

Singular Experiences (With and Without Objects) *

Angela Mendelovici

November 1, 2021

Abstract

Perceptual experiences seem to in some sense have singular contents. For example, a perceptual experience of a dog as fluffy seems to represent some particular dog as being fluffy. There are important phenomenological, intuitive, and semantic considerations for thinking that perceptual experiences represent singular contents, but there are also important phenomenological, epistemic, and metaphysical considerations for thinking that they do not. This paper proposes a two-tier picture of the content of singular perceptual experiences that is based on phenomenal intentionality theories of intentionality combined with self-ascriptivism about derived representation, a combination of views that allows mental states to have two types of contents: phenomenal contents and derived contents. On the proposed picture, singular perceptual experiences represent singular phenomenal contents, which do

not involve worldly objects, as well as singular derived contents, which do involve worldly objects. This picture accommodates and reconciles the considerations for and against thinking that perceptual experiences have singular contents.

1 Introduction

A distinction is sometimes made between general or abstract contents and objectual or singular contents. For example, consider a perceptual experience of a red, round tomato before you. The experience involves a representation of a *singular* content—the bit corresponding to the tomato—as well as the representation of one or more *proprietal* contents—the bits corresponding to the redness, the roundness, a perceiver-relative location, and perhaps even the tomato-ness. These contents are embedded within what we might call a *propositional* content, a content that represents that something is the case. In this example, the relevant propositional content might be something to the effect of $\langle o \text{ is a red, round tomato before me} \rangle$, where $\langle o \rangle$ is a singular content corresponding to the tomato.

The case of perceptual experiences with apparently singular contents is puzzling. There are strong phenomenological and intuitive considerations in favor of thinking that experiences have such contents, but there are also strong empirical, epistemic, and metaphysical considerations for thinking that they do not. In this paper, I argue that a two-tier view of content based on the phenomenal intentionality theory, a view of intentionality in terms of phenomenal

consciousness, can accommodate and reconcile all these considerations. According to the two-tier view, perceptual experiences have *phenomenal contents*, which are contents the representation of which is nothing over and above the having of particular phenomenal properties, and *derived contents*, which are derived from phenomenal contents. Phenomenal contents can exhibit a kind of “internal singularity” (Sainsbury 2010) but do not involve actual worldly objects as constituents. Derived contents can also exhibit a kind of singularity and sometimes *do* involve worldly objects as constituents.

2 Background

We are concerned with the question of whether perceptual experiences have singular contents. Let us begin by getting clear on the relevant notions of content and perceptual experience.

We can introduce the notion of *intentionality* and the related notion of *content* ostensibly by pointing to paradigm cases: We have thoughts, experiences, and other mental states or events that in some sense seem to “say” or “present” something or be “of”, “about”, or “directed at” something. A visual experience of a blue cup in some sense presents the putative blueness of a cup, is about a cup, or says that a cup is blue. A thought that havanese dogs are charming is in some sense directed at havanese dogs and says that they are charming. In these and other such mundane, everyday cases, we introspectively notice a phenomenon that we are tempted to describe using representational terms like “directness”, “aboutness”, and “saying something”. This phenomenon is *in-*

tentionality. In order for something to qualify as an instance of intentionality, it would have to be an instance of the same phenomenon we notice and are tempted to describe representationally in paradigm cases, such as those mentioned above (though it need not itself be introspectively observable or such that we are tempted to describe it representationally).¹

In the paradigm cases of intentionality mentioned above, we not only notice intentional states but also what the states are “of” or “directed at” or what they “say”. Let us call this an intentional state’s *content*. We can say that a thought or other intentional item *represents* its content.

In what follows, I will be concerned with a type of perceptual experience. *Experiences* are mental states or events exhibiting *phenomenal consciousness*, which is the felt, qualitative, subjective, or “what it’s like” (Nagel 1974) aspect of mental life. The particular felt, qualitative, subjective, or “what it’s like” aspect of a perceptual experience (or other item) is its *phenomenal character*. Perceptual experiences might also have other features, such as times, locations, causal roles, or intentional features. For example, a perceptual experience as of a red, round tomato might occur at a certain time and place, have a reddish-roundish phenomenal character, have a certain causal profile, and represent the content $\langle o \text{ is a red, round tomato} \rangle$.²

In this paper, I will not worry about what distinguishes perceptual experiences from other experiences, such as (perhaps) emotional and cognitive

¹For a defense of this kind of approach to fixing reference on intentionality, see Mendelovici (2018, chapter 1, 2010, chapter 1) and Kriegel (2011, chapter 1).

²It is an open question whether we should take mental states or, instead, subjects to be the bearers of phenomenal and intentional features. For ease of exposition, I will write as if it is mental states that bear such properties.

experiences, or even whether a sharp distinction can be drawn, since I think the view I will ultimately propose will apply to non-perceptual experiences as well. I will take perceptual experiences to include the paradigm cases of visual, auditory, gustatory, olfactory, and somatosensory experiences, as well as experiences in other perceptual modalities, such as equilibrioception. As is customary, I will focus on the more familiar case of vision, though I believe similar points can be made about other perceptual. Of course, that the proposed view applies so broadly should not be uncontroversial, but adjudicating this question is not my focus.

On our fairly neutral definition of “intentionality”, it should be fairly uncontroversial that at least some perceptual experiences represent. Perceptual experiences are among our paradigm cases of intentionality, and no requirements are imposed on intentionality that would automatically preclude them from qualifying. On our neutral definition, perceptual intentionality need not involve intermediaries between the mind and the world, need not relate us to non-mental entities, and need not even be relational. The claim that perceptual experiences represent is compatible with direct realism, sense data theory, and even adverbialism—these are all different views of what the intentionality of perceptual experiences really amounts to. In our terminology, direct realism, for instance, offers a theory on which perceptual experiences represent objects or facts in the mind-independent world by being directly related to them.³

Sometimes a distinction is made between original and derived intention-

³See also Siegel (2010), Schellenberg (2011), and Logue (forthcoming) for alternative fairly neutral ways of understanding content on which it should be uncontroversial that experiences have contents.

ality, where *original intentionality* is intentionality that is not constitutively dependent on other instances of intentionality and *derived intentionality* is intentionality that is not original, i.e., that *is* constitutively dependent on other instances of intentionality. For example, it is sometimes thought that mental items have original intentionality while words and sentences in a language have derived intentionality, which derives from the original intentionality of mental items. In practice, most theories billed as theories of intentionality primarily aim to provide an account of the intentionality of mental items, which, on most views, is original intentionality. However, we will soon see that accounting for the intentionality of perceptual experiences might also require something like a story of derived intentionality.

