothing special about R, as far as causal relations go, that makes it to play this special semantic role of linking contents to their referents.<sup>23</sup>

There is a bigger problem, however, with the strategy, one that gets closer to the heart of the matter. The problem is that the resulting view is not compatible with “voluntarism,” which we might define as follows:

(Voluntarism) Our content’s referents, truth-makers, or conditions of truth and reference are in some important sense up to us.

*Involuntarism* is the denial of voluntarism. According to voluntarism, we have ultimate control of what we refer to in various circumstances or what would make our contents true. In some cases, we might pick out our contents’ referents and truth-makers by acquaintance, thereby specifying them directly. In other cases, we might only be able to specify only our contents’ *conditions* of truth and reference, *what it would take* for them to refer to some item or to be made true. In either case, though, we are able to specify a path to our referents and truth-makers. In this way, voluntarism respects the intuitive idea that which items we refer to and which states of affairs make our contents true is in some way up to us.<sup>24</sup>

In contrast, according to involuntarism, factors beyond our intentions, stipulations, etc. can play a role in determining our referents, truth-makers, and conditions of truth and reference. These factors might include the causal relations

---

<sup>23</sup>This is, in effect, Putnam’s (1983) argument against taking causal relations to be “external” constraints on reference. Putnam argues that such a view is committed to brute semantic facts.

<sup>24</sup>One might object that a voluntaristic theory of reference is circular in that it requires semantic factors, such as referential intentions, stipulations, endorsements, or the like. While I accept that a voluntaristic theory presupposes content, it need not presuppose reference, so it is not circular. Thanks to Uriah Kriegel for prompting this clarification.

our mental states enter into, similarity relations between mental states and other items, or even primitive semantic facts about the world.

Voluntarism receives some intuitive support. Hilary Putnam writes: “[T]he world does not pick models or interpret languages. *We* interpret our languages or nothing does.” (1980, p. 482, emphasis in original) Likewise, although David Lewis ultimately rejects voluntarism, he acknowledges its intuitive force: “Take your favourite theory of reference. Let us grant that it is true. But let us ask: what makes it true? And the tempting answer is: *we* make it true, by our referential intentions.” (1984, p. 226)<sup>25</sup>

One might argue that voluntarism follows from the very notion of truth and reference. Suppose we had an intentional state with a content that bears whatever special relation R we want to some worldly item independent of our intentions, stipulations, or other voluntaristic factors. No matter what is the nature of R, bearing R to some worldly item would not by itself suffice for picking out that item as our referent or truth-maker. This is clearest if R connects us to something that is incongruent with voluntaristic factors. For example, suppose R links the content <goldfish> to a ham sandwich. There is no non-voluntaristic way of filling in the nature of R such that <goldfish> refers to the ham sandwich. This suggests that even in cases where R connects us to something that is in line with voluntaristic factors—e.g., in a case where it links <goldfish> to the property of being a goldfish—it is the voluntaristic factors that do the work of fixing upon our referent or truth-maker, not R. If all this is right, non-voluntaristically-determined relations would not count as relations of reference or truth-making. They would not satisfy our notions of truth-making

---

<sup>25</sup>Putnam and Lewis both discuss voluntarism in the context of Putnam’s model-theoretic argument against metaphysical realism, a view that Putnam takes to include a commitment to a correspondence theory of truth and reference. Although their discussion centers on the case of language, similar points can be made in the case of thought (cf. Lewis 1984, p. 222). See also n. 31.

and reference.<sup>26</sup>

The strategy for accommodating truth and reference under consideration combines PIT with an independent theory of truth and reference in terms of causality, similarity, primitive relations, or the like. The main problem with this strategy is that the resulting view is involuntaristic: what the relevant causal, similarity, or other relations end up relating us to is not up to us.

## 5 A voluntaristic theory of truth and reference

At the end of the previous section, I argued for voluntarism, the view that our referents, truth-makers, or conditions of truth and reference are up to us, perhaps because they are stipulated, accepted, or otherwise endorsed by us. This section sketches my positive proposal, which is a version of voluntarism that is compatible with PIT.

