

Relational vs adverbial conceptions of phenomenal intentionality

David Bourget

University of Western Ontario

January 1, 2019 (penultimate draft)

*To appear in A. Sullivan, ed., **Sensations, Thoughts, Language: Essays in honor of Brian Loar**, Routledge.*

Much theorizing on intentionality has proceeded on the assumption that it can be fully explained in terms of functional role and causal/informational/correlation-type relations, i.e. that it can be “naturalized”. This assumption remains widespread, but it has come under increased scrutiny over the past couple of decades. Brian Loar (1995) was among the most influential critics of the naturalization project, arguing that intentionality cannot be fully understood without the first-person perspective. He also articulated a picture of the mind on which all intentionality either is or originates from *phenomenal intentionality*, a kind of intentionality that arises from consciousness alone (2002, 2003). Since Loar’s papers on this topic began circulating in the 90s, many philosophers have joined the phenomenal intentionality camp.¹ His rich and inspiring work has been highly influential, helping to usher into the mainstream what is increasingly being perceived as a main contender for a theory of intentionality.

So far, proponents of phenomenal intentionality have been more concerned with defending their common view than adjudicating their differences. I think it is time to look more closely at internal disagreements within the phenomenal intentionality camp. One such disagreement runs particularly deep: according to *relationalists*, phenomenal intentionality is a kind of relation to mind-independent entities (for example, propositions); according to *adverbialists* (or, as I prefer to call them, *aspect theorists*), it is purely non-relational. The choice between these two fundamentally different conceptions of phenomenal intentionality is a crucial but largely overlooked one for the phenomenal intentionality research program. This paper aims to shed some light on this choice.

Loar advocated for a non-relational conception of phenomenal intentionality. I happen to favor a relational view, but I recognize that this view faces many challenges. To echo Loar’s words, “my homage to [Loar] will be expressed by my being driven to extremes” in defending a relational view of phenomenal intentionality in the face of his and others’ objections to such views. I will start by laying out the issue I want to talk about more precisely.²

1 The relation and aspect views

Phenomenal states are mental properties there is something it is like to instantiate and that are individuated by what it is like to instantiate them (their *phenomenology*). *An experience* is an event consisting in the instantiation of a phenomenal state by an individual at a time (or over a

period of time). Correspondingly, *to experience* is to instantiate a phenomenal state at a time.

Proponents of *phenomenal intentionality* (PI) take at least some phenomenal states to give rise to intentionality in that they ground or otherwise necessitate some “aboutness” or “directedness”. I will refer to phenomenal states that give rise to intentionality as *phenomenally intentional states*, or *PI states*. The *content* of a PI state is what the directedness it gives rise to is directed towards. The claim that there are PI states is neutral on whether intentionality is explanatorily or metaphysically prior to consciousness or vice versa.³

A wide range of phenomenal states plausibly give rise to intentionality. For example, the phenomenal state instantiated in a visual experience of a red ball seems to in some way be about a red ball or about there being a red ball. It is “directed” at a red ball, red-ball-ness, or there being a red ball. This seems to be an essential feature of this state that cannot be separated from its phenomenology: change the phenomenology to that of a blue square experience, and the red-ball content is gone; change the target of the directedness to something having to do with blue squares, and the red-ball phenomenology is gone. It is not just visual experiences in normal conditions that are instances of PI states. Plausibly, all perceptual experiences are, and so are episodes of imagery and conscious thoughts.

The *relation view* takes PI states to be states that consist in standing in a relation to their contents, which are distinctly existing entities. For example, one might say that the PI state of experiencing a red ball consists in standing in a relation to the property of being a red ball, whose existence is not tied to that of the experience. The relevant relation might be the same across all PI states, or there could be several relations (for example, one per sensory modality). For present purposes, I am going to assume that there is only one relation (if any) that satisfies the relation view, and I am going to call it *phenomenal representation*.⁴ Nothing here turns on the uniqueness assumption except ease of exposition.

It is non-trivial to explain what it means to say that a state consists in standing in a relation to something as opposed to instantiating a monadic property. One point that is often made is that relations differ from monadic properties in requiring the existence of two or more entities (objects or other kinds of relata), but this isn’t quite right since some binary relations are not irreflexive. However, if A’s being F implies the existence of a thing B distinct from A, and A existing itself does not have this implication, this is at least an indicator that F is a relational property involving B. We can also take it as partly definitional of relations that if A and B stand in a relation, then A and B exist.

We can read Russell (1912) as proposing a relation view in terms of “acquaintance”, though this is no doubt a debatable exegetical claim. More recent and explicit proponents include Chalmers (2006), Pautz (2007, 2009b, 2009a, 2010; forthcoming), Speaks (2015), Woodling (2016), and me (2010, 2017, 2018, forthcoming b). Tracking representationalism (Dretske 1995, Tye 1995, Lycan 1996, 2001) may also be counted as a kind of relation view, but my focus here will be on non-reductive relation views.⁵

It is crucial not to conflate the relation view of phenomenal intentionality with naïve realism, the view that experiences occurring as part of veridical perceptual episodes are relationships to external objects or facts involving external objects. Both views ascribe an aboutness-underpinning relational structure to at least some phenomenal states: PI states in the case of the relation view, states instantiated in veridical experiences in the case of naïve realism. However, PI

states and the states instantiated in veridical experiences need not (and almost certainly don't) coincide, so naïve realism and the relation view don't necessarily ascribe a relational structure to the same mental states. For this reason, neither view entails the other. In addition, naïve realism takes the relata of relational experiences to be ordinary objects like chairs and tables or facts involving such objects, whereas the relation view carries no such commitments. The relation view is consistent with *abstractivism*, the view that PI states are relationships to abstracta such as properties or general propositions. In my opinion and that of many proponents of the relation view, abstractivism is the most plausible form of the relation view, and that is precisely because naïve realism's commitment to experiences of ordinary objects is not very plausible (see Bourget forthcoming b).

We must also distinguish the relation view from externalism about PI states or consciousness. The claim that PI states have a relational structure is at least *prima facie* conceptually independent from the claim that they are wide (not supervenient on a subject's internal state). This is easiest to see if we assume an abstractivist view of the contents of PI states: there seems to be no inconsistency in the idea that facts about one's brain determine that one stands in the phenomenal representation relation to an abstract entity outside of space and time. Phenomenal representation is in this respect like some non-mental relations to abstract objects. For example, something can stand in the relation *X is the mass of Y in kg* to the number 2 simply in virtue of its intrinsic properties.

The main alternative to the relation view that I will be concerned with is the *aspect view*. According to this view, the contents of PI states are aspects of these states rather than relata. An aspect of a mental state might be the mental state itself, a monadic property of the state, a monadic property of a monadic property of the state, or some other monadic property that is intuitively an aspect of the state. Kriegel (2007, 2008, 2011), Pitt (2009), Mendelovici (2018, ch. 9), and Banik (forthcoming) endorse variants on the aspect view. Crane (2013) defends a non-relational view of "objects of thought", which is related in many ways to the aspect view of phenomenal intentionality assuming certain relationships between thought and phenomenal intentionality. For present purposes, I will focus primarily on Mendelovici's and Kriegel's views, which are the most developed competitors to the relation view of phenomenal intentionality with respect to the issues discussed here.

Let us now consider the evidence for and against the relation and aspect views. I will consider what seem to me to be the weightiest evidence on each side, starting with evidence favoring the relation view.

2 For the relation view

2.1 The argument from introspection

2.1.1 *The argument*

When I try to inspect an experience of a red ball, I find something that I can only describe using the phrase "a red ball" or a cognate expression: I can introspectively "see" that I am in a state of experiencing *a red ball*. Whatever it is that I denote with the phrase "a red ball" when I report on my introspectively apparent phenomenology, it is not my mental state itself or any intrinsic, non-relational aspect of my mental state. My words "red" and "ball" pick out properties of sorts that we expect external objects to have, not mental properties or features.^{6,7} Moreover, these words do not

seem to describe a contingent feature of my experience. For example, I am not saying that I am having an experience of a type sometimes caused by a red ball. When I say that I am experiencing a red ball, I am giving a non-contingent description of the phenomenal nature of my experience.