3 The puzzle of singular experiences

Some perceptual experiences seem to have *singular* contents, contents that in some sense involve or purport to involve a particular item. For example, when you perceptually experience your dog Lily as fluffy, the content of your perceptual experience at least seems to involve Lily herself. Your experience seems to in some sense be *about* or *directed at* Lily. Let us consider some familiar considerations for and against the view that at least some perceptual experiences have singular contents.⁴

⁴ The issue at hand is closely related to the questions of what are the “objects” of experience, whether the phenomenal characters of experience constitutively involve particular worldly objects, and what it is that we can be said to “perceive”. (See Valberg 1992 and Crane and French 2021.) Insofar as these questions are construed as pertaining to the contents of perceptual experiences, the questions are related. But this way of construing these questions is not obligatory. For example, one might think that the question of which object

There are phenomenological, intuitive, and epistemological reasons for thinking that experiences represent some kind of singular content. When you perceptually experience Lily as fluffy, your experience has a particular phenomenology as of there being a particular object present in your experience.⁵ This *phenomenological observation* is one reason to think that our perceptual experiences do represent singular contents of some sort.

There is also a strong *intuition* that Lily herself is part of the content of your experience. It almost seems morally abhorrent to suggest otherwise, to suggest that it is not your beloved dog Lily herself that you represent but merely some mental intermediary, a mere “idea” of Lily, or some cluster of features that Lily happens to have.⁶ This *intuition of object-involving-ness* might be taken to suggest that our perceptual experiences represent singular contents.

Another reason to think that perceptual experiences have singular contents is that this might best explain how perceptual experiences can enable thoughts with singular contents. For example, upon perceptually experiencing a particular dog, I can think about that particular dog. The best explanation of how perceptual experiences enable thoughts with singular contents, one might think, involves attributing singular contents to perceptual experiences. This *consideration from singular thought* is a reason to think that perceptual experiences represent singular contents.⁷

one perceives in having a perceptual experience is entirely independent of the content of that experience.

⁵See, for instance, Schellenberg (2010), Campbell (2006), French and Gomes (2019), Martin (2002).

⁶See, e.g., Martin (2002).

⁷For considerations of this flavor, see McDowell (2002), Campbell (2006), Martin (2002), Brewer (2017), Evans (1982), and Schellenberg (2010).

But there are also strong reasons to think that experiences do not have singular contents, especially if such contents involve existing worldly objects. One well-known reason is that hallucinations can be phenomenally identical to veridical experiences but, presumably, there are no worldly objects to serve as constituents of their contents. For example, consider a hallucination of an elephant walking down the street. This hallucination might be phenomenally identical to a veridical experience of an elephant walking down the street. Both cases have a phenomenology of singularity. But in the case of the hallucination, there is no elephant to be constitutively involved in the content of your perceptual experience. This suggests that particular worldly objects are not part of the contents of our experiences. The case of qualitatively identical objects can be used to make a similar point. A perceptual experience of Lily can be phenomenally identical to an experience of a qualitatively identical dog, Tily, which suggests that these experiences don't involve particular objects as parts of their contents. These *observations concerning hallucination and qualitative duplicates* might be taken to suggest that perceptual experiences do not represent singular contents.⁸

There are also epistemic reasons to think that perceptual experiences do not represent singular contents. In graduate seminars at Princeton, Frank Jackson was known to perform demonstrations such as the following: He would hold up two quarters, baptise them with separate names, and hide them both behind his back. Then he would reveal one quarter and challenge his audience to identify which one it was. If we represent particular objects in experience,

⁸ Along these lines, Mehta (2014) worries that the “phenomenal particularist” makes incorrect predictions about the similarities and differences of different experiences.

we should be able to identify the same object on multiple occasions. But we cannot, so we do not. Jackson takes this to show that the only way that we can identify particular objects is by their associated properties, and not by their actual identity, haecceity, or the like. Such mundane *epistemic observations* suggest that we do not represent singular contents after all.⁹

There are also *metaphysical considerations* against thinking that we represent singular contents. If singular contents involve particular worldly objects that make a difference to our phenomenology, it would seem that those objects would somehow have to get into our conscious minds. But it is utterly mysterious how anything beyond us could enter our consciousness.¹⁰

Another family of considerations concerns the veridicality conditions of perceptual experiences. One might argue that in order for a perceptual experience to be veridical, a particular worldly object must be related to the experience in a particular way. This might be taken to suggest that perceptual experiences have singular contents. But one might also argue that the veridicality of perceptual experiences only requires that *some* object exist and have particular properties, not that a specific object is related to the experience in some way.¹¹ It is not entirely clear how such considerations bear on the question of whether perceptual experiences can have singular contents, since even someone who accepts that a perceptual experience's veridicality requires a relation

⁹ See Jackson (1998) and also McGinn (1982).

¹⁰ Papineau (2021) argues for this point in detail. One worry is that there can be a time lag between the existence of a putatively experienced object and the experience of that object, as in an experience of a star that has gone out of existence before its light reached our retinae. This would make our singular experiences dependent not only on spatially distant events but also temporally past events, which seems exceedingly implausible.

¹¹ Soteriou (2000) interestingly argues that which way one goes here is related to whether one accepts the possibility of veridical hallucinations.

to an object might aim to accommodate this without requiring that the object itself be part of an experience's content.¹² In any case, we might say that such *veridicality considerations* are presumably relevant to the question of whether perceptual experiences represent singular contents, though how exactly they bear on this question is a matter of debate.

I have very briefly overviewed some considerations for and against the claim that we have experiences that represent singular contents:

1. The phenomenological observation: Some perceptual experiences have a phenomenology of singularity.
2. The intuition of object-involving-ness: Some perceptual experiences evoke an intuition of singularity.
3. The consideration from singular thought: Perceptual experiences enable the representation of singular contents in thought.
4. Observations concerning hallucination and experiences of qualitative duplicates: A perceptual experience putatively involving an object can be phenomenally identical to a hallucination and a perceptual experience of a qualitatively identical object.
5. Epistemic observations: A perceptual experience putatively involving a singular content does not put us in a position to discriminate objects from their qualitative duplicates.

¹² See Searle (1983), Brogaard (2010), Bourget (2019), and Shroer (2014).

6. Metaphysical considerations: A worldly object cannot be part of an experience or its content; it is not something that can be experienced or entertained.
7. Veridicality considerations: The veridicality of a perceptual experience might or might not require a specific object to be perceived or otherwise involved in the experience.

These considerations, of course, have been developed in many different ways and these developments have been contested in many different ways. For our purposes, we need not enter these debates. I will accept that there is *something* right about each of these considerations, that there is some way of construing them such that they are onto something. I will suggest that we can perfectly adequately satisfy all these considerations using a two-tier view of content based on the phenomenal intentionality theory, a view of intentionality in terms of phenomenal consciousness. On the proposed view, perceptual experiences have both phenomenal contents and (at least in some cases) derived contents. Phenomenal contents can exhibit a kind of singularity that need not involve worldly objects as constituents. Derived contents can exhibit a kind of singularity that does involve worldly objects as constituents.

In order to present this view in detail, we must first introduce the phenomenal intentionality theory and the view of derived content I propose to combine

it with, self-ascriptivism.