Let us begin by considering some ways in which we might pick out referents, truth-makers, and conditions of truth and reference that are compatible with voluntarism:

**By acquaintance.** We have already encountered one voluntaristic way of picking out referents, truth-makers, and conditions of truth and reference: by acquaintance. This way is voluntaristic because it involves us ourselves selecting a referent or truth-maker, perhaps by inwardly demonstrating it, specifically attending to it, or embedding it in a higher-order thought.

---

<sup>26</sup>Another way to evoke voluntaristic intuitions is to consider what we care about when we care about truth and reference. I want to suggest that we do not care what R—or any non-voluntaristic factors—link us to. What we care about is whether the conditions of truth and reference that we specify obtain. Suppose you want to know if God exists. It is no comfort to you to be told that <God> is related by R to something and so that therefore God exists. What you want to know is whether <God> is related to something that accords with your referential intentions, a thing the path to which you have specified—you want to know if there's something at the end of that path.

**By description.** Another voluntaristic way of picking out referents, truth-makers, and conditions of truth and reference is by constructing a description that is satisfied by them. For example, if we think a thought involving the descriptive content <the closest wallaby to me right now>, we voluntaristically refer to our referent (even if we don't happen to know which particular wallaby satisfies the description).

**By criteria.** A third voluntaristic way of picking out referents, truth-makers, and conditions of truth and reference is via criteria of truth and reference that we endorse. For example, we might endorse the following general criterion of perceptual reference:

(Causal-reference) A perceptual content C refers to O iff O causes C (by causal relation R).

If a particular perceptual state is appropriately caused by, say, an apple, then its content voluntaristically refers to that apple.

Recall that there two types of truth and reference we want to accommodate: internal-world and external-world truth and reference. In §3, I suggested that internal-world truth and reference proceeds by acquaintance: we are acquainted with internal-world referents and truth-makers, which enables us to pick them out, either by inwardly demonstrating them, selectively attending to them, or embedding them in a thought or other intentional state. If so, our story of internal-world truth and reference is voluntaristic.

Although my focus in this paper is not on internal-world truth and reference, it is worth making a few further remarks on how we might flesh out this acquaintance-based picture. We can distinguish between reductive and primitivist versions of the view: *reductive* versions take picking out referents and truth-makers by acquaintance to be ultimately grounded in, constituted by, or otherwise

nothing over and above something else, while *primitivist* versions take picking out referents and truth-makers by acquaintance to be primitive.

At least some versions of the *embedding view*—that truth and reference by acquaintance is a matter of embedding some item in a thought or other intentional state—are reductive. Embedding an item, say a pain P, in an intentional state arguably involves forming a complex mental content that includes P and various supplementary contents as constituents, e.g., the complex content <the pain P>. This picture reduces truth and reference by acquaintance to the having of complex intentional states involving internal items as constituents.

Another reductive view is the *selective attention* view, which reduces truth and reference by acquaintance to selective attention. On this view, picking some item out by acquaintance is a matter of attending to it, or perhaps to the exclusion of other items, and perhaps in a certain special way. Attention (or the relevant kind of attention) might be taken to be primitive, or it might be reduced to something else, as on representationalist theories of attention.<sup>27</sup> Either way, such a view does not require a primitive form of truth or reference in order to account for truth and reference by acquaintance, so it is reductive.

In contrast, a view that takes truth and reference by acquaintance to be a matter of inner demonstration is arguably primitivist, since inner demonstration is just a kind of reference, one that cannot be elucidated in terms of something else.

My preference is for a combination of the embedding and the selective attention views. On this view, truth and reference by acquaintance can be achieved both by embedding and by selective attention.<sup>28</sup> However, nothing

---

<sup>27</sup>See Chalmers 2004, Nanay 2010, Koralus 2014, and Bourget 2017). I am sympathetic to such views of attention.

<sup>28</sup>My reason for preferring this view is that both selectively attending to something and embedding it in the right kind of intentional state seem to be sufficient for picking it out in a way that answers to the intuitive notion of reference.

that follows hangs on this particular view. What matters for present purposes is that an account of internal-world truth and reference in terms of acquaintance is voluntaristic.