How can I give a non-contingent description of the phenomenal nature of my experience using terms such as “red” and “ball”, which refer to non-mental properties? There are two related possible explanations that I can concretely imagine. First, it could be that my experience is constituted at least in part by entities to which the terms “red” and “ball” refer (let us assume that these are the properties of being red and being a ball, respectively). Second, it could be that my experience has as a constituent something that has the properties to which these terms refers, that is, an actual red ball. Of course, both possibilities might obtain together. So, my experience is at least partly constituted by either an actual red ball or red-ball-y properties or both. From there, it is a short step to conclude that my experience consists in a relationship to these entities (or to one or more entities that have these entities as constituents), since it is hard to see how else my experience could be partly constituted by these entities. More specifically, it seems that these entities, or one or more entities they at least partly constitute, constitute the content of my experience, which is a relatum of the phenomenal state that my experience instantiates. Since these entities are plausibly mind-independent (and so exist distinctly from my experience), this establishes the relation view with respect to my experience of a red ball. Parallel observations apply to all or at least many PI states, which strongly supports the relation view.⁸

2.1.2 An objection

One reply to this argument goes something like this:

What introspection seems to reveal is an actual red ball (say, red ball 22) or that *some* red ball is a constituent of your experience. But we know that the same experience could have occurred in the absence of a red ball. So, introspection is known to be wrong. Why should we take anything it says seriously?⁹

This objection makes two main claims. First, it claims that the apparent revelations of introspection regarding red ball experiences go too far: they imply that a red ball is a constituent of one’s experience when this is not the case. In other words, introspection seems to endorse naïve realism, but naïve realism is false. Second, since introspection is wrong in this respect, we should dismiss what it says about such experiences entirely.

I disagree with both of these claims. Regarding the first claim (that introspection endorses naïve realism and naïve realism is false), note first that the fact that what introspection “says” is correctly reported by my saying that *I am experiencing a red ball* (RED_BALL) does not show that introspection is committed to a red ball being a constituent of my experience. There are three possible construals of RED_BALL and its relation to what introspection says. On the *singular reading*, this report is elliptical: what introspection says is that I am experiencing red ball 22 (say). This interpretation takes introspection’s claim to be best rendered as *Experiencing(I, o)* in predicate logic, where “o” is a name for a specific red ball. The other two possible construals of what introspection says take RED_BALL to be a literal statement of what introspection says. These interpretations are generated by two different understandings of the quantification apparently introduced by the quantified noun phrase “a red ball” in RED_BALL. On the *de re* understanding, “a red ball” is used as in “I am kicking a red ball”: $\exists x(\text{RedBall}(x) \wedge$

Experiencing(I, x)). This reading and the singular reading entail that a red ball is a constituent of my experience. However, RED_BALL can also be given a *de dicto* reading similar to the salient interpretation of “I ordered a red ball”. In this case, there need not be a red ball for the sentence to be true. Montague’s (1974) treatment of such constructions tells us that this is because (on the dicto reading) “a red ball” refers to a generalized quantifier (a property of properties) rather than being a use of such a quantifier. Without putting too much weight on this particular semantic view, “a red ball” seems to denote some abstract entity on the *de dicto* reading of RED_BALL (a generalized quantifier or some other abstract entity). The overall logical form of the *de dicto* reading is *Experiencing*($I, ARedBall$), where “ARedBall” names some abstract entity (not a red ball).

Arguably, what introspection tells me is what RED_BALL says on its *de dicto* reading, not its *de re* or singular reading. First, I do not detect any inconsistency between what introspection tells me and the hypothesis that there is no red ball. Pace some naïve realists, it seems clear that I can identify my relevant state of awareness and coherently think *this could be just like it is in the absence of a red ball*. Second, we would not recognize as introspective any observation that included the claim that there is a red ball out there (characterized generically as on the *de re* reading, or characterized specifically as on the singular reading). How could this count as introspection? It is through perception and theory that we find out about the existence of red balls outside the mind. Lastly, the *de re* and singular readings do not even capture the phenomenology. If I say *there is a red ball, and I am experiencing it*, or *I am experiencing red ball 22*, I am not saying *how* I am experiencing the ball. For example, this is consistent with my experiencing the red ball as blue or square.¹⁰ The *de re* and singular readings of RED_BALL at once say too much (that there is a red ball) and too little (nothing on phenomenology) to capture what I might find introspectively. Therefore, the correct reading of RED_BALL (understood as a phenomenological description directly supported by introspection) is the *de dicto* reading. This reading is not falsified by the fact that my experience could have occurred in the absence of a red ball.

Even if introspection did claim that a red ball is a constituent of my experience, I don’t think we should take this to discredit it to the point that we cannot get supporting evidence for the relation view from it. On my objector’s interpretation of introspection’s claim, we can break up the claim into two parts: my experience is constituted by a certain “foreign object” distinct from my mental states, and that “foreign object” is a red ball. Even if introspection is wrong about the second claim, it might be right about the first. It could be that my experience is a relationship to some property involving red-ball-ness, and I somehow get confused and take it to be about an actual red ball.

2.2 The argument from language

2.2.1 *The argument*

Another argument for the relation view is that the language we use to describe experiences encodes a relational structure. There are many reasons for thinking that RED_BALL, on the intended (*de dicto*) reading, has a relational logical form, binding together a polyadic property with two (or more) entities. Since I have developed parallel reasons for ascribing a relational structure to other *de dicto* perceptual ascriptions in detail in (Bourget 2017a, forthcoming b), I will go over this quickly.

First, notice that RED_BALL supports existential generalization with respect to “a red ball” even

on the de dicto reading. From RED_BALL, it does follow that there is something I experience. Even when I am hallucinating a red ball, it is correct to say that I am experiencing *something*. What does not follow is that I am experiencing an actual red ball, i.e. that there is a red ball such that I am experiencing it. But that is just to say that the intended reading of RED_BALL is not de re. Even on the de dicto reading, there is still something I am experiencing, but not an actual red ball. One hypothesis suggested by Montague's treatment of these constructions is that I am experiencing the following generalized quantifier: the property of being a property had by a red ball.

Relatedly, RED_BALL licenses inferences that would not be valid were it not relational. For example, suppose I know that, due to some apparatus connecting our brains, you experience *everything* that I experience. It seems to follow from this claim and RED_BALL that you experience a red ball. I can see that this follows even if I don't know what a red ball is. This suggests that this follows in virtue of the very logical form of the claims in question. This explanation requires that we take "a red ball" to fill an argument place in RED_BALL.

Consider also that some overtly relational statements seem to be made true by de dicto "experiencing" ascriptions. For example, take *COLOR*, the claim that I am experiencing *something* involving a color. *COLOR* seems to be made true by the fact expressed by RED_BALL. If so, it seems that *COLOR* and RED_BALL must involve the same use of "experiencing", referring to the same property (whether a monadic or polyadic one). But "experiencing" clearly refers to a relation in *COLOR*, so it must refer to a relation in RED_BALL as well.¹¹

Granted that RED_BALL has a relational logical form along the lines of *Experiencing(I, ARedBall)*, how does this show that experiences have a relational structure? There is a logical gap between linguistic expressions having a relational form and the mental states they ascribe having a relational form.

There is more than one way to fill in the gap. The way that I favor first points out that the logical form of RED_BALL is plausibly the logical structure of the proposition it expresses. This seems quite plausible assuming a structured view of propositions. Next, we observe that the form of the proposition expressed is also the form of the fact that putatively makes RED_BALL true. This is *prima facie* plausible, and this falls out of a natural understanding of facts as true propositions: if facts are simply true propositions, then the proposition expressed by RED_BALL just *is* the fact that makes it true, and the structure of the proposition is the structure of the fact. Since the direct object of RED_BALL refers to a property or similar mind-independent entity, this establishes the relation view with respect to the ascribed phenomenal state. We can repeat the argument with multiple ascriptions to convince ourselves that the relation view holds generally.

Another way to fill in the logical gap is inductive. The above arguments show that "to experience" is not merely apparently transitive but also genuinely transitive, contributing a relational form like other transitive verbs. All other verbs that contribute a relational form express relational facts, so it is reasonable to take the relational form of expressions such as RED_BALL to indicate that they ascribe relational properties.