4 Phenomenal intentionality

A theory of intentionality is an account of the deep nature of intentionality, an account of what it really is, metaphysically speaking. *Tracking theories* of intentionality aim to account for intentionality in terms of causal, indication, or informational relations with the environment (Fodor 1987, Dretske 1996, Neander 2017). For example, a tracking theory might take an internal state to represent the type of state of affairs that normally triggers it or that it was selected to be triggered by. Some theories also or instead invoke functional or computational ingredients. Alternative *functional role theories* take an internal state's content to be determined by its functional role, which might be its causal role in inferences or its causal role with respect to other mental states and the environment (Harman 1987, Block 1998, Chalmers forthcoming).

Another kind of view is the *phenomenal intentionality theory (PIT)*, on which original intentionality is nothing over and above phenomenal consciousness, the felt, subjective, or “what it’s like” (Nagel 1974) aspect of mental life. According to PIT, all it is to represent a content is to have the appropriate phenomenal character. On this picture, having a state with a phenomenal character (or a phenomenal character of a certain type) is not having a mere feeling, a mere non-intentional quale. The phenomenal aspects of experience are not mere “mental paint”. They don’t come apart from intentional features of the experience or stand in need of interpretation before they get to repre-

sent anything. Rather, in having an experience with a particular phenomenal character (at least of a certain type), we are already representing, already entertaining a content.¹³

For example, a visual experience with the content <blue square> might represent this content by having a bluish-squarish phenomenal character. A more complex visual experience having an entire represented visual scene as its content might have a complex phenomenal character, involving constituent phenomenal characters corresponding to various colors, shapes, textures, and locations. Thanks to having this phenomenal character, the experience represents a corresponding arrangement of colors, shapes, and textures at certain locations.¹⁴

The same goes for experiences in other sensory modalities. For example, an olfactory experience might represent a sweet and fruity smell in one's vicinity. The phenomenal intentionalist might say that the experience represents this content by having a corresponding sweet-fruity-in-this-vicinity phenomenal character. Likewise, multimodal experiences, which include constituent experiences in distinct modalities, might represent complex contents whose constituent contents are represented by distinct modalities by having correspondingly multimodal phenomenal characters.

¹³PIT and nearby views have been developed and defended by Loar (2003), Kriegel (2011), Pitt (2004, 2009), Farkas (2008, 2013), Siewert (2000), Horgan and Tienson (2002), Woodward (2019, 2016), Bourget (2010), Mendelovici and Bourget (2020), Bourget and Mendelovici (2014), and Mendelovici (2018). For overviews, see Kriegel (2013) and Bourget and Mendelovici (2016).

¹⁴ I am assuming that complex experiences have parts that are also experiences. This is consistent with a reductive view of experiential structure on which complex experiences are nothing over and above their parts arranged in a certain way, but it is also consistent with a non-reductive picture on which complex experiences are something more than the sum of their parts and their arrangements.

My proposal about singular perceptual experiences presupposes that perception has phenomenal contents. It is not my aim here to argue extensively for PIT (I've done this elsewhere).¹⁵ But I want to note that PIT is not implausible in the case of perceptual experience, and, indeed, that the case of perceptual experience provides the basis for an argument for PIT.

PIT is not implausible in the case of perceptual experiences because perceptual experiences are rich in precisely the kinds of phenomenal characters that would—if any phenomenal characters could—give rise to many of the contents perceptual states can arguably be said to represent. Experiences involving the representation of redness have a reddish phenomenal character; experiences involving the representation of squareness have a squarish phenomenal character; experiences representing something as in a particular perceiver-relative location have a phenomenal character corresponding to that perceiver-relative location. It is not implausible that representing the relevant contents is nothing over and above having the corresponding phenomenal characters. Once you have a phenomenal experience with a bluish-squarish-in-front-of-you-ish phenomenal character, you are thereby representing a blue square in front of you—nothing more needs to be added to your overall state in order for you to represent this content.¹⁶

Not only is PIT well-suited for accommodating much of the content of per-

¹⁵ See especially Mendelovici (2018).

¹⁶In contrast to the case of perception, PIT is arguably less plausible in the cases of thoughts (such as occurrent beliefs, occurrent desires, and occurrent “entertainings”), standing propositional attitudes (like standing beliefs and standing desires), and subpersonal representational states (like states representing edges in early visual processing and unconscious representations rules of grammar). Elsewhere, I've argued that PIT can plausibly handle all these cases (Mendelovici 2018, chapters 7 and 8, 2020, Bourget and Mendelovici 2020).

ceptual experience but it is also not clear that any other view can do the trick. According to tracking theories of intentionality, perceptual experiences represent the contents they're connected to via the relevant causal, informational, or other tracking relations. But, arguably, the world is not *quite* as we represent it in experience. We represent colors, but arguably only track surface reflectance profiles; we represent absolute three-dimensional space and time, but arguably only track relativistic space-time; we represent objects as solid through and through, but arguably only track largely empty space that is resistant to penetration. This is not to say that we are massively confused or mistaken (though I think this is also true) but only that the properties available to be tracked don't adequately capture the contents of our experience. If this is right, then not only is the case of perception a good case for PIT—a case it is largely able to accommodate—but it also forms the basis of an argument for PIT. The only way, it seems, that many perceptual contents can be represented is *phenomenally*.¹⁷

5 Self-ascriptivism

In Section 2, we distinguished between original and derived intentionality, where *original intentionality* is intentionality that is not constitutively dependent on other instances of intentionality and *derived intentionality* is intentionality that is not original. Many theories of intentionality recognize both types of intentionality, though the idea that there are *mental* instances of derived

¹⁷I develop these arguments in more detail in Mendelovici (2018, chapters 3 and 4, 2013, 2016) and Mendelovici and Bourget (forthcoming).

intentionality is most commonly seen in versions of PIT.¹⁸ Since phenomenal consciousness appears to be far more scarce than mental instances of intentionality, phenomenal intentionalists often appeal to derived intentionality. Let us take *derivativist* versions of PIT to be versions of PIT that accept mental instances of derived intentionality.

There are many derivativist versions of PIT. For example, Kriegel (2011) appeals to derived intentionality to account for the apparent intentionality of unconscious states, such as standing beliefs and desires and subpersonal states posited by cognitive science. On his view, the content of a subject’s unconscious states is derived from the phenomenal content of a hypothetical ideal interpreter who ascribes those contents to explain and predict the subject’s behaviors. Pautz (2021) proposes a view of a similarly interpretivist spirit, on which subjects derivatively represent the contents assigned to them by the “best interpretation” of their internal states.

There are many other examples of versions of PIT that can arguably be classified as versions of derivativism: Searle’s (1990) dispositionalist view, on which, roughly, potentially conscious states derivatively represent the contents they would phenomenally represent if conscious; Horgan and Tienson’s (2002) view on which broad contents are derived from phenomenal contents together with “grounding presuppositions”; Loar’s (2003) view on which thought contents are derived from “lateral connections” with other mental states; and Bourget’s (2010) view appealing to multiple derivation mechanisms.