Let us now turn to external-world truth and reference. We are not acquainted with external-world referents and truth-makers, so if we can pick them out voluntaristically, it must be either by description or by criteria. It is not difficult to see how various descriptions and criteria might be formed that would pick out various external-world items. There is, however, a complication: In order for our descriptions and criteria to play the requisite roles in picking out referents and truth-makers, they themselves must already refer. For example, suppose we endorse (Causal-reference). In order for this criterion to have any force, it must itself already succeed in referring to the conditions it requires for perceptual reference. Assuming a modicum of compositionality,<sup>29</sup> this would require that its component contents antecedently successfully refer, e.g., that <perceptual state>, <caused>, and other non-logical components of the criterion already refer. These component contents might refer by acquaintance, by description, or by criteria, but if they refer by description or by criteria, the problem resurfaces: the descriptions or criteria by which they refer must *themselves* antecedently refer. Ultimately, reference has to be secured by a means that does not require further reference—that is, by acquaintance.<sup>30,31</sup>

---

<sup>29</sup>How PIT can accommodate compositionally structured contents is an important question. See Mendelovici 2018a, §9.3.3 and Mendelovici forthcomingb for discussion.

<sup>30</sup>Might there be some other voluntaristic way of referring that does not require other forms of reference? I don't think so, but if there is, then we could appeal to it here.

<sup>31</sup>This is what I take to be the moral of Putnam's (1977) model-theoretic argument against metaphysical realism. Roughly, Putnam argues that any ideal theory can be interpreted (i.e., assigned referents and truth-values for its terms and sentences) in "deviant" ways such that it comes out true and that there is no way internal to the theory to specify the "intended" interpretation. Any stipulation as to how a theory is supposed to be interpreted is "just more theory"—it is just another sentence to be added to the theory, which itself can be deviantly interpreted.

(Lewis' (1984) response to Putnam is that facts about natural kinds impose external constraints on truth and reference, i.e., that an involuntaristic theory of truth and reference is true. These constraints are not "just more theory" because they are supposed to be supplied by the world itself, not by our theory. But Putnam (1983) takes such a theory to be magical,

So, then, to put things starkly, all truth and reference must ultimately be grounded in acquaintance. It is pertinent, then, to consider what exactly we are acquainted with, for these items will form our “vocabulary” of contents out of which we can build our criteria and descriptions.

Earlier, we said that we are acquainted with conscious states and (at least some of) their properties. This might seem like a highly constrained vocabulary, but it arguably includes quite a lot, such as the following:

- Token conscious states
- The property of being a conscious state
- The property of being a perceptual conscious state
- The relation of similarity in a respect (which we are arguably acquainted with in having conscious states that are similar in certain respects)
- The relation of identity
- The relation of (perhaps merely Humean) causation (which we are arguably acquainted with in having conscious states that cause further states)

We might assume that we get our logical vocabulary for free, but if we do not, we might be able to build some of it up out of logical operations we are acquainted with:

- Predication (which we might be acquainted with in being acquainted with conscious states being certain ways)
- Conjunction (which we might be acquainted with in being acquainted with multiple mental states being certain ways at the same time)

---

requiring brute and arbitrary semantic facts—see n. 25).

- Negation (which we might be acquainted with in being acquainted with failures to have certain states, e.g., in cases where we try to bring a particular content to mind but fail to)
- Existence (which we are acquainted with in having any conscious state)
- Totality (which we are acquainted with in being acquainted with our total conscious state)
- Identity (which we are acquainted with in having a conscious state that is identical to itself—i.e., any conscious state)

This logical vocabulary allows us to define all truth-functional connectives, quantifiers, and perhaps even modal operators (analysed in terms of possible worlds).

Our vocabulary has fairly strong expressive power, allowing us to construct descriptive contents such as the following, where S is a token conscious state that we are acquainted with:

- <the non-mental thing that is most similar to S>
- <the mental thing that is most similar but not identical to S>
- <the kind of thing that causes token conscious states of the same type as S>

As these examples illustrate, we can form descriptions that invoke similarity relations and causal relations and that pick out mental and non-mental things. Note also that the fact that we are acquainted with particular conscious states allows us to form descriptions that are anchored to our local environments, allowing us to pick out referents and truth-makers in our local environments rather than qualitatively similar referents and truth-makers in, say, far-off parts of the world.

