

The Rational Role of Experience*

David Bourget

Draft of June 29, 2017

It is natural to think that we are rationally and causally responsive to the contents of our thoughts, to what they “say” about the world. For example, it seems that we sometimes directly notice entailment relationships and inconsistencies between the contents of our thoughts. When we do this, we seem to reason *on* the contents of our thoughts and to be causally responding to the relationships between these contents. In order to allow that thoughts might also have contents to which we are not rationally and causally responsive, we can state the point as follows: thoughts have *cognitive contents*, contents to which we are rationally and causally responsive.

There has not been much discussion of cognitive content as such, but there is a large literature on narrow content, which is a closely associated notion (as I will argue below, cognitive content is plausibly narrow). Many philosophers

*Thanks to Adam Pautz, Declan Smithies, Dylan Hurry, Fabio Malfara, and Will Laufs for very helpful comments. Thanks also to Angela Mendelovici for her extensive feedback. This paper will appear in a special issue of *Inquiry* edited by Dimitria Gatzia and Berit Brogaard (expected 2018).

are skeptical of the existence of narrow content. One of the chief reasons for this skepticism is that it is hard to see how such a thing as narrow content (whether cognitive or not) could even be possible: how could the brain “reach out” to grab on to propositions on its own? The aim of this paper is to figure out what is the best available explanation of narrow, cognitive content. I won’t try to fully explain how we grab on to propositions, but I hope to say enough to make it more plausible that we do.

I will start with more background on the notion of cognitive content and the aforementioned concerns (section 1). I will then briefly argue (mostly drawing on discussions found elsewhere) that there is reason to take seriously the view that a mental state’s cognitive content is simply the content that enters the subject’s phenomenal consciousness (section 2). I will refer to this view as *the phenomenal theory* of cognitive content. In section 3, I address the main challenge to the phenomenal theory. The challenge is that this theory seems to imply that thoughts that seem to be about abstract or complex matters often don’t have cognitive contents that capture these matters (because cognitive phenomenology is not rich enough). I will suggest that this is an acceptable consequence of the phenomenal theory that explains some apparently irrational behavior that is otherwise hard to explain. I will also briefly discuss how the phenomenal theory fits into the scientific view of the mind as a largely unconscious information processing mechanism (section 4).

1 Cognitive content and narrow content

Consider the following quote, in which Fodor uses a passage from a Sherlock Holmes novel to give us a glimpse of what thinking is like:

Here, for example, is Sherlock Holmes doing his thing at the end of “The Speckled Band”:

I instantly reconsidered my position when . . . it became clear to me that whatever danger threatened an occupant of the room couldn't come either from the window or the door. My attention was speedily drawn, as I have already remarked to you, to this ventilator, and to the bell-rope which hung down to the bed. The discovery that this was a dummy, and that the bed was clamped to the floor, instantly gave rise to the suspicion that the rope was there as a bridge for something passing through the hole, and coming to the bed. The idea of a snake instantly occurred to me, and when I coupled it with my knowledge that the Doctor was furnished with a supply of the creatures from India I felt that I was probably on the right track.

The passage purports to be a bit of reconstructive psychology: a capsule history of the sequence of mental states which brought Holmes first to suspect, then to believe, that the doctor did it with his pet snake. (Fodor 1990, p. 20)

As the Sherlock Holmes passage illustrates and Fodor agrees, we naturally

think of reasoning or thinking as the manipulation of mental states such as *the suspicion that the rope was there as a bridge* or *the idea of a snake*. It is natural to think that there is some kind of machinery of thought, and that this machinery “processes” mental states in broadly the same way that an assembly line processes input materials and components.

As Fodor also notes, it is not unstructured mental states that are manipulated, but mental states with contents, propositions such as *that the rope was there as a bridge*. In Holmes’ monologue as in many reports that we naturally volunteer, we find the idea that thinking is responsive to relationships between the contents of mental states, understood as propositions that the mental states are in some way “about”.

This idea also seems to be part of our conception of rational norms. Very roughly, a rational person is one whose thinking process does not tolerate inconsistencies and which produces correct and relevant inferences. Thoughts are not consistent or inconsistent or justified in abstraction of their contents: it is in virtue of features of their contents and relationships among their contents that they can be said to be consistent or inconsistent or to stand in (doxastic) justificatory relationships to each other. As a result, rational norms are norms enjoining us to be appropriately responsive to the logical features of the contents of thought. Since ought implies can, these norms seem to presuppose that we *can* be responsive to these logical features: through internal activity taking place in the brain, the contents of mental states can be examined and compared, and responses are produced based on what this

inspection process reveals.

It is natural, then, to think of thinking as a process that takes as inputs various mental states and, by causally responding to the contents of these mental states, produces more mental states and behavior. When I talk about *reasoning on content*, this is the causal process that I have in mind. As I already stated in the introduction, I am going to refer to the contents of mental states that we reason on as *cognitive contents*. Note that mental states may or may not have contents beyond their cognitive contents. I am going to remain largely neutral on this question.

The process of reasoning on content occurs in the brain, so a mental state's having a certain cognitive content must be *narrow*, metaphysically determined by intrinsic features of the brain. As Hume points out in Section XV, Book I of the *Treatise*, "The cause and effect must be contiguous in space and time." With the possible exception of quantum entanglement, this observation has been confirmed over and over again by science. Even gravity, which at some point seemed to be an exception, has turned out not to involve any action at a distance. This is why a thermometer must be connected to the substance whose temperature it responds to, for example, by a wire, or by being directly immersed in it. The same goes for the mechanisms that implement reasoning in brains. We know that our reasoning is not responsive to our mental contents in virtue of causal chains that pass through the environment (this would not be reasoning proper). Therefore, the brain activity that constitutes reasoning can only be causally responsive to contents

that are in some sense in the brain. At the very least, mental states' cognitive contents must be metaphysically determined by intrinsic features of brains.¹

While it is plausible that all cognitive contents are narrow, it might be that some narrow contents are not cognitive. In principle at least, there might be mental contents that are metaphysically determined by internal features of one's brain but to which one's reasoning processes do not have access. Such contents would be narrow without being cognitive. For example, one might say that a certain kind of headache has the narrow content *high blood pressure* because it signals the presence of high blood pressure. A headache could have such content without the subject being in a position to reason on this content.

Such possibilities bring out the fact that cognitive content is what proponents of narrow content ought to be interested in, not narrow content per

¹Philosophers have offered various sophisticated replies to arguments from mental causation for narrow content. Burge (1986) finds issues with a number of arguments from mental causation for internalism, but the arguments he considers have little to do with the above argument, which is much simpler. Williamson (1998) considers an argument that is similar to, but not quite the same as, the above argument, but his argument takes a detour through questionable premises. In particular, he assumes that wide mental states need to factor into internal and external components in order for there to be inner mental states that cause behavior. Williamson then takes issue with this assumption, which is incompatible with the primeness of wide states. As Yablo (2003) points out, it is unclear why we would make this assumption: it might well be that behavior has internal mental causes that are not exactly components of wide mental states. Yablo's (1996 and 2003) own response to arguments for narrow content from mental causation relies heavily on his proposed analysis of causal relevance in terms of proportionality and naturalness, and on the assumption that the relevant kind of causation is causal relevance. I would question the second assumption. Causal relevance, as understood by Yablo, is an epistemically charged notion very closely associated with explanation. One might say that a different, more *oomphy* and less epistemic concept of causation is relevant to the above considerations, for example, a process or transmission concept of causation (Russell 1912 and Salmon 1984).

se. As Farkas (2008b) points out, there is nothing special about the skin boundary, so the thesis that there is narrow content seems arbitrary. This thesis seems interesting because it is a proxy to other, less arbitrary claims. One such claim is that mental states have contents to which internal thinking processes are causally responsive.