Elsewhere, I have defended a version of derivativism (or something near

¹⁸ For defenses of mental cases of derived intentionality, or something near enough, independent of PIT, see Neander (2017), Williams (forthcoming), and Chalmers (forthcoming).

enough) that I called “self-ascriptivism” (Mendelovici 2018, chapter 7, 2020). On this view, we count as derivatively representing the contents we ascribe to our mental states, our phenomenal contents, and ourselves. For example, your concept BACHELOR might not *phenomenally* represent the content <unmarried man>, but it might derivatively represent that content because you are disposed to accept yourself as thinking <unmarried man> when you use the concept BACHELOR and such dispositions are one way in which we can ascribe contents. Elsewhere, I suggested that such derived representation is prevalent in thought, standing states, and even perception.

I think this kind of derived representation is sufficiently different from paradigm instances of intentionality to not be properly considered a kind of intentionality. For one, it is always relative to a subject in that a mental state or other item derivatively represents a content *for* a subject and not in and of itself. Not so for paradigm cases of intentionality. Another reason for thinking that derived representation is sufficiently different from paradigm instances of intentionality is that it can be indeterminate *which* contents are self-ascribed and hence derivatively represented. But the phenomenon we notice in paradigm cases does not admit of this kind of indeterminacy. In any case, it is a terminological choice whether we choose to use the term “intentionality” so that it covers both the having of original contents and the having of derived contents when the two differ in the ways described. I will continue to use the term “derived representation” to pick out the having of derived contents, but nothing hinges on this terminological choice.

The main argument for self-ascriptivism is roughly this:¹⁹

- (P1) Metaphysically necessarily, if S ascribes a content C to an item o, then o derivatively represents C for S.
- (P2) We at least sometimes ascribe contents to our own mental states, contents, or ourselves as a whole.
- (C) At least some of our mental states, contents, or ourselves as a whole derivatively represent some contents for us.

P1 expresses a metaphysically sufficient condition for something to derivatively represent a content for a subject. It's supported by the consideration of paradigm cases of derived representation. If I demonstrate a pen and stipulate "Let this pen hereby stand for the treasure", the pen thereby derivatively represents the treasure for me. If I utter "Let 'P' stand for the proposition that physicalism is true," then, for me, P thereby derivatively represents the proposition that physicalism is true. Similar claims are true if I don't stipulate but instead merely go along with (what I take to be) someone else's stipulation or (what I take to be) a societal convention. I accept that stop signs mean <stop!>, so, simply in virtue of this fact, stop signs mean <stop!> for me. If I mistakenly thought that stop signs meant <go!>, they would mean <go!> for me. Note that the relevant ascriptions need not be realized by occurrent states that I have. I don't have to go around constantly thinking to myself <"dog" means <dog>> in order for "dog" to mean <dog> for me. It is enough that

¹⁹This is an improvement over the argument presented in Mendelovici (2020).

I am disposed to accept that “dog” means <dog> in relevant circumstances, such as when prodded as to what I mean by “dog”.

The above examples illustrate that a subject’s ascribing a content to something is sufficient for that thing to represent that content for the subject. Once a subject stipulates or in some other way accepts that one thing means or stands for something else, no other conditions have to be met in order for the first thing to mean the second for that subject. In this way derivatively representing a content is analogous to having a personal value: Just as a subject’s valuing something is metaphysically sufficient for it to be valuable to that subject, a subject’s taking something to have a content is metaphysically sufficient for that thing to have that content for that subject. When it comes to derived representation, accepting it is so makes it so.²⁰

(P2), in effect, states that we ascribe contents to mental items—to our

²⁰A further claim is that ascriptions are metaphysically necessary for derived representation. I think this further claim is true, and, elsewhere, I used it to argue against alternative versions of derivativism, which do not involve these kinds of ascriptions (2018, section 8.2). Considerations of possible cases again supports this necessity claim. Suppose I am unaware of what a particular road sign means. This road sign might have a meaning for other people or for my community, but it does not mean anything *for me*. Likewise, words in a language that I do not understand mean nothing to me.

One might object that signs and words can be said to derivatively represent a content in and of themselves and not just for specific individuals. This type of derived representation, one might claim, does not require any particular subject’s ascriptions. In the previous example, the road sign might be said to derivatively represent a particular content even though I don’t ascribe to it that content. In this case, it might be right to say that the road sign represents the content even if it’s not right to say that it represents the content *for me*.

This notion of representation arguably picks out an abstraction of the ascribed contents relative to particular subjects (perhaps of subjects belonging to a community or of relevant experts or authorities). After all, if there are no subjects ascribing the appropriate contents to a representation, then there is no sense in which it’s correct to say that they derivatively represent something. If no one ever accepted that, say, stop signs mean <stop!>, stop signs would not mean <stop!>. For present purposes, though, it does not matter whether other versions of derivativism fail to describe metaphysically sufficient conditions for derived representation. It is enough for our purposes that self-ascriptivism does.

mental states, our represented contents, and even to ourselves as a whole. Let us call ascriptions in which we ascribe contents to ourselves or such parts of ourselves *self-ascriptions*. (P2), then, states that there are self-ascriptions.

As in the non-mental examples of derived representation described above, there are many ways in which self-ascriptions might be achieved. We might stipulate that a phenomenal content—perhaps a mentally represented word, image, or symbol—stands for some further content. More commonly, we might merely implicitly accept that a phenomenal content that we entertain cashes out into some further content. For example, when I think the thought that the mental supervenes on the physical, I do not entertain a complete understanding of supervenience (which might, in this case at least, be a definition) but rather some impoverished verbal, imagistic, gisty, or partial content (which, incidentally, is a good candidate for being phenomenally represented). But I am disposed to unpack this impoverished content into my complete understanding if I happen to need this further information. For instance, when working out whether I think that the conceptual possibility of zombies shows that the mental does not supervene on the physical, I might unpack my concept of supervenience to retrieve a definition that can help me think through this case. My disposition to accept this content as standing for or meaning some further content is enough for it to derivatively represent that further content.²¹

From (P1) and (P2), self-ascriptivism follows. We derivatively represent at

²¹ In a sense, my concept SUPERVENIENCE is symbolic: I use it in place of a more complex content that it in some sense stands for. This picture is in line with several converging ways of thinking about concepts on which there are thin or impoverished representations used in cognition that are distinct from more complete understandings that can be drawn on as required by the task at hand (Barsalou 1993, 1999, Prinz 2002, Wickelgren 1979, and Eliasmith 2006).

least some contents by self-ascribing them. It is a further question, though, precisely which contents we self-ascribe. Elsewhere, I have argued that many of the contents we take thoughts and standing propositional attitudes, like standing beliefs and standing desires, to have are not phenomenally represented but derivatively represented in the way described by self-ascriptivism (2018, chapters 7–8, 2020).²²

According to PIT, all original intentionality is phenomenal intentionality. In order for the self-ascriptivist picture to be compatible with PIT, any representational states required for the relevant self-ascriptions would have to phenomenally represent the contents in virtue of which they count as ascribing contents or ultimately derive their contents from phenomenal contents. A self-ascription has three components: it represents (1) a phenomenal content, a mental state, or oneself as a whole, (2) some relation of ascription, and (3) the content ascribed.