Our vocabulary also allows us to form criteria such as these:

(Causal-reference) A perceptual content  $C$  refers to  $O$  iff  $O$  causes  $C$  (by causal relation  $R$ ).

(Similarity-reference) A content  $C$  refers to  $O$  iff  $O$  is similar to  $C$  (in respects  $R$ ).

(Similarity-truth) A propositional content  $C$  is true iff there is a fact  $P$  and  $P$  is similar to  $C$  (in respect  $R$ ).

Of course, the relevant kind of causal relation in (Causal-reference) and the relevant respects in (Similarity-reference) and (Similarity-truth) would have to be specified in terms of items we are acquainted with. The point now, however, is to illustrate the power of our vocabulary. Importantly, we can see how we have the resources to formulate both causal and similarity criteria, which are plausible candidates for criteria of truth and reference that we actually endorse.

We might also be able to formulate criteria of truth specifying how the truth-values of logically-complex contents depend on the truth-values of their parts. As I will later suggest, we might in fact endorse such criteria:

(Truth-predication) A predicative thought content  $Fa$  is true iff  $a$  refers to something that instantiates the property that  $F$  refers to.

(Truth-conjunction) A conjunctive thought content  $A\&B$  is true iff  $A$  is true and  $B$  is true.

(Truth-negation) A negated thought content  $\sim A$  is true iff  $A$  is not true.

Since (Truth-predication) invokes the notion of reference, it would have to be combined with a further criterion of reference for contents with an objectual or proprietorial form,<sup>32</sup> perhaps one appealing to similarity or causal relations along

---

<sup>32</sup>Such contents needn't actually involve objects or properties as parts—see fn. 38.

the lines mentioned above. Likewise, (Truth-conjunction) and (Truth-negation) would have to be combined with further criteria of truth, like (Truth-predication). Once all this is in place, though, we have all we need in order to provide Tarski-style (1943) recursive criteria of truth for contents with a first-order logical structure without quantification since we can define the other truth-functional predicates out of conjunction and negation.

In order for all this to be compatible with PIT, of course, we need a PIT-friendly account of where the relevant descriptions and criteria live. In most cases, we do not phenomenally represent such contents. Elsewhere, I have argued that both descriptive contents associated with thoughts, concepts, and other representations, and the contents of standing states like beliefs and desires are derivatively represented thanks to our dispositions to self-ascribe them (Mendelovici 2018a, chs. 7–8, forthcominga).<sup>33</sup> This provides a theoretical home for our descriptions and, assuming they are represented by standing states, the criteria we endorse: we derivatively represent them.

This section recommended a voluntaristic theory of truth and reference, one on which internal-world truth and reference is secured by acquaintance and external-world truth and reference is secured by descriptions or criteria that ultimately obtain their reference from acquaintance. The proposal is sketchy, but it suffices to show how truth and reference might be possible for creatures like us. The next section offers some suggestions as to the specific descriptions and criteria we might endorse.

---

<sup>33</sup>Incidentally, this account of derived representation is voluntaristic in an extended sense of the term in that what contents we derivatively represent is wholly up to us.

## 6 Matching criteria in perception and recursive criteria in thought

The previous section, if successful, shows that it is at least possible to specify descriptions and criteria that determine conditions of external-world truth and reference. It is a further question, however, which specific descriptions or criteria we endorse. In this section, I suggest that we endorse what I will call “matching” criteria of perceptual truth and reference and a combination of descriptions and recursive criteria for truth and reference in thought.

