Representational Content and the Objects of Thought

Nicholas Kirk Rimell
Midlothian, Virginia

Bachelor of Arts, Duke University, 2006
Master of Arts, University of Chicago, 2008

A Dissertation presented to the Graduate Faculty
of the University of Virginia in Candidacy for the Degree of
Doctor of Philosophy

Corcoran Department of Philosophy

University of Virginia
May, 2018

for my parents, Pat and Donna Rimell

Acknowledgements

I've benefitted from the generosity of several individuals throughout my work on this project. I am grateful to Jonathan Barker, Elizabeth Barnes, Ross Cameron, Jim Cargile, Jim Darcy, Matt Duncan, David Ingram, Derek Lam, Harold Langsam, John Mahlan, Andrew Morgan, Paul Nedelisky, and Adam Tiller for feedback on chapter drafts and for helpful discussion. I'd also like to thank the UVA Metaphysics Writing Group (some of whose members are among those just listed). I owe special thanks to Andrei Mărăşoiu, Brie Gertler, and Trenton Merricks. In addition to providing detailed feedback on one of my chapters, Andrei helped proofread multiple chapters. Brie provided detailed feedback on several chapters and engaged with me in extensive discussions over them. Trenton, who served as my dissertation advisor, read through drafts of each chapter, offered extensive feedback, and engaged with me in extensive discussions on each chapter and on the project overall.

My defense committee consisted of Ross Cameron, Brie Gertler, Trenton Merricks, and David Vander Meulen. I am grateful to them for their incisive questions and their constructive feedback.

I am grateful to Matt Andler, Galen Barry, Jeff Carroll, Jim Darcy, Ralph DiFranco, Matt Duncan, Megan Duncan, Corin Fox, Derek Lam, Teddy Lam, John Mahlan, Andrei Mărăşoiu, Andrew Morgan, Doug Reed, Stacie Thyrion, Adam Tiller, McCailin Wunder, and several others for their friendship and support as I worked on this project. I am especially grateful to Matthew Adams – for his friendship, his support, and (especially) his patience – as well as to Adam Blincoe – for opening up his home to me for a dissertation retreat, for his friendship, and for his constant encouragement.

Finally, I am grateful to my parents. I don't know what it's like *not* to feel excited to pursue my passions. I don't know what it's like *not* to feel valuable for who I am. For most of my life I've taken these feelings for granted. But surely I wasn't just born feeling that way. My parents made me feel that way. And they did it without my even knowing they were doing it. And they still do. I would not have finished – or even begun – this project if not for their love and support. I dedicate this dissertation to them.

Table of Contents

Introduction

Since Frege, it has been common to think of beliefs and related mental states as having the representational contents they have in virtue of their being attitudes towards propositions. Taken straightforwardly, this view is an ontological one: there are these entities, propositions, to which beliefs and related mental states are attitudes; and it is in virtue of a certain mental state's being an attitude toward a certain proposition that that mental state represents what it does.

My primary goal in what follows is to present a refined version of this ontological claim – which I'll call "the first tenet of the traditional doctrine of belief" – and to identify some important implications of it that stretch well beyond ontology. These include implications (i) for the debate between internalists and externalists about representational mental content, (ii) for the question of how to make sense of essentially egocentric belief, and (iii) for the question of whether thought is an essentially relational activity. My secondary goal in what follows is to provide an initial – albeit, incomplete – defense of the overall picture of mental representation that results from accepting the first tenet of the traditional doctrine of belief.

Chapter One lays the groundwork for my project. My goal in this chapter is to characterize the debate between externalists and internalists about representational mental content. I begin by distinguishing this debate from other, related debates. I then lay out the main argument for externalism about representational mental content, the crucial premise of which is *Content Fixes Truth* (CFT), the claim that, necessarily, if two thoughts share representational content (of some sort), then they have the same truth value. If CFT is true, then internalism is false and externalism is true. I devote the remainder of Chapter One to laying out the main arguments for CFT and to presenting some compelling internalist responses to these arguments.

It is not until Chapter Two that I introduce what I am calling "the first tenet of the traditional doctrine of belief." The traditional doctrine of belief, as I characterize it, consists of two tenets. The first tenet, to put it more precisely, says that, necessarily, if two thoughts share representational content (of some sort), then they are attitudes toward – and, consequently, inherit the representational content of – the same proposition. The second tenet of the traditional doctrine of belief says that, necessarily, propositions (if they exist) have their truth values absolutely. It is commonly recognized that these two tenets jointly entail CFT. I argue that the first tenet – all by itself – entails CFT, from which it follows that the first tenet – all by itself – entails externalism. I then argue that, for similar reasons, the first tenet also entails that our egocentric beliefs are attitudes toward *private* propositions. All this may lead some to deny the first tenet of the traditional doctrine of belief. But I endorse this tenet. Consequently, I endorse CFT and externalism, along with the view that our egocentric beliefs are attitudes toward private propositions. I then defend these positions against some standard objections.

In Chapter Three I argue from CFT to the claim that to think about something is to stand in a certain relation to it. This claim is controversial since, at least *prima facie*, it commits us to denying that we can think about nonexistents (e.g., Santa Claus, the fountain of youth, the event of UVA's winning the men's 2018 NCAA basketball tournament). I accept this commitment. Consequently, I deny that we can think about nonexistents. I conclude Chapter Three by providing the bare bones of an error theory meant to explain what we are doing when we think we're thinking about nonexistents.