In order for (1) to be phenomenally represented, we have to be able to phenomenally represent a content that picks out a mental item. We might achieve this with some kind of inward phenomenal demonstration, a kind of embed-

²²Unlike some other versions of derivativism, self-ascriptivism does not apply to the sorts of unconscious subpersonal contents posited by cognitive science, since we do not generally self-ascribe such contents. I think this yields the right answer. While such states can be said to in some sense represent, they do so in a very different sense than that in which personal-level states like thoughts and experiences represent. They might carry information about, causally correlate with, or have the function of indicating various worldly items or features. Or they might be dispositionally related to other states such that the functional or computational role they play is congruent with or elucidated by describing them as representing a particular content. Indeed, personal-level states presumably also represent in such less demanding senses. Despite all this, we might argue, these sorts of representation are not the Real Deal when it comes to intentionality. Representation in the sense in which unconscious subpersonal states represent is an entirely different kind of phenomenon than intentionality. (See Mendelovici 2018, chapter 8.)

ding of a phenomenal content or other mental item in a phenomenal state, the ability to pay a particular kind of attention to a mental item, or a descriptive content that singles out the relevant internal items. Since the relevant items are all internal and either phenomenal or arguably characterizable in terms of phenomenal features (e.g., as the bearer of phenomenal properties), it is not implausible that they can be picked out in one of these ways and that this requires no more than the entertaining of certain phenomenal contents.

It is also not implausible that we can phenomenally represent (2). Earlier we saw that there are many ways in which self-ascriptions can be achieved, some more demanding than others: by overt stipulation, by explicit acceptance, or by implicitly going along with a given meaning. In the case of self-ascriptions, presumably the relevant ascriptions are usually more like implicit acceptances and less like overt stipulations. It is not implausible that such relatively undemanding ascription relations are phenomenally represented. This might involve a phenomenal experience of “going along with” a supposed expansion of a content or accepting that one content is a “better” version of another.

The case of (c), the representation of the ascribed content, is the most interesting and relevant to our purposes. There are two ways in which we might specify an ascribed content, resulting in two different types of derived representation. In the case of *direct derived mental representation*, we ascribe a phenomenal content to a mental item. For example, we might ascribe the phenomenal content <unmarried man> to the phenomenal content <bachelor>. Such an ascription might consist in phenomenally representing the content <by <bachelor>, I mean <unmarried man>>. In this case, the ascription

phenomenally represents the very content it self-ascribes, <unmarried man>, just as it includes the content to which it is ascribed, <bachelor>. This kind of self-ascription requires that we can “mention” phenomenal contents in our phenomenal intentional states, which might be achieved in some of the ways mentioned in our discussion of (1) above.

In the case of *indirect derived mental representation*, in contrast, we ascribe *the referent of* a phenomenal content to a mental item. For example, an indirect ascription might ascribe the set of unmarried men themselves, or perhaps the property of being an unmarried man, to the phenomenal content <bachelor>. Such an ascription might phenomenally represent the content <by <bachelor>, I mean unmarried man>. In this case, the ascription does not phenomenally represent the content it self-ascribes. In this self-ascription, <unmarried man> is a phenomenal content (perhaps a complex descriptive content) that is “used” rather than “mentioned”. What it contributes to the self-ascription is its regular referent, which, presumably, is the set of unmarried men, the property of being an unmarried man, or some other unmarried-man-related worldly item.

Note that the distinction between direct and indirect derived representation also applies to non-mental cases. Suppose I hold up a pen and say, “I hereby stipulate that this pen represents the tallest person on Earth”. There are two things I might mean by my stipulation. First, I might mean that the pen stands for the descriptive content <the tallest person on Earth>. In this case, the pen directly derivatively represents this descriptive content for me. Second, I might mean that the pen stands for the tallest person on Earth themselves. In this

case, the pen directly derivatively represents some particular person.

As mentioned earlier, it needn't be fully determinate which precise contents we derivatively represent. Some self-ascriptions are realized by dispositions to accept certain ways of cashing out our phenomenal contents, and which precise dispositions we have can be indeterminate. Indeed, it can even be indeterminate whether a particular content is used rather than mentioned in a self-ascription. One reason to think that our self-ascriptions are often indeterminate in this way is that, absent philosophical training, we are often not very good at distinguishing between use and mention. In cases of self-ascriptions that are indeterminate between use and mention in this way, it is indeterminate whether we have an instance of direct or indirect derived representation.

It is also possible for us or a given state of ours to determinately represent multiple derived contents at the same time. For example, a given phenomenal content might be such that we are disposed to cash it out by using a phenomenal content *and* by mentioning that same content. In such a case, the phenomenal content would both directly derivatively represent one content and indirectly derivatively represent another, perhaps related, content.

6 Singular phenomenal and derived contents

The case of experiences with apparently singular contents is puzzling. In Section 3, we saw that there are phenomenological, intuitive, and semantic reasons for thinking that they represent some kind of singular contents but there are

also strong empirical, epistemic, and metaphysical reasons for thinking that particular objects do not and cannot form constitutive parts of their contents. Considerations of veridicality might be taken to pull in either direction. I want to suggest that the two-tier picture of content can respect all these considerations.

Key to the account I will recommend is a distinction that is sometimes made between “internal” and “external” singularity (Sainsbury 2010), where *external* singularity is a kind of singularity that involves particular worldly objects, while *internal singularity* does not require any particular worldly objects. There are different ways in which we might understand internal singularity. Perhaps it has something to do with the structure, psychological role, phenomenology, or other features of singular intentional states. But the basic idea is that a perceptual experience can have many of the features of an externally singular experience while lacking a relation to an object.

The two-tier picture of content described in the previous two sections provides a concrete way of making sense of the distinction between internal and external singularity that accommodates and reconciles the considerations for and against thinking that perceptual experiences have singular contents mentioned earlier. Singular perceptual experiences have singular phenomenal contents, which exhibit internal singularity, and at least sometimes have derived contents, which at least sometimes exhibit external singularity.

Let us first consider singular phenomenal contents. As noted earlier, singular experiences have a phenomenology of singularity. There is a distinctive “what it’s like” of experiencing something as being a particular, unique item,

as being its own thing. Call these phenomenal characters *singular phenomenal characters*. Singular experiences of different items might be phenomenally different from one another, but they share a phenomenology of singularity.

Thanks to having *singular phenomenal characters*, singular experiences have *singular phenomenal contents*. They phenomenally represent a particular, unique, unrepeatable item, something that is its own thing. This particular thing is a haver of properties, a locus of predication. Our singular perceptual experiences phenomenally represent their singular contents as having various properties. For example, a perceptual experience of Lily might phenomenally represent an object as being mostly black, warm, and dog-shaped. This involves having a singular phenomenal content to which the proprietorial contents <black>, <warm>, and <dog-shaped> are predicated.²³

As this example illustrates, singular phenomenal contents have singular “forms”, which allow them to play a subject role when forming part of more complex propositional phenomenal contents. What permits singular phenomenal contents to form parts of propositional phenomenal contents is their singular phenomenal character (which, presumably is reflected in various physical, functional, or “syntactic” properties of the vehicles underlying these phenomenal states, which might be physical states or functional states of subjects).²⁴

²³ Which proprietorial contents are phenomenally represented in any given case is an open question. Depending on your views, you might think that a perceptual experience of Lily represents its object as not only black, warm, and dog-shaped but also as being a dog and perhaps even as being a particular dog—as being Lily (and not her twin, Tilly). As I will suggest at the end of this section, the representation of such “high-level” features in perception might be a matter of having an impoverished phenomenal content and a richer derived content.