The criteria that I will suggest we endorse in the case of perception invoke the notion of “matching,” which I take to be a kind of similarity between mental contents and things in the world. In order to get a grip on the relevant kind of similarity, it is helpful to draw a distinction between a content’s *deep nature*, which is what it *is*, deep down, metaphysically speaking, and its *superficial character*, which is the set of features that characterize it as the particular content that it is.<sup>34</sup> We can tell independently of any theory of intentionality whether an intentional state represents, say, <grass is green> rather than <there is an octopus on my arm>. The features of our contents that allow us to make such judgements are their superficial characters. In contrast, we arguably do not have the same kind of access to facts concerning our contents’ deep natures, e.g., whether they are structured propositions, sets of possible worlds, or adverbial modifications of subjects. With this understanding of superficial character in mind, we can say that an intentional content *C* *matches* an item *O* when all the features of *C*’s superficial character are instantiated by *O*.

As a first pass, the matching criterion can be specified as follows:

(Matching-reference) Perceptual content *C* refers to external-world item *O* iff *C*

---

<sup>34</sup>See Mendelovici 2018a, §2.2.2 and 2018b.

matches O.

(Matching-truth) Propositional perceptual content C is true iff there exists some external-world fact F and C matches F.

There is some reason to think that we at least tacitly endorse criteria of this sort. Consider your total conscious perceptual state at the moment and ask yourself what it would take for that perceptual state to be true (or “accurate” or “veridical,” if you prefer). A tempting answer is that in order for the state to be true, the world would have to be “like *this*,” where “like” refers to similarity and “*this*” refers to some features of your conscious perceptual contents. Which features? Recall that we do not have clear access to our mental contents’ deep natures, so we are not requiring that the world has the same deep nature as our perceptual contents in order for those contents to be true. (Even if we did have such access, it is independently plausible that we don’t take perceptual truth to require a similarity in deep nature—we don’t think the world must be *mental* in order for it to make our perceptual contents true.) Instead, arguably, the relevant features are their superficial characters. In order for the world to be “like *this*,” then, it must share our perceptual contents’ superficial characters—i.e., it must be such that our perceptual contents *match* it.

Our formulation of the matching criteria might need some refinement. For instance, perhaps the criteria we endorse are closer to the following, which also invoke causal conditions:

(Matching-reference2) Perceptual content C refers to external-world item O iff C matches O and C causes our mental state with content O.

(Matching-truth2) Propositional perceptual content C is true iff there exists some external-world fact F and C matches F and F causes our mental state with content O.

Whether we endorse these amended criteria or the original criteria depends on whether we think external world items must cause our mental states in order to be our referents or truth-makers. Suppose the external world happens to match your total conscious perceptual state but that it did not play a role in causing your perceptual state, perhaps because there is a pre-established harmony between the external world and your total perceptual state. If you think that your total perceptual state would then be false, you arguably endorse the amended criteria. If not, you arguably endorse the original matching criteria. (Note that the question of whether we intuitively think *perception* involves a causal relation is distinct from that of whether we intuitively take the *truth* of perceptual contents to involve a causal relation.)<sup>35</sup>

Another possible refinement might be to allow perceptual truth and even reference to be a matter of degree. As in the case of the causal condition, such a refinement would have to be motivated by appeal to considerations concerning what we at least implicitly take to be required for perceptual truth and reference, since what we're after are the criteria we in fact endorse.

One might object that the matching criteria are too demanding, yielding widespread perceptual error and perceptual reference failure. This is because, one might claim, the superficial characters of mental states are not properties that can be instantiated by the external world—they are properties that can only be had by conscious states. In reply, it is important to note whether superficial characters can be had by anything outside of the mind depends on what exactly they end up being. It is not out of the question that at least some superficial characters are abstract or functional in nature and so that they can be instantiated outside the mind.

But suppose it does turn out that superficial characters cannot actually be

---

<sup>35</sup>Thanks to Brian Cutter for urging me to take seriously the possibility that we endorse criteria with a causal condition.

instantiated outside the mind. Then, on the proposed criteria, our perceptual states are never true and never refer. There are two things to note about this outcome. First, this is arguably the right answer—if the external world is indeed wildly different from the superficial character of our perceptual contents, then we really do not perceive the world as it is. Second, this does not mean that we’ve failed at the task of providing a story of external-world truth and reference. Such a story does not require that any given content in fact be true or refer; it only requires that contents have *conditions* of truth and reference. On the scenario under consideration, our perceptual states do have conditions of truth and reference, even though they are never fulfilled. Lest this situation seem too unhappy, note that we are always free to adopt less demanding criteria, such as causal criteria or similarity criteria only requiring a kind of structural similarity.<sup>36</sup> This is the advantage of voluntarism!