Many contemporary theorists reject the preceding picture of reasoning. There are two main reasons for this. First, arguments by Putnam (1975), Burge (1979), and Kripke (1980) have convinced many that the contents of mental states are generally *wide* (metaphysically determined in part by factors in the environment). Wide contents cannot be reasoned with in the kind of way required by the above picture.

Second, as Fodor puts it, if “graspings of propositions” are real, “they must be really [sic] something else”, because they are not “plausible candidates for ultimate stuff” (1990, p. 13). At least until recently, there was no plausible account of narrow content in terms of “something else” (there were formalizations of narrow content along broadly two-dimensionalist lines, but no plausible account of what *grounds* narrow content). The only real candidate on the table was short-arm conceptual role semantics (CRS). According to this view, mental states have their narrow contents in virtue of the complex web of causal links that tie them to each other, sensory stimuli (physically conceived), and behavioral outputs (physically conceived). This view suffers from an explanatory gap: it is just not clear why a big mesh of causal connections between physical things would result in any thinking at all. It also suffers

from well-known underdetermination problems.² In the absence of a plausible explanation of putative narrow contents in terms of something else, many have concluded that we don't reason on mental content. The predominant view seems to be that thinking is responsive to properties of the vehicles of mental states that fall short of determining intentional contents.

It is important not to conflate the second line of reasoning against narrow content with the obvious point that reasoning has to be based on concrete features of the vehicles of mental states, and so cannot operate on abstract objects such as propositions per se. Even if we grant that propositions, qua abstract objects, are not causally efficacious in reasoning, we can still hold that a mental state's *having* of a certain proposition as content (a grasping of a proposition, as Fodor puts it) is a concrete event that can be localized in the brain and causally efficacious qua grasping of a certain proposition. Alternatively, we might compromise a little and say that we can be causally responsive to a mental state's content (or something close enough to this) by being causally responsive to concrete features of a mental state that metaphysically determine its content. Neither of these kinds of causal contact with contents can be ruled out at the outset. Skepticism regarding causal contact with propositions arises not from the obvious impossibility of this kind of contact, but from pessimism regarding the prospects of giving a satisfactory account of the grounds of narrow content.

²Underdetermination objections to such views are raised by Putnam (1981, ch. 2, appendix), Kripke (1982), Horgan and Tienson (2002a), Strawson (2008), and Goff (2012), among others.

This skepticism about narrow, and, by extension, cognitive content has prevailed for many years, but the tide is starting to turn. Today, many theorists take the arguments for externalism to show not that mental content is generally wide, but that mental states generally have both wide and narrow contents. Furthermore, a theory of narrow (or cognitive) content that skeptics about narrow content have largely overlooked has come to be widely endorsed: the phenomenal intentionality theory (PIT).³ According to this theory, there is a kind of intentionality that arises solely from phenomenal consciousness and that grounds all other kinds of intentionality (if any).⁴ The next section discusses the application of this theory to cognitive content. This will lead us to identify a core challenge for this application of PIT, which the rest of the paper aims to address.

2 The phenomenal intentionality theory

PIT is promising as an approach to intentionality in general because consciousness and intentionality seem to be intimately related. The paradigm

³Chudnoff (2015, ch. 6) suggests that arguments for externalism also show that phenomenal intentionality is wide, which would undercut attempts to shore up narrow or cognitive content using PIT. I cannot go into this here, but I disagree that the arguments for externalism apply to phenomenal intentionality.

⁴Proponents of PIT, or something close to it, include Strawson (1994, 2000, 2004, 2008), Siewert (1998), Horgan & Tienson (2002a), Horgan et al. (2004), Loar (2003), Kriegel (2003, 2011b), Pitt (2004, 2011), Pautz (2008, 2013), Farkas (2008a, 2008b), Mendola (2008), Bourget (2010), Chudnoff (2011b, 2011a and 2011b, 2015), and Smithies (2011a, 2011b, 2013b). See especially Mendelovici forthcoming for a systematic development and defense of a version of PIT very congenial to the account of cognitive content I defend in this paper. Bourget and Mendelovici 2016 introduces the acronym “PIT” and offers a detailed overview of the debates surrounding PIT.

cases of phenomenal states (states of consciousness) are states that seem to present something. For example, when I perceive (or hallucinate) a red cube in front of me, there is a certain “what it’s like” to my mental state, and there is also a way the world could be that is being presented. I am “told” that the world is a certain way. I experience a red cube, and I also represent a red cube. These two facets of my mental state are intimately related. In fact, it seems inconceivable that I could experience a red cube in just this way, yet not be presented with a red cube. In virtue of having this experience, there is some intentional content (*there is a red cube*) that is there before my mind. We can refer to the intentional content that we entertain simply in virtue of having an experience with a certain phenomenal character as the *phenomenal content* of the experience.⁵

Phenomenal states appear to be narrow. One reason to think that phenomenal states are narrow is that they cause bodily movements. When you are in excruciating pain, for example, it seems that your screams are triggered by the pain without the intermediary of environmental factors. Barring any action at a distance, this can only be the case if the pain is in your brain, where it is triggering the screaming mechanism.⁶

⁵For more on phenomenal intentionality, see Siewert 1998, Crane 2003, Thompson 2003, Chalmers 2004, Siegel 2005, 2007, 2010, Pautz 2010, Smithies 2013a,b, Bourget and Mendelovici 2016, and Mendelovici forthcoming.

⁶This example makes it clear that experiences are not merely structuring causes of bodily movements, which is a kind of causation that is more plausibly compatible with externalism (see Dretske 1995). In “Anomalous Panpsychism” (forthcoming c), I question this kind of evidence for mental causation, but I offer additional evidence that is no less supportive of an internal causal role for experience.

The narrowness of phenomenal states and the fact that they are by nature intentional make their phenomenal contents good candidate ingredients for explaining cognitive contents. It also seems independently plausible that some phenomenal contents are cognitive contents. For example, when we perceptually experience a way the world might be, we seem to have an excellent cognitive grasp on that way the world might be. This is so even if we don't assume that immediate perceptual beliefs themselves involve reasoning: they at least put us in a position to reason about their objects. It seems, then, that at least some cognitive contents are phenomenal contents.

It seems to me that such an account of cognitive content adequately addresses the Fodorian challenge to cognitive content. Fodor's challenge arises from the claim that *if it exists, narrow (or cognitive) content should be explainable in terms of something else*. In the absence of a suitable something else, this can lead one to doubt that cognitive content exists. If we can identify cognitive contents with phenomenal contents, we safeguard the former from elimination, because phenomenal content is the "something else"⁷. Crucially, Fodor's challenge cannot be re-iterated at the level of the explanans, because consciousness is not something that we are prepared to eliminate if it proves impossible to reduce: we do not believe that in order for consciousness to exist it must really be something else.⁸

The fact that we seem able to adequately account for some cognitive

⁷Phenomenal content is at least conceptually distinct from cognitive content, and I take it that this is the relevant standard of otherness, since reduction is consistent with identity.