²⁴The proposal that we represent impoverished singular contents is congruent with psychological models of visual processing invoking “object files” or “visual indices” (Pylyshyn 2001). Such models are supported by observations concerning our ability to represent oc-

Singular phenomenal contents are internally singular, but they need not be externally singular. Their internal singularity at least partly consists in their phenomenal and phenomenal intentional features and the fact that these features give them a form that allows them to play the subject role in propositional contents. But singular phenomenal contents need not involve any external objects—we can have the relevant phenomenal characters, and hence the relevant phenomenal contents, regardless of whether corresponding external objects exist. So, singular phenomenal contents need not be externally singular. Except perhaps in the special case in which we represent our own phenomenal or phenomenal intentional states, they do not include the objects they refer to or in some sense target as constituents.

Let us turn now to the derived contents of singular perceptual experiences. Although our singular perceptual experiences do not phenomenally represent external objects, they do in some sense target external objects. When you perceptually experience Lily, your mental state in some sense aims at or purports to represent the worldly Lily herself. In this way, then, like many of our

cluded objects over time (Carey and Xu 2001), our ability to represent the same object as enduring changes in its properties (Scholl 2007), and our ability to represent objects prior to representing their features (Pylyshyn 2001). (See Skrzypulec (2015) for discussion of these and other empirical considerations for and against the representation of singular contents in vision.) However, some of these models include further claims about the way “object files” or “visual indices” obtain their semantic features, which we need not take on board. For example, Pylyshyn (2001) claims that visual indices secure reference to the items they visually track. Similarly, Cohen (2004), in a discussion of Clark’s (2000) view that visual objects are spatial locations, suggests that Pylyshyn’s work supports the idea that visual objects are instead ordinary objects like cars, tables, and people. These further claims do not simply fall out of the empirical observations but are motivated in part by theoretical considerations concerning our ability to refer to the world. These claims go well beyond the idea that we have internal representations of a singular form and that such internal representations are involved in tracking external items, which are the claims most clearly supported by the empirical findings and the claims that are most clearly supportive of the overall picture I am proposing.

other phenomenal contents, singular phenomenal contents are impoverished—they do not include all the contents that we in some sense target with our intentional states.

Self-ascriptivism allows us to derivatively represent contents that we in some sense aim at but that we do not phenomenally represent. In the previous section, we saw that there are two types of derived representation: direct and indirect. Perceptual experiences with singular phenomenal contents might exhibit derived representation of either kind. For example, you might be disposed to cash out your singular perceptual phenomenal content corresponding to Lily in terms of some descriptive content: <by <Lily>, I mean <the black-and-white tuxedo dog who is causing my present experience with the content <Lily>>>. ²⁵ In this case, your singular content <Lily> directly derivatively represents the full descriptive content that you are disposed to thus produce.

More interestingly for our purposes, our singular experiences can also exhibit indirect derived representation. Suppose you are disposed to cash out your singular perceptual phenomenal content <Lily> in terms of a description that you *use* rather than mention: <the black-and-white tuxedo dog who is causing my present experience with the content <Lily>>>. Then, if your description manages to pick out some actual dog, Lily, that very dog is indirectly derivatively represented. ²⁶

²⁵This might require that we can have intentional states that involve both perceptual and “cognitive” contents or at least that we can have cognitive states that refer to perceptual contents. How to classify such a state needn’t concern us; all that matters is that we can have such states.

²⁶I have suggested that reference to particular objects goes by description, but it need not—self-ascriptivism can be combined with a causal or other theory of perceptual reference. Even if we go with a description theory of perceptual reference, the relevant descriptions need not include all the information we have about an item. We might simply pick out

In this way, perceptual experiences with singular phenomenal contents can derivatively represent further descriptive or object-involving contents, depending on our self-ascriptions. Insofar as we intuitively take our singular perceptual experiences to be “about” worldly objects themselves and not just our best understanding of them, the contents we self-ascribe presumably include indirectly derivatively represented worldly objects themselves. Indeed, our intuitions that we represent such contents arguably reflect our dispositions to accept that we do, which realize our self-ascribing such contents.^{27,28}

Derived contents constitutively involving particular objects are externally singular, though they need not be internally singular in all the ways that singular phenomenal contents are. In particular, they aren’t singular phenomenal contents and they don’t come with a singular phenomenal character. And, while they may correspond to phenomenal contents that can play the subject role in propositional contents, worldly objects themselves, presumably, can’t combine with proprietal contents or form parts of propositional contents, at

the relevant object by a description like <the thing causing *this* perceptual experience>. Also note that nothing that we have said bears on the question of what are the “objects” of perception, i.e., what determines which specific object is “perceived”. Presumably, the answer to this question depends on what we mean by “perceive” and need not directly bear on what are the phenomenal contents, derived contents, and referents of perceptual experiences.

²⁷In general, we should expect our intuitions about particular cases to closely align with our self-ascriptions. In contrast, our theoretical intuitions may be relatively independent of our self-ascriptions. This is one way in which we can be wrong about which contents we derivatively represent—our intuitions about particular cases and hence our self-ascriptions, on the one hand, and our theoretical intuitions, on the other, can come apart.

²⁸In the previous section, I suggested that in some cases it might be indeterminate whether we represent a direct or an indirect derived content. I also suggested that we might determinately represent both a direct and an indirect derived content. In the case of singular phenomenal contents, these scenarios are not unlikely. We might accept both that there are various associated descriptive contents that go beyond our phenomenal contents and that in perception we target particular individuals, or our self-ascriptive states might be indeterminate between the two.

least not proprietarial and propositional contents that we can entertain. In any case, it does not really matter for the success of the overall proposal whether singular derived contents bear some of the marks of internal singular. What is interesting for our purposes is that they at least sometimes are externally singular.

This completes the presentation of the two-tiered picture of singular experiences based on PIT and self-ascriptivism. In section 3, we considered several considerations for and against the claim that perceptual experiences can have singular contents. I now want to briefly suggest that this two-tiered picture can accommodate and reconcile the relevant considerations:

1. The observation that some perceptual experiences have a phenomenology of singularity is accommodated by singular phenomenal contents, the representation of which is nothing over and above the having of singular phenomenal characters. The having of such phenomenal characters does indeed support the claim that perceptual experiences have singular contents, but it only supports the claim that they have singular *phenomenal* contents; it does nothing to support the claim that they have singular *derived* contents, since our having singular derived contents make no difference to our overall phenomenology.
2. The intuition that we take some of our perceptual experiences to have particular worldly objects as parts of their contents is accommodated by singular derived contents, which do at least sometimes involve worldly objects as constituents. The fact that we have this intuition reflects the

fact that we self-ascribe such contents, so it is indeed a good reason for thinking that we do represent them—derivatively. This is the best way to accommodate this intuition while at the same time respecting the observation that hallucinations also yield the same kinds of intuitions and have a phenomenology of singularity and that pairs of experiences that correspond to different objects can have the same kind of phenomenology of singularity (see below).