Let us now turn to the case of thoughts, which are intentional mental states like those of thinking, occurrently believing, and occurrently desiring. We arguably don’t endorse a mere matching criterion in the case of thought—we don’t require the world to be *similar* to our thought contents in order for those thought contents to be true or to refer. One reason to think this is that thought content is often “symbolic,” involving words, concepts, or images that don’t involve our full intended meanings but that in some sense stand for them. For example, in thinking <the mental supervenes on the physical>, we do not occurrently entertain our full understanding of supervenience but rather a word, image, or partial grasp of the full understanding that nonetheless stands for our full understanding.<sup>37</sup> What, then, do we take to be required for truth and

---

<sup>36</sup>We might instead refine the matching criteria so as to allow truth and reference to be a matter of degree. Then, even if our contents don’t completely match the world, we might nonetheless achieve some degree of truth and reference, perhaps corresponding to certain structural features of our contents’ superficial characters.

<sup>37</sup>See Mendelovici 2018a ch. 7, 2019, and forthcominga.

reference in thought?

As suggested in the previous section, we might have the resources to specify the criteria involved in a recursive story of truth in terms of items we can refer to by acquaintance. This arguably goes a long way towards capturing our intuitive requirements for truth. The fact that such criteria seem correct—even trivial—suggests that they capture what we take to be required for the truth of structured contents.

Such criteria ultimately presuppose reference for objectual and proprietarial contents.<sup>38</sup> How do these contents refer? For example, how does the content <cat> refer to the kind *cat* and <Eleni> refer to the person Eleni?

I want to suggest that these contents refer by descriptions whose components ultimately refer by acquaintance or by satisfying a matching criterion. We have some inkling of which referents we are targeting with the relevant objectual and proprietarial contents, an idea that is captured by descriptions that are associated with these contents. For example, <Eleni> might be associated with the description <my first child>. The contents involved in the description might themselves refer by description. Ultimately, our descriptions must bottom out in contents that refer in some other way—perhaps by acquaintance or by satisfying a matching, causal, or other criterion. I want to suggest that descriptive bedrock includes both contents that refer by acquaintance and contents that refer by matching. The contents that refer by acquaintance might mention particular token conscious states (e.g., the content <my> might ultimately be cashed out in terms of belonging to the bearer of a particular token conscious state that I am acquainted with). The contents that refer by matching might be perceptual contents (e.g., the content of the concept RED might ultimately be cashed out

---

<sup>38</sup> *Objectual contents* are contents with an objectual superficial character—contents like <Justin Trudeau> and <Melbourne>—and proprietarial contents are contents with a proprietarial superficial character—contents like <round> and <red>. See Mendelovici 2018b for this way of understanding objectual and proprietarial contents in terms of superficial characters.

in terms of a range of perceptual contents, which refer by satisfying matching criteria of perceptual reference). Conceptual contents that don't seem to admit of further elucidation in terms of further descriptive contents (perhaps, e.g., <good>) might also refer by matching, perhaps because we endorse a matching criterion of reference for such basic conceptual contents.

## 7 Conclusion

My main aim has been to provide a proof of concept, a sketch of a way in which PIT might provide a voluntaristic theory of truth and reference and thereby meet the challenge set out at the start of this paper. On the picture I've proposed, truth and reference are voluntaristic in that our referents and truth-makers are up to us, either because we are acquainted with them or because we pick them out by description or endorsed criteria. Sometimes a distinction is drawn between original and derived intentionality, where original intentionality is intentionality "in the first instance" and derived intentionality is intentionality that ultimately derives from original intentionality. We can present the overall view in terms of a similar distinction between original and derived truth and reference: original truth and reference—truth and reference that does not require any other truth and reference—is secured by acquaintance, while derived truth and reference—truth and reference that ultimately depends on original truth and reference—proceeds by descriptions or criteria.