⁸These issues are discussed in greater detail in Mendelovici and Bourget 2014.

contents in terms of consciousness is very promising for a PIT-like view of cognitive content. So far, no other approach has succeeded at explaining *any* cognitive content. The only alternative approach is short-arm conceptual role semantics, which suffers from an explanatory gap, and, as a result, cannot explain *any* cognitive content.

There are two ways that a PIT-like view of all cognitive contents could be developed. First, we could try to identify all cognitive contents with phenomenal contents. I will refer to this as *the phenomenal theory* of cognitive content. Proponents of such a view of narrow content (and, by extension, cognitive content) include Strawson (2008), Mendelovici (forthcoming), and myself (Mendelovici & Bourget forthcoming). Second, we could try to identify some cognitive contents with phenomenal contents, or otherwise ground them directly in phenomenal contents, while accounting for other cognitive contents less directly in terms of consciousness. A relatively widespread view along these lines is *phenomenal functionalism*.⁹ On this view, some cognitive contents are grounded directly in consciousness, while others are grounded in functional roles that involve states with phenomenal content. For example, one might say that sensory experiences of colors directly represent specific colors in virtue of their phenomenal characters, but the abstract concept of a color gets its content through functional connections to experiences of colors, shapes, and perhaps other things. A number of theorists have put forward versions of phenomenal functionalism for narrow content or closely

⁹This label is borrowed from Pautz (2013), who uses it a little more narrowly.

related views, including Horgan and Tienson (2002a), Loar (2003), Horgan et al. (2004), Mendola (2008), Kriegel (2011a,b), Smithies (2012a, 2013a,b, 2014), Horgan and Graham (2012), and Pautz (2013). Chalmers (2012) is also sympathetic.

The main reason behind phenomenal functionalism's relative popularity (among proponents of PIT-like views) is that a view such as the phenomenal theory seems to require more cognitive phenomenology than there is. In order for the phenomenal theory to account for the cognitive contents of occurrent thoughts in the same way that it accounts for the cognitive contents of a perceptual experience in the above example, there would have to be *cognitive experiences*, states of consciousness that occur as part of thought processes. These cognitive experiences would need to have phenomenal contents that span the complete range of cognitive contents we can have in thought. They would have to be part of all thoughts with cognitive content, or to *be* these thoughts. In other words, there would have to be a lot of cognitive phenomenology. The problem is that there does not seem to be that much cognitive phenomenology. The phenomenal theory also faces a challenge with respect to standing propositional attitudes, which have no phenomenology at all, but I take it that this challenge is derivative on the challenge for occurrent thoughts.¹⁰

¹⁰Most proponents of PIT-like views hold that standing propositional attitudes are dispositional: they are dispositions to have certain occurrent thoughts and/or behave in certain ways. On such a view of standing attitudes, it seems plausible that they do not in themselves have cognitive content, because we don't reason on them at all—we reason on occurrent mental states that they are dispositions to token. In any case, it is plausible

It is important not to exaggerate the challenge with occurrent thoughts. It is arguable that many concepts are associated with a phenomenology that is characteristic of their contents. Consider, for example, the concepts *circle*, *above*, *object*, *wrong*, *animal*, *person*, *meaning*, and *causation*. For each of these, there seems to be something characteristic it is like to use the concept (at least in optimal conditions, if one is putting in enough effort to grasp the idea). The phenomenology that goes with these concepts seems to be non-imagistic and genuinely non-sensory. No doubt, the phenomenal contents associated with these concepts are not equivalent to the dictionary definitions of “circle”, “above”, etc., but this is to be expected. It is clear that the contents of our thoughts only roughly match the conventional meanings of the words we use to express them.¹¹ Bearing this in mind, the phenomenal theory seems to offer a promising account of the cognitive contents of many thoughts, including in particular many thoughts involving relatively basic, plausibly innate concepts such as the above.¹²

However, even granting that there is considerable cognitive phenomenology and associated phenomenal content, it is not pre-theoretically and prima facie obvious that the phenomenal theory can account for *all* the apparent cognitive contents of thoughts. The greatest challenge for the phenomenal theory is that thoughts about highly complex and/or abstract matters often

that we could account for the cognitive contents of propositional attitudes (if any) given a satisfactory account of the cognitive contents of occurrent thoughts.

¹¹I talk about this a little more below.

¹²For more detailed defenses of cognitive phenomenology, see Strawson 1994, Siewert 1998, Horgan and Tienson 2002b, Goff 2012, Smithies 2013a,b, Chudnoff 2015.

seem to have cognitive contents that far outstrip associated phenomenology. Take for example thoughts about democracy. On the face of it, there is no phenomenology characteristic of thinking about democracy. Typically at least, thoughts about democracy are associated with visual and verbal imagery, and perhaps other kinds of sensory imagery, but it doesn't seem that there is a phenomenal state whose phenomenal content has to do with democracy per se. Recall that phenomenal content is content that a mental state has simply in virtue of its phenomenal character. This means that a mental state's phenomenal content is metaphysically determined by its phenomenal character. The visual and verbal imagery that reliably accompanies thoughts about democracy might be sufficient for thinking about things such as *people*, *boxes with slits*, *folded papers*, *check marks*, *the word "democracy"*, *the word "vote"*, etc., but it does not seem to capture the full-blown idea of a *system of government in which free votes by a large, inclusive segment of the population is used to select representatives or decide governance questions directly* (or something like this). Yet it does seem that the above definition, or something similar, has to be the cognitive content of typical thoughts about democracy, because we apply such definitions when assessing the rationality of thoughts about democracy.¹³

Consider, for example, Fred. Fred seems to have a good understanding

¹³Pautz (2008) is also skeptical of complex cognitive phenomenology. In Pautz 2013, he offers an argument against such phenomenology that goes beyond introspection. Mendelovici (forthcoming) is as skeptical of complex and abstract cognitive phenomenology as I am and fleshes out a detailed response to this problem that is closely related to the response I give later in this paper.

of the idea of democracy, as witnessed by his giving us the above definition when prompted. However, Fred insists that the kingdom of Louis XIV was a democracy in which the king had a birth right to rule absolutely until his death. Without knowing anything more about Fred, it is natural to think that he is being irrational: he is having incoherent thoughts. Fred seems so blatantly inconsistent, we have trouble imagining what is on his mind. This understanding of Fred takes the cognitive content of the thought that he expresses using the term “democracy” to include something like the content that he gives us when we ask him for a definition of “democracy”. This view of Fred’s cognitive content seems inconsistent with Fred’s cognitive content being a phenomenal content, because Fred does not have phenomenology that captures this definition. So it seems that democracy thoughts can have complex cognitive contents that outstrip their phenomenal contents.