3. The observation that perceptual experiences can enable thoughts with singular contents is accommodated by taking perceptual experiences to have both singular phenomenal contents and singular derived contents. A perceptual experience's singular phenomenal content provides an internally singular content that can then be represented by thought. There are different stories of how this transfer might occur: perhaps in thought we demonstrate or otherwise refer to the singular phenomenal content represented in perception, or perhaps an entirely new singular phenomenal content is represented in thought and derivatively represents the singular content represented in perception. A perceptual experience's derived contents might also carry over to related thoughts, allowing thoughts to represent singular derived contents, which might involve particular worldly objects. Again, there are different stories of how this might go: perhaps in the derived contents of thought we demonstrate or refer to the derived contents of a related perceptual experience, or perhaps we ascribe the same contents to the thought as those ascribed to the perceptual experience.

4. The observation that perceptual experiences in which we are related to a particular object can be phenomenally identical to hallucinations in which we are not related to any relevant objects and perceptual experiences in which we are related to qualitatively identical objects, all involving an indistinguishable phenomenology of singularity, is accommodated by the fact that all can be alike with respect to singular phenomenal characters and, hence, singular phenomenal contents. Since phenomenal contents are internally but not externally singular, worldly objects make no contribution to the phenomenology of singular experiences, so an experience can be phenomenally identical to another experience related to a different worldly object or no worldly object at all.
5. The observation that a perceptual experience putatively involving a singular content does not put us in a position to discriminate objects from qualitatively identical ones is accommodated by the fact that neither a perceptual experience's singular phenomenal contents nor its singular derived contents allow us to discriminate between objects and qualitative duplicates of those objects. Singular phenomenal contents don't allow us to perform such discriminatory feats because they do not even involve worldly objects in the first place. Perceptual experiences related to the same worldly objects are alike with respect to their singular phenomenal contents. Having perceptual experiences with singular derived contents also does not put us in a position to discriminate between the objects related to our experiences and their qualitative duplicates. This is because we derivatively represent these contents *indirectly*. In indirect

derived representation, we self-ascribe a content without involving the content itself in a self-ascription, allowing us to derivatively represent contents that we can refer to but that we may not be able to recognize perceptually.

6. The claim that a worldly object cannot form a constitutive part of a perceptual experience is accommodated by the fact that singular phenomenal characters and singular phenomenal contents do not involve worldly objects as constitutive parts. So we never do experience or “entertain” worldly objects. We *derivatively* represent them, but this is unproblematic since we don’t experience or “entertain” our derived contents; they form no constitutive part of our perceptual experiences. We merely (at least sometimes) indirectly target them thanks to our self-ascriptions.
7. We noted that it is not clear what are the veridicality conditions of purportedly singular perceptual experiences and so that it was not clear what conclusion to draw from considerations of veridicality conditions. The two-tier view sheds some light on the debate, and might even allow for the explanation and resolution of conflicting intuitions. Singular phenomenal contents and singular derived contents can differ in their veridicality conditions. In particular, singular phenomenal contents might have less demanding veridicality conditions, requiring only that *some* object exist having various phenomenally-represented features. In contrast, singular derived contents might require that a particular worldly object has particular features, which might even include

relations to the subject, depending on which precise contents are self-ascribed and hence derivatively represented. Since it can be indeterminate which derived contents a perceptual experience represents and since a perceptual experience can represent more than one derived content, this affords more leeway in accommodating conflicting intuitions about the veridicality conditions of singular experiences.

I've argued that the two-tier picture of the contents of singular experiences accommodates the above-mentioned considerations for and against thinking that perceptual experiences have singular contents, reconciling them with one another. Indeed, one might suggest, this picture explains why the case of singular experiences is so controversial to begin with. The competing considerations for and against perceptual experiences representing singular contents cannot be reconciled with a single layer of content. This would require a single content to both involve and not involve particular objects, to both make a phenomenal difference and not differ between duplicates and in the case of hallucination, and to play the myriad of other roles the various considerations describe. By accepting two types of singular contents—which are singular in different ways and represented in different ways—we can make good sense of the pattern of considerations for and against singular experiences.

A similar kind of two-tiered picture might help resolve disputes concerning other contested contents of experience, such as various *high-level* perceptual contents like natural and artifactual kinds, meanings, causation, and other contents going beyond the standard and relatively uncontroversial *low-level* perceptual contents like colors, shapes, textures, locations, temperatures (Siegel

2011, Bayne 2009, Brogaard 2018).²⁹ There are phenomenological, metaphysical, epistemic, and empirical reasons for and against thinking that such contents are represented. A two-tiered picture allows us to say that we phenomenally represent some impoverished high-level-related contents but derivatively represent our full understandings of the relevant high-level features or high-level features themselves.

7 Conclusion

Do perceptual experiences represent singular contents, and do these singular contents involve particular worldly objects? I have argued that perceptual experiences represent two different types of contents—phenomenal contents and derived contents—both of which can be singular but only one of which involves particular worldly objects. I have also argued that this picture accommodates all the observations, intuitions, and other considerations concerning singular perceptual experiences—and that it might explain why the question of whether perceptual experiences have singular contents is so controversial to begin with.³⁰

Bibliography

Aizawa, K. (1996). Fred dretske, naturalizing the mind. *Minds and Machines*, 6:425–430.

²⁹ On some ways of categorizing singular contents, they would qualify as high-level contents. How they are categorized does not matter for our purposes.

³⁰ Thanks to David Bourget and participants at Pär Sundström and Alex Sandgren’s workshop “Particularity and Generality in Mind” (Umeå University) for helpful comments and discussion.

- Barsalou, L. W. (1993). Flexibility, structure, and linguistic vagary in concepts: Manifestations of a compositional system of perceptual symbols. In Collins, A., Gathercole, S., Conway, M. A., and Morris, P. E., editors, *Theories of Memory*, page 1. Lawrence Erlbaum.
- Barsalou, L. W. (1999). Perceptual symbol systems. *Behavioral and Brain Sciences*, 22(4):577–660.
- Bayne, T. (2000). Charles sievert, the significance of consciousness. *Philosophy in Review*, 20:217–221.
- Bayne, T. (2009). Perception and the reach of phenomenal content. *Philosophical Quarterly*, 59(236):385–404.
- Block, N. (1998). Conceptual role semantics. In Craig, E., editor, *Routledge Encyclopedia of Philosophy*, pages 242–256. Routledge.
- Bourget, D. (2010). Consciousness is underived intentionality. *Noûs*, 44(1):32–58.
- Bourget, D. (2019). Implications of intensional perceptual ascriptions for relationalism, disjunctivism, and representationalism about perceptual experience. *Erkenntnis*, 84(2):381–408.
- Bourget, D. and Mendelovici, A. (2016). Phenomenal intentionality. In *The Stanford Encyclopedia of Philosophy*.
- Brewer, B. (2017). The object view of perception. *Topoi*, 36(2):215–227.
- Brogaard, B. (2010). Strong representationalism and centered content. *Philosophical Studies*, 151(3):373–392.
- Brogaard, B. (2018). In defense of hearing meanings. *Synthese*, 195(7):2967–2983.
- Carey, S. and Xu, F. (2001). Infants’ knowledge of objects: Beyond object files and object tracking. *Cognition*, 80(1-2):179–213.
- Chalmers, D. J. (forthcoming). Inferentialism, australian style. *Proceedings and Addresses of the American Philosophical Association*.
- Clark, A. (2000). A theory of sentience. *Philosophy*, 77(299):135–138.
- Clark, A. (2006). Attention and inscrutability: A commentary on john campbell, reference and consciousness. *Philosophical Studies*, 127(2):167–193.