In §4.1, we considered two reasons for being skeptical about external-world truth and reference. First, the very idea of external-world truth and reference might be confused and defective—it is unclear what it would be to select something we are not acquainted with. Second, we have no phenomenological or psychological evidence for external-world truth and reference. The account

provided offers responses to these worries: We can see how descriptions and criteria allow us to make sense of how we might select something in the external world even if we are not acquainted with it. And even though we have no direct phenomenological or psychological evidence for external-world truth and reference, we do have at least indirect evidence that the conditions for their existence—namely, those set out by voluntarism—are satisfied.<sup>39</sup>

## References

- BonJour, L. (1998). *In Defense of Pure Reason*. Cambridge University Press.
- Bourget, D. (2010a). Consciousness is underived intentionality. *Noûs*, 44(1):32–58.
- Bourget, D. (2010b). *The Representational Theory of Consciousness*. PhD thesis, Australian National University.
- Bourget, D. (2017). Why are some phenomenal experiences ‘vivid’ and others ‘faint’? representationalism, imagery, and cognitive phenomenology. *Australasian Journal of Philosophy*, 95(4):673–687.
- Bourget, D. (2019a). Implications of intensional perceptual ascriptions for relationalism, disjunctivism, and representationalism about perceptual experience. *Erkenntnis*, 84(2):381–408.
- Bourget, D. (2019b). Relational vs adverbial conceptions of phenomenal intentionality. In Sullivan, A., editor, *Sensations, Thoughts, Language: Essays in honor of Brian Loar*, pages 137–166. Routledge.

---

<sup>39</sup>Thanks to David Bourget, Uriah Kriegel, Walter Ott, Torrance Fung, and audiences at University of Notre Dame, University of Crete, Ruhr-Universität Bochum, and the Central European University for helpful feedback.

- Bourget, D. and Mendelovici, A. (2016). Phenomenal intentionality. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University.
- Brentano, F. (1973/1874). *Psychology from an Empirical Standpoint*. Routledge and Kegan Paul, London.
- Chalmers, D. (2004). The representational character of experience. In Leiter, B., editor, *The Future of Philosophy*, pages 153–181. Oxford University Press, Oxford.
- Chalmers, D. J. (1995). Facing up to the problem of consciousness. *Journal of Consciousness Studies*, 2(3):200–19.
- Crane, T. (2001). Intentional objects. *Ratio*, 14(4):298–317.
- Dretske, F. (1995). *Naturalizing the Mind*. MIT Press, Cambridge.
- Farkas, K. (2008). Phenomenal intentionality without compromise. *The Monist*, 91(2):273–293.
- Fodor, J. A. (1987). *Psychosemantics*. MIT Press, Cambridge.
- Gow, L. (MS). Perceptual experience: Non-relationalism without adverbialism.
- Horgan, T. and Tienson, J. (2002). The intentionality of phenomenology and the phenomenology of intentionality. In Chalmers, D. J., editor, *Philosophy of Mind: Classical and Contemporary Readings*, pages 520–533. Oxford University Press, Oxford.
- Koralus, P. (2014). The erotetic theory of attention: Questions, focus and distraction. *Mind and Language*, 29(1):26–50.
- Kriegel, U. (2007). Intentional inexistence and phenomenal intentionality. *Philosophical Perspectives*, 21(1):307–340.