Thoughts about abstract mathematical objects also illustrate the challenge for the phenomenal theory. Consider, for example, Josie’s occurrent thought to the effect that a 6-cube has 32 vertices (a 6-cube is the six-dimensional analog of a cube). It seems that Josie and anybody else with a minimum of background in geometry can think about a 6-cube and entertain various propositions about it. At the same time, it seems that in entertaining the proposition that a 6-cube has 32 vertices, Josie is making an error of logic or reasoning (broadly construed), because it is logically impossible for a 6-cube to have 32 vertices. So it seems that the cognitive content of a thought about a 6-cube captures the nature of 6-cubes, including the fact that they

necessarily have 64 vertices. Yet it is doubtful that the phenomenology of such thoughts captures the nature of 6-cubes. A normal person, in any case, cannot visualize a six-dimensional figure in the same way that she can visualize a three-dimensional figure, and there seems to be no non-visual phenomenology of six-dimensional figures.¹⁴

I am going to refer to thoughts with apparently highly abstract or complex cognitive contents as *high-level thoughts*. Phenomenal functionalism is a natural response to the challenge that high-level thoughts pose to the phenomenal theory. However, it is a kind of tactical retreat. We start off with a clear, precise, and perfectly intelligible account of cognitive contents: they are phenomenal contents. When this seems to fail, we introduce a qualification: perhaps the contents we can't account for in this way somehow arise from the phenomenally grounded ones and something else. There is really nothing other than functional role to appeal to, so phenomenal functionalism seems to be our only option. This motivation for phenomenal functionalism seems quite ad hoc.¹⁵ As a result, I think we can plausibly say that *if* the problem of high-level thoughts could be addressed consistently with the simple, pure phenomenal theory, this theory would be more attractive.

Another reason for giving the simple phenomenal theory a try first is that

¹⁴In a recent paper (forthcoming b), I argue that there are no differences in kind between perceptual, imagistic, and cognitive phenomenology: only differences in content. Given such a view, the absence of imagistic phenomenology for visual contents implies the absence of phenomenology simpliciter.

¹⁵Pautz (2013) might have a better motivation for phenomenal functionalism from a broadly Lewisian view of propositional attitudes.

phenomenal functionalism seems to suffer from the same explanatory gap that plagues its non-phenomenal cousin, short-arm conceptual role semantics: why should a large tangle of causal relations between phenomenal experiences that have relatively simple or concrete contents somehow *amount to* thinking a more abstract or complex content? Mendelovici (forthcoming) develops this argument in more detail. This prima facie explanatory gap problem, and the additional simplicity and elegance of the simple phenomenal theory, suggest that the latter theory would be much more attractive if its problem with high-level thoughts could be solved. This is what I hope to make progress on in the rest of this paper.

3 A deflationary view of cognitive content

While I am primarily concerned with the problem of high-level *occurrent* thoughts, some key facts about standing propositional attitudes are relevant. Beliefs, desires, and other standing propositional attitudes are not states that we simply see for what they are: in order to figure out what our own beliefs and desires are, we have to engage in some theorizing and reflection. This theorizing makes various assumptions. For example, it assumes that the thoughts that come to mind on any given occasion tend to reflect what we believe, and that what we believe and desire explains what we do. When ascribing standing propositional attitudes to others, we make even more assumptions. For example, we rely heavily on the principle of charity and

on the assumption that people generally say pretty much exactly what they think.

This means that, in general, we can only expect the contents of standing propositional attitudes ascribed by folk psychology to correspond very roughly to what is really going on in the subject's mind. Consider in particular the implications of the assumption that people generally say what they think, which is clearly central to our practice of propositional attitude ascription (in many cases, our grounds for saying that someone believes P is simply that they have said so). When I say that democracy is a good system of governance, the proposition that I am stating is arguably a proposition made up in large part of the linguistic meanings of the terms "democracy", "good", and so on. The linguistic meanings of these terms are somehow grounded in conventions, which are plausibly grounded ultimately in what various people are inclined to say or think about such terms. When "what is said" is seen in this light, it becomes obvious that we should only expect a very rough match between what is said and what is thought. This holds true of utterances such as "I believe that P" as well: we have no more ability to select words that exactly convey our internal state of mind when we say "I believe that P" than when we merely assert "P" on its own.

In sum, propositional attitudes rely heavily on rough heuristics and must necessarily cut corners due to the nature of public language. This goes a long way toward resisting the difficult cases for a phenomenal theory of cognitive content. This is because our conception of our occurrent thoughts is quite

closely connected with our conception of our standing propositional attitudes: we generally think of our occurrent thoughts as simple manifestations of our standing attitudes. As a result, we can expect the accuracy of ascriptions of occurrent thoughts to be negatively affected by the general inaccuracy of standing attitude ascriptions. This gives us reason to take intuitive ascriptions of occurrent thoughts and related properties (such as rationality) with a grain of salt. In particular, it doesn't seem unreasonable to suggest that the case where we hold Josie to be irrational because she thinks 6-cubes have 32 vertices *might* be one where we are projecting more cognitive content than there really is. Similarly, we *might* be projecting more cognitive content than there really is in Fred's case. We picture him having thoughts whose cognitive contents correspond fairly closely to the meanings of sentences he is disposed to utter, but this might be a gross oversimplification of the reality. If we could look directly into Josie's and Fred's heads, we might see that the contents that they reason on are not the high-level contents that we are inclined to ascribe to them.

In some respects, this picture is revisionary and runs against naïve ascriptions, but common sense also provides support for it. Reflection on cases suggests that we can distinguish between two different ways of thinking a proposition: we can *grasp* it, which puts us in a position to reason competently about it, and we can think it without grasping it, which leaves us unable to reason competently about it. For example, consider the difference between these claims:

SUN The volume of the Sun is $1.412 \times 10^{18} \text{ km}^3$.

SQUARE There is a square.

Intuitively, I have an excellent grasp of SQUARE, but a poor grasp of SUN. One might put the point like this: I *really know* what a square is, but I don't *really know* how big $1.412 \times 10^{18} \text{ km}^3$ is. This intuitive, felt difference is reflected in my reasoning capabilities. Given SQUARE, I can draw various inferences, such as that something has four sides. I can do so independently of the language used to convey SQUARE to me. In contrast, when thinking about SUN (or at least its numerical component), I find myself almost entirely dependent on symbol manipulation: whether or not I can easily infer a given proposition from SUN seems to depend largely on the language in which SUN and the other proposition are expressed, as well as the nature of the rules that I know for manipulating the relevant expressions. For example, I can easily tell that SUN entails SUN/10, but I cannot as easily tell that it entails SUN/8.

SUN/10 If the Sun's volume were ten times smaller, it would be $1.412 \times 10^{17} \text{ km}^3$.

SUN/8 If the Sun's volume were eight times smaller, it would be $1.765 \times 10^{17} \text{ km}^3$.

We can restate the preceding claims in base 8 (rounding the volume of the Sun in base 8, as we did in base 10):

SUN₈ The volume of the Sun is 1.163×10^{24} km³.

SUN/8₈ If the Sun's volume were eight times smaller, it would be 1.163×10^{23} km³.

SUN/10₈ If the Sun's volume were ten times smaller, it would be 7.65×10^{22} km³.

Now the relationship between SUN₈ and SUN/8₈ is obvious, but the relationship between SUN₈ and SUN/10₈ is not.¹⁶ This illustrates the fact that, at least when thinking about large numbers, we seem to reason entirely with numerals, following numeral-handling rules, not with the numbers themselves; otherwise, the base of representation should make no difference to the obviousness of inferences. I am using very large numbers for effect here, but this is true of pretty much all numbers except the smallest natural numbers.¹⁷ Again, this contrasts with SQUARE. In the case of SQUARE, it does not matter what words we use to state the proposition (so long as we understand it). The difference between SUN and SQUARE illustrates the contrast between grasping a content, which allows us to reason about it, and merely grasping symbols that stand for the content, which only allows us to reason using syntactic rules about symbols.¹⁸

¹⁶SUN/10₈ is rounded, so there is strictly speaking no entailment, but the example nevertheless illustrates my point; the relationship would be even less clear if I showed the precise result.

¹⁷I am not sure where to draw the line, which might be subject-dependent, but we plausibly have some kind of grasp of natural numbers up to 5.