2.2.2 Objections

One might object that the preceding argument adds nothing to the argument from introspection of

the preceding section. The preceding argument from language is based on how we describe experiences, whereas the argument from introspection is based on how they strike us introspectively. Since the manner in which we describe experiences is likely directly guided by how they strike us, one might say that I am double counting the evidence: all that I have offered evidence for is that experience strikes us as relational in introspection.

I don't think this is quite right. Introspection is subject to various influences, including in particular the influence of theory. It seems probable that theorists are prone to introspectively see different things depending on their theories. In contrast, language might be based in some way on introspection, but it encodes how experiences have reliably struck countless users of the language before any theorizing about the question under discussion took place.¹² In this way, the structure of language offers objective evidence regarding how experience appears to us naively and collectively, which one might argue is more valuable and weighty than specific introspective observations made by us today, now that we have adopted views on the question at hand.

Another objection is that the relational analysis of experiential ascriptions given above is not clearly correct because such ascriptions can also be glossed using monadic language, as suggested by adverbialists in response to similar arguments. For example, one might say that "I experienced a red ball" should be glossed as "I experienced-red-ball-y".

This reply misses the point of the above argument from language. The mere possibility of paraphrasing perceptual or experiential ascriptions without using monadic expressions would not show that such ascriptions are not correctly glossed using explicitly relational terms (it is trivial to produce such paraphrases¹³). What one needs to show is that the explicitly relational glosses of perceptual and experiential ascriptions given above and in (Bourget forthcoming b) are *incorrect*, not just *non-mandatory*. Showing this would require taking on the arguments given above that purport to show that at least one correct gloss of the logical form of the relevant expressions is relational. A paraphrase reply to the above arguments would be like saying that "John eats", notwithstanding evidence that it is correctly glossed as *Eats(John)* in predicate logic, doesn't have predicative structure because it can be glossed as "P" in propositional logic.¹⁴

Mendelovici (2018, pp. 268-9) suggests that the argument from language offered in (Bourget forthcoming b), which is similar to the present argument, is susceptible to a debunking reply. The reply casts doubt on the significance of the relational form of experiential ascriptions on the grounds that we would describe mental states using relational language whether or not the relation view was true. According to Mendelovici, we would employ relational language to describe our mental states even if the aspect view were true because this would allow us to economize on words and concepts by reusing words and concepts that purportedly pick out external things.

I don't think it is very plausible that we would speak (and think) relationally even if the aspect view were true. One reason is that it is unclear how we could speak relationally *and* succeed in picking out mental aspects (so it is not clear that speaking relationally is an option if the aspect view is true). If the aspect view is true, every phenomenal state is a purely monadic property of subjects. No phenomenal state is partly constituted by a red ball, the external property of being red, or the external property of being a ball. It is unclear how we might succeed in talking about such monadic states using grammatically and logically relational ascriptions such as RED BALL, in which "red" and "ball" purport to refer to the properties of being red and being a ball. Recall that "red" and "ball" do not merely provide a contingent description of the mental

state in question. If “red” and “ball” refer to properties that are not merely contingently related to the mental state picked out, their referents have to be essential constituents of the mental state picked out, but this is inconsistent with the aspect view unless they refer to mental features, not external properties. It seems clear that “red” and “ball” refer to external properties—if those terms don’t refer to external properties, we are unable to refer to *any* external properties. The only way that RED_BALL could avoid ascribing a relational state is by “red” and “ball” failing to refer, but this would clearly lead to a failure to ascribe a mental state, not an ascription of an aspectual state. It is at best unclear how relational expressions whose constituents purport to refer to mind-independent properties could be used to refer to mental aspects. (Of course, this reply relies on my claim that RED_BALL has a relational form, and premises I have used to motivate this form, but this does not beg the question against Mendelovici’s objection, because she is trying to show that these claims are insufficient to justify the relation view.)

Another problem with the debunking reply is that it is unclear that it would be possible to speak relationally at all on the aspect view (let alone refer to monadic states *by* speaking relationally). On this view, we don’t have the kind of awareness of the world that we naively think we have: we are not in contact with external properties like being red or being a ball, much less external objects. As we will see in the next section, it is unclear how experiences could have truth conditions or refer to (or even purport to refer to) external entities on this view. But Mendelovici’s debunking argument rests on the premise that we start by referring to (or purporting to refer to) external properties, then we start thinking about our experiences by reusing conceptual resources geared toward thinking about a mind-independent reality. This is why we are supposed to find it economical to describe phenomenal states using relational constructions. Insofar as the aspect view makes it unclear how we can think and talk about a mind-independent reality in the first place, it undermines the main premise of the debunking argument.

2.3 Accounting for truth-aptness

PI states seem intrinsically, essentially capable of being true or false, or accurate or inaccurate. (For ease of exposition, I will talk about “truth conditions” and “truth-aptness”, but this is intended to cover accuracy conditions and satisfaction conditions as well.¹⁵ The important point is that there are worldly conditions that experiences are by nature suited to “aim at”.¹⁶) Assuming that PI states are essentially truth-apt, we might ask whether the relation view or the aspect view best accounts for this feature. In this section, I suggest that the relation view offers the most plausible explanation.

2.3.1 *The relational account*

One relational account of truth-aptness claims that experiences are relationships to propositions, where propositions just *are* truth conditions in the form either of sets of possible worlds or intensions (functions from possible worlds to truth values). On this view, it seems trivial to explain how experiences can be true or false. For example, if we take propositions to be sets of possible worlds, we can say that an experience is true when the actual world is a member of the proposition to which it relates the subject.

However, this view is phenomenologically implausible: experiences don’t introspectively seem like relationships to sets of possible worlds or intensions; introspection does not reveal sets of possible worlds or functions from possible worlds (Kriegel 2007 and Mendelovici 2018 make this point). Moreover, this view seems ontologically problematic because experiences being relationships to sets of possible worlds or intensions ties the existence of experiences to that of

relevant merely possible worlds.¹⁷

Typically, sets of possible worlds and intensions are taken to *model* mental contents, not to be constituents of mental states. This is how the preceding concerns are avoided by most philosophers who appeal to possible worlds. But if sets of possible worlds and intensions merely model the truth conditions of experiences, they cannot explain these conditions in the way described above. We need a substantive explanation of how experiences acquire the modeled truth conditions.

An alternative to sets of possible worlds and intensions is to take the contents of PI states to be *possible states of affairs* (PSOAs), which are logical combinations of universals and (optionally) particulars.¹⁸ This view is much more phenomenologically plausible than the preceding views. At least, it relates us to entities that we might reasonably be thought to encounter introspectively, such as the property of being red and the property of being a ball. I will come back to the PSOA view's phenomenological plausibility in section §3.3.4.

The PSOA view naturally goes with the following account of truth conditions: an experience that relates one to PSOA X is true just in case X *is* a fact (or *obtains*). This simple *identity theory of truth* assumes that PSOAs and facts are the same kind of thing.¹⁹

Mendelovici (2018, §9.3.4) anticipates this account and rejects it on the grounds that obtaining and non-obtaining states of affairs are fundamentally different: the obtaining ones are concrete states of the world, whereas the non-obtaining ones are mere abstract objects (at best).

I disagree that obtaining and non-obtaining states of affairs are fundamentally different. Obtaining is an accidental feature of facts (though it is a feature of all facts). When X is a fact, it is that very combination of objects and properties, whose nature allows that it not be a fact, that is a fact, not some other state of affairs.²⁰ Otherwise, every PSOA that does not obtain would necessarily not obtain, but everyone should agree that it is a contingent fact that it is not a fact that zebras eat aluminum. *So long as* we admit possible states of affairs in our ontology, we have to agree that some that don't obtain could have obtained, and so it cannot be part of their nature to be non-obtaining, and the difference between obtaining and non-obtaining states of affairs has to be an accidental feature.

In any case, we can allow that PSOAs are a different kind of thing than facts. This requires us to specify the obtaining property as a kind of correspondence, which we can do. For example, the PSOA <there is a red ball bouncing> consists in the properties *being red*, *being a ball*, and *bouncing* bound together by the existential quantifier. For it to obtain is for there to be something that instantiates these three properties. It is easy to provide a recursive definition of the obtaining property along these disquotational lines for all PSOAs that we might reasonably be thought to experience.