- Kriegel, U. (2011). *The Sources of Intentionality*. Oxford University Press.
- Kriegel, U. (2013). The phenomenal intentionality research program. In Kriegel, U., editor, *Phenomenal Intentionality*, pages 1–26. Oxford University Press.
- Lewis, D. (1984). Putnam’s paradox. *Australasian Journal of Philosophy*, 62.
- Loar, B. (2003). Phenomenal intentionality as the basis of mental content. *Reflections and Replies: Essays on the Philosophy of Tyler Burge*.
- Mendelovici, A. (2013). Reliable misrepresentation and tracking theories of mental representation. *Philosophical Studies*, 165(2):421–443.
- Mendelovici, A. (2016). Why tracking theories should allow for clean cases of reliable misrepresentation. *Disputatio*, 8(42):57–92.
- Mendelovici, A. (2018a). *The Phenomenal Basis of Intentionality*. Oxford University Press.
- Mendelovici, A. (2018b). Propositionalism without propositions, objectualism without objects. In Grzankowski, A. and Montague, M., editors, *Non-propositional intentionality*. Oxford University Press, Oxford, UK.
- Mendelovici, A. (2019). Immediate and reflective senses. In Shottenkirk, D., Curado, M., and Gouveia, S., editors, *Perception, Cognition, and Aesthetics*, pages 187–209. New York: Routledge.
- Mendelovici, A. (forthcominga). Propositional attitudes as self-ascriptions. In Oliveira, L. L. G. and Corcoran, K., editors, *Commonsense Metaphysics: Essays in Honor of Lynne Rudder Baker*. Oxford University Press.
- Mendelovici, A. (forthcomingb). Reply to woodward. *Philosophical Psychology*.
- Mendelovici, A. and Bourget, D. (2014). Naturalizing intentionality: Tracking theories versus phenomenal intentionality theories. *Philosophy Compass*.

- Mendelovici, A. and Bourget, D. (2020). Consciousness and intentionality. In Kriegel, U., editor, *The Oxford Handbook of the Philosophy of Consciousness*, pages 560–585. New York, USA: Oxford University Press.
- Montague, M. (2016). *The Given: Experience and its Content*. Oxford University Press Uk.
- Nagel, T. (1974). What is it like to be a bat? *The Philosophical Review*, 83(4):435–450.
- Nanay, B. (2010). Attention and perceptual content. *Analysis*, 70(2):263–270.
- Neander, K. (2017). *A Mark of the Mental: A Defence of Informational Teleosemantics*. Cambridge, USA: MIT Press.
- Ott, W. (2016). Phenomenal intentionality and the problem of representation. *Journal of the American Philosophical Association*, 2(1):131–145.
- Papineau, D. (2014). The presidential address: Sensory experience and representational properties. *Proceedings of the Aristotelian Society*, 114(1pt1):1–33.
- Pautz, A. (2007). Intentionalism and perceptual presence. *Philosophical Perspectives*, 21(1):495–530.
- Pautz, A. (2010). Why explain visual experience in terms of content? In Nanay, B., editor, *Perceiving the World*, pages 254–309. Oxford University Press.
- Pautz, A. (2013). Does phenomenology ground mental content? In Kriegel, U., editor, *Phenomenal Intentionality*, pages 194–234. Oxford.
- Pitt, D. (2004). The phenomenology of cognition or what is it like to think that P? *Philosophy and Phenomenological Research*, 69(1):1–36.
- Pitt, D. (2009). Intentional psychologism. *Philosophical Studies*, 146(1):117–138.

- Putnam, H. (1977). Realism and reason. *Proceedings and Addresses of the American Philosophical Association*, 50(6):483–498.
- Putnam, H. (1980). Models and reality. *Journal of Symbolic Logic*, 45(3):464–482.
- Putnam, H. (1983). *Why There Isn't a Ready-Made World*, pages 205–228. Cambridge University Press, Cambridge.
- Russell, B. (1912/1997). *The Problems of Philosophy*. Oxford University Press, New York.
- Schellenberg, S. (2011). Ontological minimalism about phenomenology. *Philosophy and Phenomenological Research*, 83(1):1–40.
- Siewert, C. (1998). *The Significance of Consciousness*. Princeton University Press, Princeton.
- Stoljar, D. (2008). The deflationary theory of truth. In *Stanford Encyclopedia of Philosophy*.
- Strawson, G. (1994). *Mental Reality*. MIT Press, Cambridge.
- Tarski, A. (1943). The semantic conception of truth and the foundations of semantics. *Philosophy and Phenomenological Research*, 4(3):341–376.
- Thau, M. (2002). *Consciousness and Cognition*. Oxford University Press.
- Woodward, P. (2016). Conscious intentionality in perception, imagination, and cognition. *Phenomenology and Mind*, (10):140–155.