¹⁸I discuss a variant on the preceding argument in more detail in (Bourget forthcoming a).

Of course, it is possible to grasp some facts that are related to SUN. For example, there are infographics that compare the relationship between the volume of the Sun and the volume of the Earth to the relationship between the volume of a basketball and the volume of an apple seed. It seems that being presented with such a visual comparison instantly enables one to grasp the relative sizes of the Sun and the Earth. This also gives us some inferential abilities: applying one's knowledge and grasp of facts about basketballs and apple seeds, one can infer, for example, that the diameter of the Sun is at least 20 times the diameter of the earth (you can see this simply by visualizing apple seeds lined in front of a basketball). The analogy, however, does not allow one to grasp the absolute, precise volume of the Sun, which is what is at issue here. We also have a fuzzy sense of large numbers, but this does not confer us a grasp of precise claims such as SUN.¹⁹

Chudnoff (2015, ch. 2) also gives mathematical examples that illustrate the intuitive difference between grasping and not grasping a content. In one example, he notes that the proposition *if $a < 1$, then $2 - 2a > 0$* is a fact that we can simply “see” (grasp). In contrast, the Goldbach conjecture, that *every even integer greater than 2 is the sum of two primes*, seems to elude this sort of immediate intellectual grasping. The first proposition is not just easy to grasp, it is also easy to reason with, and the second is not just hard to grasp, it is also hard to reason with. This further reinforces the idea that not all contents that we can in some sense think are cognitive contents: it seems that

¹⁹See Carey 2009.

a content must be grasped to be cognitive.

Consider also this example.²⁰ Most people know that millions of individuals die of preventable poverty every year. The typical response to such statistics is not very strong. Yet take a child who is dying of simple malnutrition and bring them to the average North American's dwelling, and you can expect an empathic and generous response. As Peter Singer (2009) stresses, we seem to be egregiously inconsistent in our moral responses: we should be far more moved by the fact that *millions* are dying of preventable poverty than by the fact that *one* child is, but many of us won't give *a single dollar* to the millions. A natural explanation of this discrepancy is that we cannot fully cognize, grasp, or have as cognitive content the proposition that millions are dying of poverty, whereas we can grasp the plight of the one child. On this view of the matter, we are not systematically, egregiously inconsistent in our moral responses, or at least not as far as the rational mind is concerned: in cases where we grasp the needs of others, we do act to help, but in cases where we merely grasp words and numerals, we are not moved to help. This makes sense, since, obviously, words and numerals don't need help. To the extent that we merely grasp words and numerals, we are unable to apply our moral sense to the relevant facts.

(I hasten to add that this is no excuse for egoism. The real ethical upshot of these observations is that we ought to ensure that we grasp the facts that are relevant to our decisions, not that we are excused from ethical norms because

²⁰Also discussed in my (forthcoming a).

we don't really know what we're doing. When important decisions need to be made, we ought to look past the veil of symbols. This duty is similar to, and in line with, the duty to try to gather appropriate information when making decisions, as well as the duty to avoid making important decisions when unable to think clearly. While not excusing anyone, the fact that arguably immoral behavior can be explained by a lack of grasping rather than a lack of moral sensitivity suggests a different approach to morality than a philosopher might be inclined to take: we don't need more arguments, we need a better grasp of the stakes. This is not in general something that can be acquired through reasoning alone.)

Examples such as the preceding suggest that not all the contents that we can in some broad sense think are cognitive contents. They also illustrate the fact that the contents that we have most trouble grasping tend to be abstract or complex. In general, the contents that we have trouble grasping and reasoning with seem to be just the contents that we are least likely to be able to phenomenally represent, such as contents involving large magnitudes, large numbers of people, or abstract, non-perceptible properties. As a result, the phenomenal theory's prediction that highly abstract or complex contents cannot be cognitive contents seems quite plausible. It is an initially surprising prediction that arguably confirms the theory.

To return to the examples used earlier, it seems plausible, in light of the above discussion, that typical democracy thoughts and typical thoughts about six-dimensional figures fall short of involving a grasp of democracy or

six-dimensional figures, respectively. In the case of six-dimensional figures, it seems that one can never grasp this. It is almost completely opaque what follows from the claim that something is a 6-cube. One has to reason semi-formally on the basis of what one knows about cubes. In the case of democracy, it might be that one can form a phenomenal representation of democracy with great difficulty, but I would not rest my theory on this. Against the background of the limitations on cognitive content highlighted above, any putative counterexample of this sort can plausibly be seen as a case where there is not just no relevant phenomenal content, but also no relevant cognitive content.

The phenomenal theory does not merely predict the correct limitations on cognitive content, but it sheds new light on these limitations. On the assumption that cognitive content is phenomenal content, it is easy to see how the limitations on cognitive content noted above arise. For example, it is plausible that only small natural numbers can be precisely represented in consciousness. This seems introspectively obvious, and there is also empirical evidence for this claim.²¹ If the phenomenal theory is true, this means that we can only reason on propositions involving small integers. Why is it that the apple seed/basketball model enables one to (approximately) grasp the relation that obtains between the volume of the Sun and the volume of the Earth? The phenomenal theory predicts that this is the case because the model allows

²¹See the discussion in Bourget forthcoming a, which draws on the empirical work discussed in Carey 2009 and Dehaene 1999.

us to visually experience that very relation, or something close to it: while the absolute volumes of the Sun and the Earth are beyond our grasp, their ratio is not so large as to be impossible to perceive in a scale model. Why is it that we behave as we do in the face of poverty statistics? Assuming that cognitive content is phenomenal content, it is easy to explain why we have weak responses to poverty statistics compared to directly perceived suffering: a million of people dying of poverty just cannot enter our consciousness. As a result, when we consider poverty statistics, we only bring into consciousness words and numerals. In contrast, one person's suffering might be able to enter our consciousness, especially with the benefit of direct perception, which sustains the content in consciousness. The mind only has access to cognitive contents (phenomenal contents), so, from its point of view, it is perfectly rational to respond more vigorously to the needs of the one individual than to the needs of the millions reported as general statistics. The phenomenal theory explains this apparently irrational behavior.

Contrast the phenomenal theory's predictions regarding high-level thoughts with those of a view such as phenomenal functionalism, understood as a theory of cognitive content. It is not clear exactly how contents get fixed on the latter view, but the general idea is that complex inferential connections can determine the cognitive contents of thoughts. Presumably, when a subject's thought T is so connected with other thoughts that the subject is (by human standards) excellent at drawing all and only the inferences that would be valid if the content of T were P, then T's content is P. The problem is that we can

be excellent at reasoning with a thought *as if* it enabled us to grasp P, but without grasping P. For example, people that are good with numbers seem to be able to reason very competently with large numbers. Intuitively, they only reason with numerals, but their numeral handling skills are such that, if one can represent P in virtue of implementing the right inferential dispositions, they should represent numbers in virtue of their inferential dispositions. In this way, it seems that phenomenal functionalism predicts too much cognitive content.

In summary, the phenomenal theory, together with the fact that cognitive phenomenology is relatively limited, has the initially surprising implication that many abstract or complex propositions that we are inclined to describe as our mental contents are not cognitive contents. In this section, I have suggested that we should not see this as refuting the theory. On the contrary, it seems that this implication of the phenomenal theory explains otherwise puzzling facts about human reasoning and rationality. Unlike phenomenal functionalism, the phenomenal theory of cognitive content seems to predict just the right limitations on cognitive content. This is a strong point in favor of the phenomenal theory.