2.3.2 *The aspectual account*

Let us now turn to the aspect theorist's account of truth conditions. If experiences are not relationships to PSOAs, universals, or anything of the sort, they are not composed of properties or other non-mental items that occur in the external world. This makes it hard to see how we can account for the truth conditions of experiences.

Mendelovici (2018) offers an aspect-theoretic account of the truth conditions and reference conditions of PI states (I will continue to focus on truth conditions, but everything I say below extends almost unchanged to her account of reference conditions). Her proposal hinges on a distinction between the superficial character and the deep nature of mental states. The *deep nature* of a mental state is what it “really is”, *au fond*. For example, on some views it is part of the deep nature of mental states that they carry information about things in the environment. In contrast, “[t]he *superficial character* of an intentional state or content is the set of superficial features that characterize it as the intentional state or content that it is” (Mendelovici 2018, p. 25; original emphasis). I had some exchanges with Mendelovici regarding the interpretation of this definition, and we have arrived at the following clarifications. First, superficial characters are properties of contents or mental states. Second, they are the sorts of properties of mental states and contents that we have a pre-theoretic, introspective grip on. Third, superficial characters are uniquely identifying of mental states and contents: if two mental states or contents have the same superficial character, they are the same mental state or content. Understood in this way, it is not trivial that mental states and contents have superficial characters, but I am happy to grant this assumption to Mendelovici.

Mendelovici offers a *matching theory* of the truth conditions of contents: a content C is true just in case there is an item (presumably a fact) distinct from C that instantiates C’s superficial character. Superficial characters, being properties, can be shared between two or more things, including potentially between mental states and external, non-mental items. Importantly, this account is offered as part of a broader internal theory of truth, according to which it is in some sense up to us to decide what are the truth conditions of mental states or contents. Mendelovici takes the matching theory to be what most of us have decided to adopt as our criterion of truth, but the internal theory allows variations.

One consideration in favor of the matching theory (noted by Mendelovici) is that it seems to be in line with how we naively think of truth: we naively think of truth as the world being *like* what we think, where this “like” seems to be the “like” of similarity. The relevant similarity does not seem to be similarity with respect to deep nature (the world need not be mental in order for a mental state to be true!) but rather similarity with respect to superficial character.

The matching theory faces at least two significant challenges. One challenge is that superficial characters can be shared between mental states; as a result, Mendelovici’s theory seems to predict that two experiences with the same superficial character make each other true even if their shared content intuitively has nothing to do with mental states. It is tempting to try to avoid this problem by requiring that the matching item be something non-mental, but this would have the unwelcome consequence that mental states representing mental states could never be true. It seems that we need a conditional criterion that goes something like this: if C is about a mental state, then C’s truth requires a (distinct) matching mental state or fact; otherwise, C’s truth requires a matching non-mental fact (Mendelovici anticipates that this may be required; see appendix H). A worry with this proposal is that it does not accommodate mixed contents such as <an experience caused a movement>. Adding a third condition for such mixed contents along the lines of the above conditions would not work because we need to guarantee that the mental and non-mental parts are made true by mental and non-mental aspects of reality, respectively. Another worry with the conditional criterion is that it seems viciously circular, since being about a mental state seems to be a matter of referring to a mental state. We are here considering only the case of truth conditions, but Mendelovici must and does offer a parallel criterion for reference (a mental state refers to what matches its superficial character), and that criterion for reference also has the

problem of being satisfied by mental items when it should not. This criterion of reference cannot without circularity be improved by saying that all and only mental states that refer to mental items must match mental items. To make the conditional criterion viable, it seems that Mendelovici must specify an understanding of “being about a mental state” that differs from its natural reading in terms of reference. It is not clear how this should be done.

A second, deeper challenge for the matching theory is that we don’t seem to have the resources to apply it to even a single case. On Mendelovici’s view, when you have an experience there is an intrinsic aspect of your experience that is its content. Contents are ‘[t]hings of the same kind as what we are tempted to describe as what our mental states are “directed at” or what they “say” when we introspect on paradigm cases of intentionality’ (Mendelovici 2018, glossary). Your experience is true if the superficial character of its content is matched by a relevant item. The difficulty here is that locutions such as “what my mental state is directed at” or “what my mental state says” are relational through and through, and their relata introspectively seem to be non-mental things such as PSOAs or facts (to naïve realists). So, it seems that their superficial characters should be PSOAs. *If* this were the case, it would be easy to see how the superficial characters of contents match worldly facts, for reasons adduced above. But Mendelovici denies that the superficial characters of contents are PSOAs or facts. On her view, this understanding of the superficial characters of contents as worldly or abstract entities must be cast aside. Here the best I can do is to focus on the content of my mental state as it strikes me (a PSOA that I have in consciousness) and imagine that its nature as a PSOA is illusory: in fact, it is an aspect of my mental state. When I do this, I need to think of certain features of my content as illusory while preserving others. The problem is that *everything* about the PSOA that I seem to have in consciousness is non-mental. I don’t know what the result of imagining that *this* is an aspect of my mental state should be. This is a little bit like trying to imagine that winter tires are a kind of beverage—what sort of beverage is this? Should I imagine a dark beverage? Why should I suppose that basically everything about my idea of winter tire beverages is mistaken except the color? I can conceive in the abstract that the superficial characters of my contents are features of my mental states that are shared with external items, but I cannot verify this in any particular case because I don’t know how to imagine my contents as mental aspects.²¹

3 For the aspect view

I now turn to considerations that seem to support the aspect view.

3.1 Arguments from externalism and physicalism

Loar (2002, §5) considers and rejects a view on which certain phenomenal states are relationships to property complexes, a variation on the PSOA view. Such phenomenal states, Loar seems to hold, would refer to properties in virtue of being relationships to entities that have these properties as constituents. I will call this kind of reference *constitutive reference*. Loar suggests that thought experiments such as the brain in a vat (BIV) and inverted spectrum show that phenomenal states cannot constitutively refer to properties in this way. According to him, a BIV’s mental states cannot refer to properties such as being red or being a ball because reference is determined by external relations. Since a BIV can share phenomenal states with a normally embedded individual, it follows that phenomenal states are not relationships to properties such as being red or being a ball.

From the point of view of the phenomenal intentionality theory, the weakest premise of this argument seems to be the assumption that all reference depends on external relations. It is hard to

see how the traditional arguments for externalism extent to the kind of constitutive reference in question.

A related argument takes its starting point from physicalism: i) physicalism is true; ii) if physicalism and the relation view are true, consciousness can only be explained in terms of informational-causal relations; iii) consciousness cannot be explained in terms of informational-causal relations (because it is narrow); therefore, the relation view is false.²²

The overall problem with this argument is that there is a tension between physicalism and endorsing (ii) or (iii). It seems to be widely accepted that defending physicalism against conceivability arguments requires endorsing the phenomenal concept strategy advocated by Loar (1990, 1997), which centrally involves denying that we have a sufficient grip on the nature of consciousness through our introspectively-grounded phenomenal concepts to be able to assess its relationships to physical facts a priori. Accordingly, a physicalist should not put much weight on thought experiments purporting to show that consciousness is narrow, which undermines much of the motivation for (iii). One should also not be confident that one is able to imagine how consciousness can be explained (against (ii)). For these reasons, the argument is dialectically ineffective.

3.2 Problems for Aristotelian abstractivism

For reasons briefly discussed in §2.1, it is not very plausible that PI states have particulars as constituents of their contents. This means that we are left with an abstractivist view, on which the contents of PI states are properties, property complexes, general PSOAs composed of properties and quantifiers (which can themselves be seen as properties), or other sorts of abstract objects that don't involve concrete particulars. The rest of this paper focuses on objections to abstractivism. I will focus more specifically on the combination of abstractivism and the PSOA view, i.e. the view that PI states relate us to PSOAs composed of properties and quantifiers only.

A key choice point for the abstractivist (especially the proponent of the PSOA view) is what general outlook on the metaphysics of properties to adopt. There are two main views of properties: the Aristotelian view and the Platonic view. The Aristotelian view takes properties to exist in their instances. This makes the existence of properties dependent on the existence of instances. In contrast, the Platonic view takes properties to exist independently of their instances. The choice between the Aristotelian and Platonic conceptions of properties seems to pose a dilemma for the abstractivist, as each kind of abstractivism faces objections. The rest of this section discusses the problems that Aristotelian abstractivism faces; the following section discusses the Platonic alternative.