- Cohen, J. (2004). Objects, places, and perception. *Philosophical Psychology*, 17(4):471–495.
- Crane, T. and French, C. (2021). The problem of perception. In *Stanford Encyclopedia of Philosophy*.
- Eliasmith, C. (2006). How to build a brain: From function to implementation. *Synthese*, 153(3):373–388.
- Evans, G. (1982). *The Varieties of Reference*. Oxford University Press.
- Farkas, K. (2008). Phenomenal intentionality without compromise. *The Monist*, 91(2):273–93.
- Farkas, K. (2013). Constructing a world for the senses. In Kriegel, U., editor, *Phenomenal Intentionality*, pages 99–115. Oxford University Press.
- Fodor, J. A. (1987). *Psychosemantics: The Problem of Meaning in the Philosophy of Mind*. MIT Press.
- French, C. and Gomes, A. (2019). How naïve realism can explain both the particularity and the generality of experience. *Philosophical Quarterly*, 69(274):41–63.
- Harman, G. (1987). (Nonsolipsistic) conceptual role semantics. In LePore, E., editor, *New Directions in Semantics*, pages 55–81. London: Academic Press.
- Horgan, T. and Tienson, J. (2002). The phenomenology of intentionality and the intentionality of phenomenology. In Chalmers, D. J., editor, *Philosophy of Mind: Classical and Contemporary Readings*, pages 520–533. Oxford University Press.
- Jackson, F. (1998). Reference and description revisited. *Philosophical Perspectives*, 12:201–218.
- Kriegel, U. (2011). *The Sources of Intentionality*. Oxford University Press USA.
- Kriegel, U. (2013). The phenomenal intentionality research program. In Kriegel, U., editor, *Phenomenal Intentionality*. 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.

- Logue, H. (forthcoming). Experiential Content and Naive Realism: A Reconciliation. In Brogaard, B., editor, *Does Perception Have Content?* Oxford University Press.
- Macpherson, F. (2011). *The Admissible Contents of Experience*. Wiley-Blackwell.
- Martin, M. G. F. (2002). Particular thoughts & singular thought. *Royal Institute of Philosophy Supplement*, 51:173–214.
- McGinn, C. (1982). *The Character of Mind*. Oxford University Press.
- Mehta, N. (2014). The limited role of particulars in phenomenal experience. *Journal of Philosophy*, 111(6):311–331.
- Mendelovici, A. (2010). *Mental Representation and Closely Conflated Topics*. PhD thesis, Princeton University.
- 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. (2018). *The Phenomenal Basis of Intentionality*. New York, USA: Oxford University Press.
- Mendelovici, A. (2020). Propositional attitudes as self-ascriptions. In Oliveira, L. R. G. and Corcoran, K., editors, *Common Sense Metaphysics: Themes From the Philosophy of Lynne Rudder Baker*, pages 54–74. Oxford, UK: Routledge.
- Mendelovici, A. and Bourget, D. (2014). Naturalizing intentionality: Tracking theories versus phenomenal intentionality theories. *Philosophy Compass*, 9(5):325–337.
- 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.
- Mendelovici, A. and Bourget, D. (forthcoming). Kolors without colors, representation without intentionality. *Philosophy and Phenomenological Research*.

- Nagel, T. (1974). What is it like to be a bat? *Philosophical Review*, 83(October):435–50.
- Neander, K. (2017). *A Mark of the Mental: A Defence of Informational Teleosemantics*. Cambridge, USA: MIT Press.
- Papineau, D. (2021). *The Metaphysics of Sensory Experience*. Oxford: Oxford University Press.
- Pautz, A. (2021). Consciousness meets Lewisian interpretation theory: A multistage account of intentionality. In Kriegel, U., editor, *Oxford Studies in Philosophy of Mind*.
- Pautz, A. (forthcoming). Review of williams *the metaphysics of representation*. *Mind*.
- 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.
- Prinz, J. J. (2002). *Furnishing the Mind: Concepts and Their Perceptual Basis*. MIT Press.
- Pylyshyn, Z. W. (2001). Visual indexes, preconceptual objects, and situated vision. *Cognition*, 80(1-2):127–158.
- Sainsbury, R. M. (2010). Intentionality without exotica. In Jeshion, R., editor, *New Essays on Singular Thought*.
- Schellenberg, S. (2010). The particularity and phenomenology of perceptual experience. *Philosophical Studies*, 149(1):19–48.
- Schellenberg, S. (2011). Perceptual content defended. *Noûs*, 45(4):714–750.
- Scholl, B. J. (2007). Object persistence in philosophy and psychology. *Mind and Language*, 22(5):563–591.
- Schroer, R. (2014). The goldilocks problem of the specificity of visual phenomenal content. *Canadian Journal of Philosophy*, 44(3-4):476–495.
- Searle, J. R. (1983). *Intentionality: An Essay in the Philosophy of Mind*. Cambridge University Press.

- Searle, J. R. (1990). Consciousness, explanatory inversion and cognitive science. *Behavioral and Brain Sciences*, 13(1):585–642.
- Siegel, S. (2010). Do visual experiences have contents? In -Nanay, B., editor, *Perceiving the World*. Oxford University Press.
- Skrzypulec, B. (2015). Two types of visual objects. *Studia Humana*, 4(2):26–38.
- Smith, N. H. (2002). *Reading McDowell: On Mind and World*. Routledge.
- Soteriou, M. (2000). The particularity of visual perception. *European Journal of Philosophy*, 8(2):173–189.
- Valberg, J. J. (1992). *The Puzzle of Experience*. New York: Oxford University Press.
- Wickelgren, W. A. (1979). Chunking and consolidation: A theoretical synthesis of semantic networks, configuring in conditioning, s-r versus cognitive learning, normal forgetting, the amnesic syndrome, and the hippocampal arousal system. *Psychological Review*, 86(1):44–60.
- Woodward, P. (2016). Conscious intentionality in perception, imagination, and cognition. *Phenomenology and Mind*, (10):140–155.
- Woodward, P. (2019). Phenomenal intentionality: Reductionism vs. primitivism. *Canadian Journal of Philosophy*, 49(5):606–627.