4 Beyond cognitive content

The phenomenal theory paints a picture of reasoning on which genuine processing of contents can be difficult to achieve because it depends on

conscious awareness. As I have tried to show, this accords with familiar observations regarding what we can and cannot grasp, and this explains certain striking facts about ourselves. However, this also seems to run counter to the received wisdom that much thinking is unconscious. If cognitive content is content that we reason on and cognitive content is phenomenal content, it might seem that a lot of what we naturally call “reasoning” or “thinking” is not reasoning at all, because it does not involve phenomenal representations of relevant contents. This section develops my answer to this objection.

My answer has two parts corresponding to two different kinds of cases in which we seem to reason on a content P without phenomenally representing P. In the first kind of case, there is some consciousness involved in the alleged reasoning process, but it is not consciousness of P. In the other kind of case, there is no consciousness involved. I will first consider the first kind of case.

In many cases in which we seem to reason about P without consciously grasping P, our thinking involves a kind of internal monologue, that is, a series of experiences in which we hear ourselves speaking certain words. Such internal monologues are a prominent feature of reasoning about abstract or complex matters such as democracy, numbers, etc. As a result, consciously reasoning about democracy plausibly involves entertaining thoughts whose phenomenal contents are uninterpreted or largely uninterpreted symbols. We can refer to such thoughts as *symbolic thoughts*. Note that symbolic thoughts can be involved in reasoning consistently with the phenomenal theory, because they have phenomenal contents just like other cognitive

experiences. Therefore, nothing in the phenomenal theory requires us to say that we don't in fact reason in cases that we are tempted to describe as instances of reasoning about abstract matters such as democracy, so long as such reasoning is reflected in conscious symbolic thoughts. This already goes a long way toward accommodating the received wisdom.

The phenomenal theory can also accommodate the intuition that such instances of reasoning are properly described as instances of reasoning *about* matters such as democracy, at least in some loose sense of "about". It can accommodate this by recognizing that the terms we represent in symbolic thoughts are in some sense about non-linguistic matters such as democracy. For this reason, manipulating these symbols in consciousness is a way of thinking about democracy. What it is not is a way of thinking or reasoning *on* democracy, because we cannot bring democracy in consciousness.

One might want to know what makes an uninterpreted symbol such as "democracy" represented in inner speech be "about" an entity such as democracy, which we cannot phenomenally grasp. There are different possible answers here. In effect, any theory of reference could satisfactorily account for the fact that the term "democracy" used in inner speech by an English speaker is in some sense about democracy. However, one might think that cognitive content ought to, in some sense, fix reference. For if cognitive content is what we reason on but it does not fix what our thoughts are about, then it seems that what our thoughts are about, and by extension, their truth conditions, will potentially come apart from reasoning. What good is reasoning that does

not track truth conditions?

We can give *something like* an account of the reference of the inner speech term “democracy” in terms of cognitive content, but note first that reasoning would track truth conditions even if we could not give such an account. For cognitive contents, even if phenomenal, are bona fide propositions, which either are or determine truth conditions (depending on one’s view on propositions). What is at stake is not whether reasoning tracks truth conditions but whether it tracks the truth conditions specified by contents of high-level thoughts that are too abstract or complex for us to grasp. We already know that reasoning is not causally responsive to such truth conditions (see the preceding section).

As I just said, we can give something like an account of the reference of “democracy” in inner speech in terms of cognitive content. We can do so in terms of what Mendelovici (forthcoming) terms *derived representation*. The general idea is that the term “democracy” used in inner speech *cashes out* into a descriptive content D that rigidly denotes democracy. By “descriptive content”, I mean roughly what a linguistic description of the form “the F” contributes to the proposition expressed by a sentence in which it figures. “Cashing out” is a technical term that we can define as follows: a content C1 *cashes out* into another content C2 (for a given subject in certain circumstances) just in case either the subject is disposed to consciously represent that *by C1, I [the subject] mean C2*, or she is disposed to consciously represent that *by C1, I [the subject] mean C3*, where the parts of C3 cash out into such contents that, when they are brought together in the logical form specified by C3, they

form an expression that is equivalent to C2. For example, we are disposed to recursively cash out the symbolic content “democracy” and the terms in whatever definition we might give of it into further linguistic expressions until we reach non-symbolic contents that we can directly grasp. When a thought’s content cashes out into a content such as *the actual F* (a rigid description) and the actual F is x, it is natural to say that the subject is in some sense thinking about x. This is how thoughts that merely represent the word “democracy” end up being in some sense thoughts about democracy, at least in some loose sense.

We cannot typically cash out all the symbols represented by all our symbolic thoughts at the same time, but we can do it in a piecemeal fashion, and, if we were to combine the different, maximally cashed out contents together, we would have descriptive contents that pick out non-linguistic items. The descriptive contents that pick out non-linguistic items can take two main forms. In many cases, they are mere metalinguistic descriptive contents such as *the thing actually called “democracy” around here*. In other cases, they can be qualitative descriptive contents such as *the actual long orange vegetable*.²²

This view of how symbolic thoughts get to be “about” non-linguistic items borrows from descriptivism and functionalism. However, there is a crucial difference, which is that I don’t claim that this dispositionalist-

²²Jackson (1998a,b) and Chalmers (2002, 2012) explain in detail how broadly-speaking descriptive conditions can account for reference.

descriptivist story is sufficient to account for the reference of “democracy” either in public language or in inner speech. The above is a (simplified) story that reflects how we juggle terms in inner speech, as well as how we end up regarding our thoughts as being about things such as democracy, but I don’t think we can reasonably take this account to characterize an objective, determinate reference for all or most contents that we don’t grasp in consciousness. The problem is that exactly how one is disposed to cash out one’s use of “democracy” is circumstance-dependent, where the relevant circumstances encompass *anything* that might influence what one does, from drugs to one’s level of hydration. In one circumstance, I might cash out this term one way, and in another I might cash it out another way. We might be tempted to regard some circumstances as abnormal and irrelevant, but I don’t think we can draw a principled line between the conditions that are relevant and those that are not.²³ For this reason, I prefer to think of the foregoing not as an account of the unique, determinate meanings and referents of terms in public language or inner speech, but as an account of why we are inclined to *think* that a term such as “democracy” has a well-defined meaning and referent. Put differently, I would agree that “democracy” is about democracy only in some very loose sense that allows that it is also about myriad other things that we might in some circumstance or other pick out using some cashing out of this term.

²³This is at bottom the same as the problem of collateral information for conceptual role semantics (Block 1986) and Kripke’s (1982) plus-quus problem.

While symbolic thoughts fall short of being about anything non-linguistic in any robust sense, and while thinking symbolically about something is very different from grasping it consciously in thought, we should not underestimate the power of symbolic thought. When the system of symbols is well designed, many logical connections between contents can be encoded as syntactic connections between symbols. This is just what we aim to do with logical systems. This is also why the arabic numerals are so well suited to certain forms of reasoning about numbers. Clever usage of symbolic manipulation is a crucial way in which we expand our cognitive abilities beyond our raw rational abilities grounded in consciousness.