3.2.1 *Inventory problems*

One problem with Aristotelian abstractivism is briefly touched upon by Loar (2002, 2003). Any experience, it seems, is of a phenomenal type that can occur quite independently of the presence of matching objects. In the extreme case, a BIV that spontaneously appeared in an empty universe could have experiences of the same phenomenal type as ours. This seems to show that the properties that one is related to in experience, if any, are “abstract objects that are unanchored in real resemblance” (2002, §8), that is, not Aristotelian. As noted in the preceding section, proponents of the phenomenal concept strategy (including most physicalists) should reject the BIV intuitions that this argument rests on, but this is an argument that we need to contend with if we do not endorse

this strategy.

Mendelovici (2018) offers a different but related argument against the Aristotelian relation view. She points out that some properties that are good candidates for being the contents of experiences are not good candidates for being Aristotelian. She refers to this as *the inventory problem*. Take the red ball example again. On the face of it, the content of this experience involves a certain color-like property, red*. Red* seems to be a warm, vivid quality that can spread homogeneously over surfaces.²³ Plausibly, the only instantiated properties that have to do with colors (or colors*) are broadly electromagnetic properties, whether monadic or relational, for example, the property of reflecting electromagnetic radiation of about 650 nm (EM650). The problem is that EM650 and all other physical properties of this sort are *prima facie* distinct from red*. On the face of it, no actually instantiated property (or related fact) matches the phenomenology of experiences of red.²⁴

This argument, too, is one that a phenomenal concept strategist would probably reject. In fact, one of the most plausible developments of the phenomenal concept strategy, due to Tye (1999), argues that phenomenal concepts get their special character from our concepts of presented qualities. If one takes this view, one should reject appeals to intuitions of non-identity between red* and EM650.²⁵ However, a relationalist who does not have in her arsenal this general protection against a priori arguments seems forced to endorse Platonic abstractivism.²⁶

3.2.2 *Ontological objections*

Aristotelian abstractivism might also seem to require a dubious ontology. Aristotelian properties are not ontologically questionable, but one might think that false PSOAs or propositions constituted of such properties are questionable, and it seems that relationalism requires that such entities be the contents of inaccurate experiences. Accounting for experiences of impossible scenes seems to be especially problematic, since we don't want to accept that all impossible states of affairs exist (Kriegel 2011a).

I can think of two promising responses to this objection that are consistent with Aristotelianism about properties. First, there is Russell's (1912) response. Russell was moved by this concern, and this led him to deny the existence of mind-independent propositions. Instead, he suggested, what seem like states of acquaintance with propositions are complex mental acts relating us to more than one universal at a time.²⁷

There is a natural way of developing this view in more detail. The PSOAs or propositions that we experience are arguably limited to certain types. For example, there do not seem to be any actual disjunctive experiences, for example, an experience of red or green. The same goes for contents involving negation or universal quantification.²⁸ More controversially, it is not clear that we can experience relations. It is natural to say that we experience an object moving another and other simple relations of this sort, but this might be glossed as experiencing one object instantiating the monadic property of being a mover at space-time position L while experiencing (as part of the same unified experience) another object instantiating the monadic property of being moved at an adjacent space-time position.²⁹ If the contents of experiences really are subject to the preceding limitations, we can naturally think of their contents as set-theoretic constructions from a state space of experienceable monadic properties. Imagine a state space whose dimensions are experienceable spatial and temporal dimensions and other basic experienceable properties such as colors. A (fully or partially defined) position in this state space represents a simple experienced

object with some basic experienceable properties. Volumes or sets of positions in this space represent complex experienced objects constituted of more basic objects. It could be that what seem like experiential representations of complex PSOAs are in fact sets of unified experiences, each of which is a relationship to a volume in the phenomenal state space. Even if we have to allow that all the positions and volumes within the state space in some sense exist, this seems like an ontologically parsimonious picture.

Admittedly, this response to the ontological concerns about PSOAs and propositions depends on substantive and debatable claims regarding the nature of the contents of experiences. However, it seems to me that the state-space story is sufficiently plausible on its face that anyone who is moved by parsimony considerations might reasonably consider endorsing the required assumptions about the nature of the contents of experiences in order to achieve parsimony. There might also be other ways of giving an ontologically minimalistic account of all relevant PSOAs that can accommodate a wider variety of PSOAs.

The alternative to Russell's approach is to argue that PSOAs built from a sparse ontology of properties are ontologically acceptable. So long as we start with a relatively sparse ontology of the constituting properties, it might be OK to admit into our ontology all the PSOAs (true or false) constituted by such properties through a finite set of quantifiers and operators. Since quantifiers and operators can be construed as higher-order properties, all that we need to admit in our ontology are set-theoretic constructions out of (existing) properties. Seen in this light, PSOAs have a very thin existence, somewhat like mereological sums of objects.

Even if neither of the two preceding responses can be made to work, we should not put too much weight on the ontological objection. Physicists talk about branching universes and superpositions of concrete states. For all we know, the world really has a kind of modal structure, with mere possibilities somehow interacting with actual, "concrete" states of the world and each other. We should keep an open mind on ontology.

Relatedly, we should bear in mind that considerations of parsimony can only recommend a theory over another when the two are at least roughly matched in explanatory power (including the theories' fit with observations and ability to make accurate new predictions). We are far from having a complete relational theory of experience, but our sketch of a theory explains the linguistic and introspective evidence discussed earlier, as well as PI states' truth-aptness. It fits fairly well with all observations so long as we allow false PSOAs to exist. Making a case from parsimony against the relation view requires articulating an alternative theory with comparable explanatory power. So far, the aspect view does not seem to have anywhere near as much explanatory power: it can at best explain away the linguistic and phenomenological evidence using auxiliary hypotheses, and it is not clear that it can explain the truth-aptness of experience.

3.3 Problems for Platonic abstractivism

In the preceding section, we saw that Aristotelian abstractivism faces objections that can all plausibly be addressed, but only from within a physicalist framework committed to the phenomenal concept strategy (the inventory problems can be avoided only if we can appeal to the phenomenal concept strategy). For theorists who do not endorse the phenomenal concept strategy, Platonic abstractivism is the only option. This view comes with its own problems.

3.3.1 Additional ontological objections

Traditionally, a central objection to Platonism is that it seems ontologically extravagant. In this vein, Platonic abstractivism seems forced to posit a vast world of causally impotent Platonic entities just to accommodate the apparent relational structure of experience; isn't it more plausible that this apparent relational structure is illusory? This objection from parsimony might seem to have even more bite than that faced by Aristotelian abstractivism.

The first thing to note in response is that the set of properties that we can in practice experience (colors, shapes, etc.) is quite limited compared to the number of possible properties. Add the properties that other kinds of actual conscious beings plausibly experience, and we still have only a finite and relatively small sample of possible properties. One might argue that giving a relational explanation of actual experiences requires only these Platonic properties to exist and that there is no major ontological extravagance there.

One might think that Platonic abstractivism cannot make do with a limited set of Platonic properties because it must accommodate all possible experiences, and all (or a lot of) possible properties are possible contents of experience. I agree that we want a theory such as relationalism to be applicable to merely possible experiences as well as actual experiences, but I disagree that this requires that the contents of possible experiences exist. Take the (correct) theory that high-fiving is relational. This theory's applicability to merely possible high-fivings doesn't require that every person that one could possibly high-five be actual. Possible applications of a theory merely require possible objects.³⁰

Of course, the preceding remarks at best establish that Platonic abstractivism only commits us to a limited set of Platonic properties. Platonic abstractivism might still incur unacceptable ontological commitments when it comes to the representation of full propositions or PSOAs. However, the points made about the ontology of PSOAs and propositions in the preceding section apply here too. It might be possible to analyze PSOAs and propositions away when giving a more fundamental description of experience, and one might reasonably claim that PSOAs and propositions composed of a limited set of properties are parsimonious enough. As before, we should also bear in mind that ontology is hostage to its applications.³¹

One might object that a handful of Platonic properties is as bad as an infinite number of them. I think this depends on whether there is a plausible explanation for why only some possible Platonic properties exist. The reconstruction of experienced (and possibly Platonic) properties as dimensions in a phenomenal state space suggested above seems to provide a possible justification for a sparse view of Platonic properties. Perhaps other explanations are possible.