In this respect, there is an illuminating analogy between the role of consciousness in the mind and the role of a CPU in a modern computer. To make the analogy vivid, let me briefly outline how a CPU does its work. A CPU “understands” a limited set of machine language instructions, in that it is only responsive to these instructions. These instructions tell the CPU to make simple transformations to the arrays of bytes in its memory. The interesting part is how the CPU is made to emulate an understanding of higher-level instructions such as *print “hello”*. For the CPU, a high-level instruction such as *print “hello”* starts up as a simple array of bytes in memory. The CPU only “sees” the bytes, not their “printy” meaning. However, the CPU has in memory programs (series of machine language instructions) that it can follow to compile (we could also say “cash out”) a series of bytes such as this into a series of machine language instructions. The result of this compiling process

is a different series of bytes that is a series of machine-level instructions that will result in a certain output device printing “hello” when followed. Note that the CPU can also perform operations on the original *print “hello”* bytes without compiling them. In its memory, *print “hello”* is just an array of about 14 bytes, and it can do anything with them. For example, it can copy the bytes around in memory or send them out as-is to a peripheral.

On the picture I am suggesting, consciousness is the CPU of the mind. Our machine language is made up of the different contents that can enter our consciousness, and we are only truly responsive to those contents. However, some of the things that we represent in consciousness (orthographic shapes) have a privileged role because we program ourselves to cash them out into machine-language instructions (further conscious contents). Through a scaffolding process not unlike the scaffolding of programming languages, we can also build complex symbolic concepts out of simpler symbolic concepts, which allows us to simulate a grasp of much more complex languages than we truly understand. Like computers, we also sometimes manipulate high-level instructions without compiling them, which is useful in all kinds of ways, some of which I sketched above. In brief, the phenomenal theory makes sense of the odd, pervasive practice of talking to ourselves: this is a computational technique used to simulate manipulation of contents that we cannot consciously grasp.

So far, I have talked about how the phenomenal theory can accommodate the point that we are able to reason about democracy when we do so

consciously but without consciously representing democracy: these are real episodes of reasoning because we are manipulating contents in consciousness, and they are in some loose sense episodes of reasoning about democracy because the contents that we manipulate are loosely speaking about democracy, but they are not episodes of reasoning *on* democracy. This does not account for episodes of reasoning that seem completely unconscious. For example, much of our intelligence lies in work done seemingly unconsciously by our perceptual systems, which seem to “infer” various features of objects. None of the difficult work of searching for relevant memories, evaluating thoughts for relevance, and figuring out the meanings of words in context involves much consciousness, as far as we can tell. It is also a familiar fact that we can solve problems without consciously thinking about them. To a large extent, we owe our cognitive abilities to apparently completely unconscious processes. One might object that all this unconscious work involves a lot of thinking and reasoning.

This might seem like a serious objection, but I think it merely reflects a terminological discrepancy, because there is nothing substantive to disagree on regarding fully unconscious “reasoning”. First, everyone can agree on what the unconscious processes are like. I do not deny that there is a lot of information processing done completely unconsciously, nor do I deny that it can be helpful to talk as if this processing was just like conscious thinking (plausibly, this is so because it is designed to extend and mimic conscious thinking). Relatedly, I am happy to grant that unconscious processes can be in

many respects similar to conscious processes (but not in all respects, because they are not conscious). Second, everyone agrees that unconscious mechanical processing is probably not truly responsive to contents, because the only available, non-consciousness-involving account of narrow or cognitive content is the short-arm conceptual role theory, which no one believes. These seem to be all the core facts about unconscious processing. The only disagreement that we can have is on whether or not we should refer to such processing as “thinking”. I am not particularly attached to my terminological choice, but it seems reasonable to reserve this term for processes that are in all important ways the same as the paradigm instances of thinking, which we have seen all involve consciousness as CPU. We have also seen that the causal powers of conscious thoughts are quite different from the causal powers of other forms of representation in the brain (witness the poverty statistics case). For these reasons, it seems reasonable to reserve the terms “thinking” and “reasoning” for processes that operate on phenomenal contents.^{24,25}

Empirical evidence regarding unconscious processes is not only consistent with the phenomenal theory, but it also seems to support it. Many authors have argued on empirical grounds that consciousness acts as a global workspace

²⁴Mendelovici (forthcoming) develops a similar response to objections from the unconscious to PIT.

²⁵Smithies’ (2011a, 2011b, 2012a, 2012b and 2014) also argues that unconscious information processing plays no rational role because it is disconnected from consciousness. This view coheres nicely with the picture presented here, though Smithies ultimately defends a different kind of constitutive connection between consciousness and rationality than the one I defend here. One key difference between Smithies’ view and mine is that Smithies seems to have a more liberal view of cognitive phenomenology.

or central information stream that controls high-level decision making and reasoning.²⁶ This is by far the most widely agreed upon way of distinguishing between conscious and unconscious processes in the brain. The phenomenal theory explains why consciousness plays such a central role in the brain: if consciousness is our only means of being truly responsive to intentional contents, it makes sense to harness it as a sort of central processing unit. In this way, empirical evidence regarding the architecture of the mind supports the hypothesis that consciousness is our means of grasping propositions.

5 Conclusion

In this paper, I argued that the phenomenal theory, according to which cognitive content is phenomenal content, offers a promising explanation of cognitive content, the content to which we are rationally responsive. I first noted that the phenomenal theory is well suited to explaining the cognitive contents of perceptual experiences, and that a case can be made that it can account for the cognitive contents of some thoughts. I then considered whether the theory can account for the cognitive contents of high-level thoughts, which seem to pose a major challenge for this theory, and I argued that the phenomenal theory can accommodate high-level thoughts. The phenomenal theory does not only accommodate high-level thoughts, but it also sheds

²⁶For example, Baars (1997), Dehaene & Naccache (2001), Dennett (1991), Tye (1995 and 2000), Dretske (1995). The work of Milner & Goodale (1995), is also relevant as Smithies (2011a, 2014) argues. See the discussion in Bourget forthcoming a, section 4.2, for more details.

new light on the limitations of the rational mind. This theory also makes sense of the apparently central roles of consciousness and inner speech as components of reasoning processes. I did not argue at length that there is cognitive content, but I made the case that the phenomenal theory offers a promising explanation of cognitive content if there is such content. This helps support the claim that there is cognitive content because, as I noted in section 1, one of the main reasons for resisting the intuitive view that there is cognitive content is that such content is in need of an explanation in terms of something else. I argued that consciousness might well be this something else.

References

- Baars, B. J. (1997). *In the Theater of Consciousness: The Workspace of the Mind*. Oxford University Press.
- Block, N. (1986). Advertisement for a semantics for psychology. *Midwest Studies in Philosophy*, 10(1):615–78.
- Bourget, D. (2010). Consciousness is underived intentionality. *Noûs*, 44(1):32–58.
- Bourget, D. (2017). Representationalism and sensory modalities: an argument for intermodal representationalism. *American Philosophical Quarterly*, 53:251–267.

- Bourget, D. (forthcoming a). The role of consciousness in grasping and understanding. *Philosophy and Phenomenological Research*.
- Bourget, D. (forthcoming b). Why are some phenomenal experiences "vivid" and others "faint"? Representationalism, imagery, and cognitive phenomenology. *Australasian Journal of Philosophy*.
- Bourget, D. (forthcoming c). Anomalous panpsychism: A novel approach to the mind-body problem. In *Routledge Handbook of Panpsychism*.
- Bourget, D. and Mendelovici, A. (2016). Phenomenal intentionality. In *Stanford Encyclopedia of Philosophy*.
- Burge, T. (1979). Individualism and the mental. *Midwest Studies in Philosophy*, 4(1):73–122.
- Burge, T. (1986). Individualism and psychology. *Philosophical Review*, 95(January):3–45.
- Carey, S. (2009). Where our number concepts come from. *Journal of Philosophy*, 106(4):220–254.
- Chalmers, D. (2012). *Constructing the World*. Oxford University Press.
- Chalmers, D. J. (2002). On sense and intension. *Philosophical Perspectives*, 16(s16):135–82.

- Chalmers, D. J. (2004). The representational character of experience. In Leiter, B., editor, *The Future for Philosophy*, pages 153–181. Oxford University Press.
- Chudnoff, E. (2011a). The nature of intuitive justification. *Philosophical Studies*, 153(2):313–333.
- Chudnoff, E. (2011b). What intuitions are like. *Philosophy and Phenomenological Research*, 82(3):625–654.
- Chudnoff, E. (2015). *Cognitive Phenomenology*. Routledge.
- Crane, T. (2003). The intentional structure of consciousness. In Smith, Q. and Jokic, A., editors, *Consciousness: New Philosophical Perspectives*, pages 33–56. Oxford University Press.
- Dehaene, S. (1999). *The Number Sense: How the Mind Creates Mathematics*. Oxford University Press.
- Dehaene, S. and Naccache, L. (2001). Towards a Cognitive Neuroscience of Consciousness: Basic Evidence and a Workspace Framework. *Cognition*, 79(1):1–37.
- Dennett, D. C. (1991). *Consciousness Explained*. Penguin.
- Dretske, F. (1995). *Naturalizing the Mind*. MIT Press.
- Farkas, K. (2008a). Phenomenal intentionality without compromise. *The Monist*, 91(2):273–93.

- Farkas, K. (2008b). *The Subject's Point of View*. Oxford University Press.
- Fodor, J. A. (1990). *A Theory of Content and Other Essays*. MIT Press.
- Goff, P. (2012). Does mary know i experience plus rather than quus? a new hard problem. *Philosophical Studies*, 160(2):223–235.
- Horgan, T. and Graham, G. (2012). Phenomenal intentionality and content determinacy. In Schantz, R., editor, *Prospects for Meaning*. De Gruyter.
- Horgan, T. and Tienson, J. (2002a). 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. and Tienson, J. (2002b). 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.
- 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*. Walter De Gruyter.
- Jackson, F. (1998a). *From Metaphysics to Ethics: A Defence of Conceptual Analysis*. Oxford University Press.
- Jackson, F. (1998b). Reference and description revisited. *Philosophical Perspectives*, 12(S12):201–218.

- Kriegel, U. (2003). Is intentionality dependent upon consciousness? *Philosophical Studies*, 116(3):271–307.
- Kriegel, U. (2011a). Cognitive phenomenology as the basis of unconscious content. In Bayne, T. and Montague, M., editors, *Cognitive Phenomenology*, pages 79–102. Oxford University Press.
- Kriegel, U. (2011b). *The Sources of Intentionality*. Oxford University Press.
- Kripke, S. A. (1980). *Naming and Necessity*. Harvard University Press.
- Kripke, S. A. (1982). *Wittgenstein on Rules and Private Language*. Harvard 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.
- Mendelovici, A. (MS). *The Phenomenal Basis of Intentionality*.
- 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. (forthcoming). Consciousness and intentionality. In Kriegel, U., editor, *Oxford Handbook of Consciousness*. Oxford University Press.
- Mendola, J. (2008). *Anti-Externalism*. Oxford University Press.

- Milner, A. D. and Goodale, M. A. (1995). *The Visual Brain in Action*. Oxford University Press.
- Pautz, A. (2008). The interdependence of phenomenology and intentionality. *The Monist*, 91(2):250–272.
- Pautz, A. (2010). An argument for the intentional view of visual experience. In Nanay, B., editor, *Perceiving the World*. Oxford University Press.
- Pautz, A. (2013). Does phenomenology ground mental content? In Kriegel, U., editor, *Phenomenal Intentionality*, pages 194–234. Oxford.
- Pitt, D. (2004). The phenomenology of cognition, or, what is it like to think that *P*? *Philosophy and Phenomenological Research*, 69(1):1–36.
- Pitt, D. (2011). Introspection, phenomenality, and the availability of intentional content. In Bayne, T. and Montague, M., editors, *Cognitive Phenomenology*, page 141. OUP.
- Putnam, H. (1975). The meaning of 'meaning'. *Minnesota Studies in the Philosophy of Science*, 7:131–193.
- Putnam, H. (1981). *Reason, Truth, and History*. Cambridge University Press.
- Russell, B. (1912). On the notion of cause. *Proceedings of the Aristotelian Society*, 7:1–26.
- Salmon, W. (1984). *Scientific Explanation and the Causal Structure of the World*. Princeton University Press.

- Siegel, S. (2005). The contents of perception. In Zalta, E. N., editor, *Stanford Encyclopedia of Philosophy*.
- Siegel, S. (2007). How can we discover the contents of experience? *Southern Journal of Philosophy*, 45(S1):127–42.
- Siegel, S. (2010). *The Contents of Visual Experience*. Oxford.
- Siewert, C. (1998). *The Significance of Consciousness*. Princeton University Press.
- Singer, P. (2009). *The Life You Can Save: Acting Now to Stop World Poverty*. Random House.
- Smithies, D. (2011a). Attention is rational-access consciousness. In Mole, C., Smithies, D., and Wu, W., editors, *Attention: Philosophical and Psychological Essays*, pages 247–273. Oxford University Press.
- Smithies, D. (2011b). What is the role of consciousness in demonstrative thought? *Journal of Philosophy*, 108(1):5–34.
- Smithies, D. (2012a). The mental lives of zombies. *Philosophical Perspectives*, 26(1):343–372.
- Smithies, D. (2012b). A simple theory of introspection. In Smithies, D. and Stoljar, D., editors, *Introspection and Consciousness*. Oxford University Press.

- Smithies, D. (2013a). The nature of cognitive phenomenology. *Philosophy Compass*, 8(8):744–754.
- Smithies, D. (2013b). The significance of cognitive phenomenology. *Philosophy Compass*, 8(8):731–743.
- Smithies, D. (2014). The phenomenal basis of epistemic justification. In Kallestrup, J. and Sprevak, M., editors, *New Waves in Philosophy of Mind*, pages 98–124. Palgrave MacMillan.
- Strawson, G. (1994). *Mental Reality*. MIT Press.
- Strawson, G. (2000). The phenomenology and ontology of the self. In Zahavi, D., editor, *Exploring the Self*, pages 23–39. John Benjamins.
- Strawson, G. (2008). Real intentionality 3: Why intentionality entails consciousness. In Strawson, G., editor, *Synthesis Philosophica*, pages 279–297. Oxford University Press.
- Thompson, B. (2003). *The Nature of Phenomenal Content*. PhD thesis, University of Arizona.
- Tye, M. (1995). *Ten Problems of Consciousness: A Representational Theory of the Phenomenal Mind*. MIT Press.
- Tye, M. (2000). *Consciousness, Color, and Content*. MIT Press.
- Williamson, T. (1998). The broadness of the mental: Some logical considerations. *Philosophical Perspectives*, 12(12):388–410.

Yablo, S. (1996). Wide causation. *Philosophical Perspectives*, 11(11):251–281.

Yablo, S. (2003). Causal relevance. *Philosophical Issues*, 13(1):316–28.