Regarding this last issue as well as the overall ontological question, we should bear in mind that our views on these questions might be largely shaped by the limits of human categories and imagination. We don't really know what a Platonic object *is*. We know that they are supposed to constitute the contents of experiences, but this leaves their nature largely open. The term "Platonic" evokes imagery such as translucent objects floating in a dark space, but clearly that is not what a Platonic object is. When it comes to grasping the nature of Platonic objects, we seem to be stuck between incomplete and misleading ideas. It could be that these objects make sense ontologically, but we have to form a better understanding of them to see this. It seems wise to keep an open mind on Platonic objects until we are confident that we are not merely butting against the limits of human understanding.

3.3.2 *Efficacy and explanation*

One objection to the ontology-minimizing strategy proposed in the preceding section is that this response seems to require some kind of interaction between abstract objects and concrete mental states: how does nature make sure that we only represent those abstract objects that exist? Given that the Platonic properties in question exist outside of space and time, the suggestion that mental states are restricted to representing those that exist seems mysterious.

This objection is a variant on a general objection to Platonic abstractivism to the effect that abstracta cannot play any role in explanations of concreta (whether causal or constitutive explanations). Papineau writes: “My conscious sensory feelings are concrete, here-and-now, replete with causes and effects. How can their metaphysical nature essentially involve relations to entities that lie outside space and time?” (p. 7) Kriegel (2007, 2011) similarly suggests that concreta cannot even be partly explained by abstracta.

Here it is important to bear in mind that the relationalist is only required to adopt a full-blown Platonic view of the contents of experience if she cannot appeal to the phenomenal concept strategy to solve the inventory problems. Since physicalism requires this strategy, it is reasonable to answer the present objection assuming dualism about consciousness.

Within a dualist framework, we might reasonably suppose that there are psychophysical laws governing interactions between experiences and physical states.³² We should also expect that the laws predict and constrain what can be experienced. Plausibly, if we knew the relevant laws, we would find that there is no residual mystery regarding how nature selects contents that exist. The laws could simply be laid down to match existing Platonic objects, or they could quantify over the domain of existing Platonic objects. There would also be no mystery how states that consist in relationships to abstract objects can be causally efficacious: they would figure in psychophysical laws, which is enough for efficacy. Similarly, abstracta would be part of explanations of concreta in virtue of figuring in psychophysical laws. Of course, we don’t currently have plausible psychophysical laws and the fact that it is very unclear how such laws could be formulated does speak against dualism, but this is not the relationalist’s problem. If dualism can’t be made to work, she will have to accept physicalism, which is consistent with the relation view and supports Aristotelian abstractivism.

3.3.3 *Mystery*

A related objection to Platonic abstractivism is that it seems to posit a mysterious reaching out of the mind towards abstract objects. How exactly can concrete mental states incorporate Platonic objects? This seems wholly mysterious.³³

The objection is simple, and so is the response: of course, it is mysterious! The relation view points out that consciousness has a relational structure, but it is not an explanation of consciousness. It does not purport to close the explanatory gap nor to shed any light on how experiences arise from non-experiential goings-on. Given that it is already maximally mysterious how consciousness arises, it does not seem to be a significant cost if the relation view posits a *prima facie* mysterious relation that bridges the world of concreta and the world of abstracta.

Without closing the explanatory gap, abstractivism plausibly makes the problem of consciousness just a little more tractable. It allows us to redescribe PI states (which arguably include all

phenomenal states) in terms of properties that have nothing to do with the mind (properties such as colors and shapes). This reveals a built-in connection between phenomenal states and the world of mind-independent objects, which cannot hurt in explaining how consciousness arises from the interactions of mind-independent objects. The relation view also articulates the internal structure of the phenomenal to some extent, which is promising. Presumably, the first step in explaining consciousness will be a suitable articulation of the internal structure of phenomenal states. Without some way of breaking up phenomenal states into simpler components, it is hard to see how they could ever be explained. The relation view is a small but significant step in the right direction.³⁴

3.3.4 Phenomenology

Kriegel (2007) raises a phenomenological objection against (Platonic) abstractivism about conscious thoughts: “phenomenologically, the entities we are aware of when we think of dragons and parrots present themselves to us, from the first-person perspective, as external concreta, not as abstracta or mental concreta.” This objection is even more plausible when applied to perceptual experiences as suggested in (Kriegel 2011b).

To a first approximation, concrete objects (concreta) are objects that have (non-trivial) spatial and temporal properties (they are in space and time). Abstracta are things that don’t have (non-trivial) spatial and temporal properties. Kriegel claims that the objects of our thoughts (or experiences), such as the dragons that dragon thoughts are about, seem to be concreta. That is to say that when I introspect a thought (or experience) of a dragon, I find myself thinking of it as being about a dragon in space and time.

It is not clear that this claim is in tension with Platonic abstractivism. On the Platonic view, phenomenally representing that there is an F is a matter of standing in a certain phenomenal representation relation to the PSOA \langle there is an F \rangle . In some cases, F might be a spatiotemporal property such as *being a dragon at position L now*. This makes the representation a representation of a concrete dragon, not an abstract dragon. Since, on the relation view, representing a concrete object is a matter of being related to a spatiotemporal property as opposed to a non-spatiotemporal property, it seems that the theory can accord with the revelations of introspection regarding the concrete nature of represented dragons.

However, there is a further concern that runs deeper than Kriegel’s objection as we have been interpreting it. This is a concern about the abstractness of the contents of experiences (e.g. propositions) rather than the things that are part of the contents of experiences (e.g. dragons). The point of the relation view is to account for the nature and feel of experiences in terms of the relation of a special relation, the phenomenal representation relation. If so, it seems that the nature of what one is related to in the phenomenal representation relation should be reflected in the phenomenal nature of one’s experience. Presumably, it is part of the nature of Platonic entities to be abstract, but we do not find any abstractness when we introspect our experiences.

I think one can reasonably deny that we find nothing abstract in introspection. When I introspect my current visual experience, there is a sense in which I am aware of a way the world could be. But I cannot find anything non-repeatable as part of my experience—no trace of a specific time, location, or even a particular object (as Kriegel [2011b, pp. 163] argues). As a result, it doesn’t seem implausible to say that what I am aware of are abstract ways the world could be and/or Platonic properties that constitute such PSOAs. Of course, it takes a little reflection to

conceptualize what one is aware of in that way (to come to put it in terms of abstractness and so on). But this only goes to show that we don't have a perfect innate reflective grasp of the nature of our own mental states, which is to be expected.

3.3.5 Epistemology

Kriegel (2011a) suggests that Platonic abstractivism (or a more general form of the view targeted at all intentionality) has unacceptable epistemological consequences, whereas adverbialism (his version of the aspect theory) does not. The first claim is justified by pointing out that Platonic abstractivism is inconsistent with a view on which beliefs about particulars can be acquired and (in some sense) justified simply by endorsing perceptual experiences. The reason Platonic abstractivism is supposed to be inconsistent with the endorsement model for perceptual justification is that, on this abstractivist view, experiences have nothing to do with particulars. This seems to require that perceptual beliefs about particulars be justified by some kind of inference. This is reminiscent of naïve realists' claim that singular (and so object-dependent) perceptual contents are necessary to account for singular thought.

I am happy to grant that we do not have singular thoughts about particulars that arise simply by endorsing perceptual experiences. I think it is far from obvious that Russell's view that we know external particulars by description has disastrous epistemological consequences. The part of Kriegel's argument that I take issue with is the claim that adverbialism is better off than abstractivism.

Kriegel correctly notes that adverbialism can preserve an endorsement account of perceptual justification by postulating that perceptual beliefs (not just experiences) are adverbial. This preserves the endorsement model because it is easy to see how beliefs can arise through endorsement of perceptual experiences if they are essentially the same in nature as these experiences. However, this preserves endorsement at the cost of extending the alleged problem to beliefs, which is something that the relationalist can also do. Kriegel suggests that we can restore parity between perceptual beliefs and experiences by supposing that beliefs have the same adverbial (and non-object-involving) nature as perceptual experiences, but a parallel move is also open to the abstractivist, who might say that the contents of perceptual beliefs are abstract like those of perceptual experiences. What the abstractivist can't do is get singular contents out of perceptual experiences, and the adverbialist has the same problem. In fact, the adverbialist has a much larger problem with respect to accounting for the apparent features of the contents of perceptual beliefs: she not only has trouble accounting for singular contents through an endorsement model, she has trouble accounting for the truth conditions of perceptual experiences generally (as we saw above), which makes it hard to see how *any* beliefs could arise through endorsement.

There is a large literature surrounding the idea that views along the lines of the abstractivism defended here have disastrous epistemological consequences; however, this is not the place to discuss these arguments because all of them apply in spades to the aspect view.³⁵

4 Conclusion

This paper explored central considerations for and against the relation and aspect views of phenomenal intentionality. On the one hand, the relation view seems to best fit with the

phenomenological and linguistic evidence we started with. It also enables a more promising approach to truth conditions. On the other hand, physicalism and externalism might seem to militate against the relation view, and the relationalist is forced to choose between Aristotelian and Platonic abstractivism, each of which faces objections. I have argued that the challenges from externalism and physicalism can be defused, and that relationalists who endorse Loar's celebrated phenomenal concept strategy can successfully defend Aristotelian abstractivism, whereas those who don't (and should therefore endorse dualism) can successfully defend Platonic abstractivism.

References

- Banick, Kyle (forthcoming). How to be an adverbialist about phenomenal intentionality. *Synthese*:1-26.
- 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. (2017a). Intensional perceptual ascriptions. *Erkenntnis*, 82(3):513–530.
- Bourget, D. (2017b). Representationalism and sensory modalities: An argument for intermodal representationalism. *American Philosophical Quarterly*, 54(3):251–268.
- Bourget, D. (2017c). The role of consciousness in grasping and understanding. *Philosophy and Phenomenological Research*, 95(2):285–318.
- Bourget, D. (2017d). Why are some phenomenal experiences 'vivid' and others 'faint'? representationalism, imagery, and cognitive phenomenology. *Australasian Journal of Philosophy*, 95(4):673–687.
- Bourget, D. (2018). The rational role of experience. *Inquiry: An Interdisciplinary Journal of Philosophy*, 61(5-6):467–493.
- Bourget, D. (forthcoming a). Anomalous dualism: A new approach to the mind-body problem. In Seager, W., editor, *The Handbook of Panpsychism*. Routledge.
- Bourget, D. (forthcoming b). Implications of intensional perceptual ascriptions for relationalism, disjunctivism, and representationalism about perceptual experience. *Erkenntnis*.
- Bourget, D. and Mendelovici, A. (2014). Tracking representationalism. In Bailey, A., editor, *Philosophy of Mind: The Key Thinkers*, pages 209–235. Continuum.
- Bourget, D. and Mendelovici, A. (2016). Phenomenal intentionality. In *The Stanford Encyclopedia of Philosophy*.
- Chalmers, D. J. (2006). Perception and the fall from eden. In Gendler, T. S. and Hawthorne, J., editors, *Perceptual Experience*, pages 49–125. Oxford University Press.
- Chalmers, D. J. (2010). *The Character of Consciousness*. Oxford University Press.
- Crane, Tim (2009). Is Perception a Propositional Attitude? *Philosophical Quarterly* 59 (236):452-469.
- Crane, Tim (2013). *The Objects of Thought*. Oxford: Oxford University Press.
- Dinges, Alexander (2015). The Many-Relations Problem for Adverbialism. *Analysis* 75 (2):231-237.
- Dretske, F. (1995). *Naturalizing the Mind*. MIT Press.
- Farkas, K. (2008). Phenomenal intentionality without compromise. *The Monist*, 91(2):273–93.
- Fine, K. (1994). Essence and modality. *Philosophical Perspectives*, 8:1–16.
- Fish, William (2009). *Perception, Hallucination, and Illusion*. Oxford University Press.
- Grzankowski, Alex (2018). The determinable-determinate relation can't save adverbialism. *Analysis* 78 (1):45-52.
- Harman, G. (1990). The intrinsic quality of experience. *Philosophical Perspectives*, 4(n/a):31–52.
- 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. Oup Usa.
- Horgan, T. E., Tienson, J. L., and Graham, G. (2004). Phenomenal intentionality and the brain in a vat. In Schantz, R., editor, *The Externalist Challenge*. De Gruyter.
- Kriegel, U. (2003). Is intentionality dependent upon consciousness? *Philosophical Studies*, 116(3):271–307.
- Kriegel, U. (2007). Intentional inexistence and phenomenal intentionality. *Philosophical Perspectives*, 21(1):307–340.
- Kriegel, U. (2008). The dispensability of (merely) intentional objects. *Philosophical Studies* 141 (1):79-95.
- Kriegel, U. (2011a). The Veil of Abstracta. *Philosophical Issues* 21 (1):245-267.
- Kriegel, U. (2011b). *The Sources of Intentionality*. Oxford University Press USA.
- Loar, B. (1990). Phenomenal states. *Philosophical Perspectives*, 4(n/a):81–108.
- Loar, B. (1995). Reference from the first person perspective. *Philosophical Issues*, 6:53–72.
- Loar, B. (1997). Phenomenal states II. In Block, N., Flanagan, O., and Güzeldere, G., editors, *The Nature of Consciousness: Philosophical Debates*. MIT Press.
- Loar, B. (2002). Transparent experience and the availability of qualia. In Smith, Q. and Jovic, A., editors, *Consciousness: New Philosophical Perspectives*. Oxford University Press.
- Loar, B. (2003). Phenomenal intentionality as the basis of mental content. In Hahn, M. and Ramberg, B., editors, *Reflections and Replies: Essays on the Philosophy of Tyler Burge*, pages 229–258. MIT Press.
- Lycan, W. G. (1996). *Consciousness and Experience*. MIT Press.
- Lycan, W. G. (2001). The case for phenomenal externalism. *Philosophical Perspectives*, 15(s15):17–35.
- McGinn, C. (1989). *Mental Content*. Blackwell.
- Mendelovici, A. (2010). *Mental Representation and Closely Conflated Topics*. PhD thesis, Princeton University.
- Mendelovici, A. (2018). *The Phenomenal Basis of Intentionality*. New York, USA: Oxford University Press.
- Mendelovici, A. and Bourget, D. (forthcoming). Consciousness and intentionality. In Kriegel, U., editor, *Oxford Handbook of Philosophy of Consciousness*. New York, USA: Oxford University Press.
- Mendola, J. (2008). *Anti-Externalism*. Oxford University Press.
- Montague, M. (2016). *The Given: Experience and its Content*. Oxford University Press UK.
- Montague, R. (1974). Towards a proper treatment of quantification in english. In Thomason, R. H., editor, *Formal Philosophy*. Yale University Press.
- 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):495–541.
- Pautz, A. (2009a). A simple view of consciousness. In Koons, R. C. and Bealer, G., editors, *The Waning of Materialism*, pages 25–66. Oxford University Press.
- Pautz, A. (2009b). What are the contents of experiences? *Philosophical Quarterly*, 59(236):483–507.
- 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 University Press.
- Pitt, D. (2004). The phenomenology of cognition, or, what is it like to think that p? *Philosophy*

and *Phenomenological Research*, (1):1–36.

Pitt, D. (2009). Intentional psychologism. *Philosophical Studies*, 146(1):117–138.

Russell, B. (1912). *The Problems of Philosophy*. Barnes & Noble.

Searle, J. R. (1990). Consciousness, explanatory inversion and cognitive science. *Behavioral and Brain Sciences*, 13(1):585–642.

Searle, J. R. (1992). *The Rediscovery of the Mind*. MIT Press.

Siewert, C. (1998). *The Significance of Consciousness*. Princeton University Press.

Smithies, D. (2012). The mental lives of zombies. *Philosophical Perspectives*, 26(1):343–372.

Soames, S. (2002). *Beyond Rigidity: The Unfinished Semantic Agenda of Naming and Necessity*. Oxford University Press.

Speaks, J. (2015). *The Phenomenal and the Representational*. Oxford University Press UK.

Tye, M. (1995). *Ten Problems of Consciousness: A Representational Theory of the Phenomenal Mind*. MIT Press.

Tye, M. (1999). Phenomenal consciousness: The explanatory gap as a cognitive illusion. *Mind*, 108(432):705–25.

Woodling, C. (2016) The limits of adverbialism about intentionality. *Inquiry*, 59:5, 488-512.

¹ Proponents of this view (or something close to it) include McGinn (1989), Searle (1990, 1992), Siewert (1998), Horgan and Tienson (2002), Kriegel (2003, 2011b), Loar (2002, 2003), Pitt (2004), Horgan et al. (2004), Farkas (2008), Mendola (2008), Mendelovici (2010, 2018), Chalmers (2010), Bourget (2010a, 2017c, 2018), Smithies (2012), Pautz (2013), Speaks (2015), and Montague (2016).

² Loar uses the quoted passage in (Loar 2003), which was written as an homage to Tyler Burge.

³ For an introduction to phenomenal intentionality and more discussion of the priority question, see (Bourget and Mendelovici 2016).

⁴ I defend this assumption in (Bourget 2010b, 2017b,d).

⁵ By “non-reductive” I mean “not committed to reduction”.

⁶ This echoes Harman’s (1990) transparency claims, but only to a limited extent. In particular, I am not claiming that I cannot introspect my experience itself or anything intrinsic about it.

⁷ One might say that “red” and “ball” have indirect reference like the objects of “that”-clauses on Frege’s theory of indirect reference, but a) this would not really make a large difference to my argument since abstract objects are not mental aspects, b) Frege’s theory of indirect reference has essentially been refuted (Soames 2002), and c) any expression that normally refers to the property of being red or a ball can be substituted *salva veritate* for “red” or “ball” in “I am experiencing a red ball”, which shows that these terms don’t have indirect reference in this context.

⁸ Chalmers (2006) makes a similar argument for (relational) Russellian contents.

⁹ Something like this objection can be found in (Mendelovici 2018, §9.3.1).

¹⁰ One might say that what I find is that there is a red ball and I am experiencing it *as* a red ball. This captures the phenomenology, but this seems revisionist because we do not feel compelled to add the “as” qualification when reporting on our experiences.

¹¹ For more arguments along these lines, see (Bourget 2017, forthcoming b).

¹² I am not sure that the use of the verb “to experience” exhibited above is entirely pre-theoretic. However, it predates theorizing about the structure of consciousness. In any case, the argument can also be made using fully pre-theoretic perceptual language (Bourget 2017, forthcoming b).

¹³ Define the *dash-ly* operator as follows: any sentence containing expressions of the form “A–B...–Z–ly” (with any number of “–”-separated terms) is true iff the ordinary English sentence obtained by removing all occurrences of “–” and the trailing “–ly” is true. With the help of dash-ly, we can construct arbitrarily complex adverbial expressions that leave nothing out. For example, “I am experiencing–a–red–ball–next–to–a–blue–triangle–ly” is true iff “I am experiencing a red ball next to a blue triangle” is true in ordinary English.

¹⁴ Recent discussions of the paraphrase strategy can be found in (Kriegel 2007, 2018), (Dinges 2015), (Grzankowski 2018), and (Banik forthcoming).

¹⁵ Crane (2009) rightly points out that accuracy conditions are quite different from truth conditions. However, I think Crane is wrong in saying that experiences cannot be true. This strains ordinary usage, but we can easily see how the term “true” applies to experiences (see the next footnote).

¹⁶ I don’t think experiences essentially have a direction of fit. Experiencing is a kind of representation that is neither

on its own desire-like nor belief-like. For this reason, it can be a little strained to talk about “truth conditions”; desire-like or other non-belief-like states are not naturally described as having truth conditions. However, I think all experiences are such that they can constitute belief-like states. Add the right attitude-conferring ingredient (perhaps functional features) to an experience, and you have a belief-like state. When I talk about experiences having truth-conditions, I am thinking of the truth/satisfaction conditions they would have were one to add a belief/desire-like attitude to them.

- ¹⁷ This may not be entirely clear in the case of intensions, but the existence of intensions implies the existence of possible worlds if functions are sets of ordered pairs, which is a common view of functions.
- ¹⁸ I include particulars to remain neutral on certain questions, but I don’t think the contents of experiences involve particulars.
- ¹⁹ This account needs to be qualified because not all experiences have PSOs as contents. Only *basic* experiences have PSOs as contents, where a basic experience is one that does not consist simply in having one of a set of distinct experiences. In my view, all experiences that we ascribe using locutions of the form “experiencing NP”, where NP is a noun phrase, are non-basic. For example, experiencing a red ball is non-basic. This is because this ascription is silent on the what the red ball is doing (bouncing, sitting there, etc.). I don’t think it is possible to merely experience a thing doing nothing in particular. So, experiencing a red ball is really a matter of having one of a wide range of distinct experiences, those experiences that involve a red ball doing something. In contrast, experiential ascriptions that have clauses as objects (e.g. “I am experiencing a red ball bouncing”) plausibly ascribe basic experiences. See (Bourget forthcoming b) for more discussion.
- ²⁰ Following Fine (1994), I’m assuming that not all properties that something has necessarily are essential properties it has.
- ²¹ For more arguments for the relation view, see Pautz (2007) and Speaks (2017, ch. 2, ch. 3-9).
- ²² One might be tempted to attribute something like this reasoning to Loar given that he was a staunch physicalist and seemed to endorse the other premises of this argument. However, I am reluctant to attribute the argument to Loar because he does not make this argument explicitly and the argument is flawed for reasons that Loar should have been aware of.
- ²³ I am setting aside the question whether red* is what we normally call “red” in English.
- ²⁴ See also (Chalmers 2006).
- ²⁵ The point here is that *if* one is wedded to physicalism and its implications, one should find it easy to resist the inventory problem using the phenomenal concept strategy. I don’t actually think it is reasonable to be wedded to physicalism, and I think the mismatch between red* and EM650 is a case that brings this out. This is why in (Mendelovici & Bourget forthcoming) I take the position that this mismatch speaks strongly against tracking theories of intentionality.
- ²⁶ Kriegel (2011b) raises a version of the inventory problem involving necessarily uninstantiated properties, such as being an Escher triangle. If experiences represent such properties, the relation view seems committed to their existence. But such properties cannot be Aristotelian. As far as I can tell, all necessarily uninstantiated properties apparently represented by experiences are complex properties: they seem to reduce to combinations of simpler properties. For this reason, I think these properties can be dealt with in the same way as PSOs.
- ²⁷ The overall strategy is the same as the reductionism advocated by Crane (2013).
- ²⁸ Famously, we experience absences, but this can be glossed as a certain thing being “missing” instead of “not being here”.
- ²⁹ I have previously argued that we experience relations (2017b), and I have built an argument against intramodal representationalism on this basis. I had not considered the alternative gloss suggested here. I think the ontological concerns under discussion might be sufficient reason to favor this gloss.
- ³⁰ One might say that Platonic properties are necessary existents, so the mere possibility that a Platonic property is experienced entails that it exists in the actual world. I don’t think we need to assume that Platonic properties exist necessarily. Anyways, this way of being committed to the existence of vast numbers of Platonic properties is independent of the relation view, as an aspect theorist might also have to agree that it is possible that they exist; the question whether necessarily existent Platonic properties are *possible* seems to go beyond the present debate.
- ³¹ An additional answer that is available to the Platonist is that PSOs are simply properties of everything or the world as a whole (Speaks 2015). If Platonic properties of objects are ontologically unproblematic, then so should be properties of the world.
- ³² Those might be stochastic laws as envisaged in (Bourget forthcoming a).
- ³³ Mendelovici (2018, pp. 264–6) makes this argument in an indirect way. She suggests that it is mysterious how content can be “psychologically involved” on the relation view, where one sufficient condition for being psychologically involved is being manifested in phenomenology. This is indirectly a complaint that it is mysterious how content can be manifested in phenomenology on the relation view. This is essentially the problem of consciousness if we take the view that phenomenally intentional states are constituted by their contents.

³⁴ For more on this topic, see (Bourget and Mendelovici 2014) and (Bourget forthcoming a) on the mapping problem.

³⁵ Bill Fish (2009) distills such arguments, and I address his and John McDowell's arguments in Bourget (2010b).