Chapter Four centers on a particular argument for necessitism put forward by Timothy Williamson. Necessitists claim that, necessarily, whatever exists exists necessarily. The argument for necessitism that I discuss runs as follows: first, for a singular proposition to exist, whatever

that proposition is directly about must exist as well; second, it follows from this requirement that, necessarily, whatever exists exists necessarily. It turns out that, if my arguments in Chapter Three are sound, then some of the most initially attractive options for resisting this argument for necessitism fail. This may seem to be particularly troubling for me since, as it also turns out, necessitism threatens to undermine the motivation for the error theory I sketch at the end of Chapter Three. But even those who are persuaded by my arguments in Chapter Three should not be persuaded by Williamson's argument. For Williamson's argument is self-undermining. Or so I argue. I then show how we can extend my argument in order to argue that necessitism is false.

One

Narrow Representational Content

The last twenty some years have seen a resurgence of the view that there is narrow mental content, i.e., mental content that is determined wholly by what is internal to the thinker. A notable part of this resurgence has been the development of two related accounts of narrow content, offered by Frank Jackson and by David Chalmers.

I have three goals in this initial, largely expository chapter. The first is to settle on a way of interpreting the question 'Is there narrow mental content?' that allows us to see this question as a substantial and difficult one. To this end, I shall suggest that we interpret this question as a question about the *representational* contents of mental states. My second goal is to present the accounts of narrow mental content offered by Jackson and by Chalmers as, in particular, accounts of narrow representational mental content, and to do so in a way that reveals just how compelling these accounts are. A key to accomplishing this goal will be to illustrate the way in which these accounts provide internalists with the resources to resist two standard arguments for the externalist-friendly claim – which I call "Content Fixes Truth" – that thoughts cannot be alike in (representational) content while differing in truth value. My third goal will be to make clear that proponents of internalism must follow Jackson and Chalmers in rejecting this claim about the relationship between content and truth value. In accomplishing these goals I'll be laying the groundwork for my arguments in future chapters.

1 – Narrow Content and Wide Content: The Initial Picture

Many mental states – if not all mental states – represent. For instance, my belief that snow is white represents. In particular, it represents snow's being white. It is, as such, a *representational state* or, simply, a *representation*.

Anything that is a representational state is also, as such, an intentional state, which is just to say that it is about something. For instance, my belief that snow is white is about snow (and perhaps *whiteness*, and perhaps snow's being white).

Moreover, anything that is a representational state also has a semantic value (and anything that has a semantic value represents). My belief that snow is white is true. My guess that there are exactly nine hundred fifty-seven sheets of paper in the copy machine is, I imagine, false. My desire two Novembers ago for Clinton to win the general election for the U.S. presidency went unsatisfied.

Not everyone agrees on which mental states represent and which – if any – do not. To keep things simple, I'll assume that all mental states represent. (None of my arguments turn on this assumption.) But I'll focus, in what follows, on those mental states that have truth values – i.e., those mental states that are either true or false.[1] I shall call such mental states "thoughts."[2]

Let us say that a representational feature of a mental state (or of a speech act, sentence, picture, etc.) is any feature of that state that plays a role in (or involves) its representing something. Analogously, let us say that an intentional feature of a mental state (or of a speech act, sentence, picture, etc.) is any feature of that state that a plays a role in (or involves) its being about something.

[1] For simplicity's sake, I'll assume that, necessarily, for any thought T, either T is determinately *true* or T is determinately *false* (but it isn't both, and it doesn't have any other truth value). None of my arguments ride on this assumption. (In Chapter Two, there may seem to be an exception. At that point, I'll briefly suspend the assumption.)

[2] Here I follow Chalmers (2002a) but with one important exception. Chalmers uses the term 'thought' to refer, specifically, to mental *tokens* that are true or false (e.g., to belief tokens). To avoid crosstalk with those philosophers who use 'thought' more liberally, I shall *not* at the onset restrict my use of 'thought' (or, for that matter, my use of 'belief', 'desire', 'concept', or any other mental-state term) to simply picking out mental *tokens*. On a couple occasions, I'll need to make explicit exactly how I am using the term 'thought' (or another mental-state term). But, in most cases, either context will make it obvious whether I am using the term to pick out a token or a type, or the ambiguity will be harmless. (I am also going to avoid, for the time being, saying anything committal about how mental state types should be individuated.)

While different philosophers use the term 'content' in subtly different ways, they generally agree on the following: to identify mental states in terms of content is to identify certain representational – and, as such, intentional – features of those mental states, which features play a role in determining those mental states' semantic values.[3]

Narrow mental content – if there is such a thing – is mental content that is fully determined by a thinker's *internal* features, where a feature counts as internal to a thinker just in case it is a *qualitative* feature and is *intrinsic* to the thinker. So, if there is narrow mental content, then – necessarily – the mental states of two thinkers x and y differ with respect to content of this sort only if they differ (in the relevant respects) with respect to what is internal to them. Wide mental content – if there is such a thing – is mental content that is not narrow, i.e., mental content that is not fully determined by what is internal to a thinker. (From now on, when characterizing mental content of some sort or other, I'll leave the qualifier 'if there is such a thing' implicit.)

Content internalism (henceforth, simply "internalism") is the view that there is narrow mental content; content *externalism* (henceforth, simply "externalism") is the view that mental content is, in general, wide. Internalism comes in two varieties. Strong internalists say that all mental content is narrow, whereas dual content theorists say that, in general, for any mental state with content, that mental state both has narrow content and has wide content. Let us say that for two mental states to share content is for these two states to have exactly the same content, with respect to content of at least one sort. For instance, as I shall use the expression 'shares mental content', my belief that grass is green shares content with my (epistemically idiosyncratic) uncle's

[3] Part of what I'll be doing in this chapter is distinguishing between different uses of 'content' and, in light of this, getting clear on how *I* will be using the term 'content' throughout the majority of this dissertation. To this end, to avoid undue wordiness, I'll sometimes use the term 'content', somewhat loosely, to mean something like *features of the sort philosophers often use 'content' to pick out.* In these cases, I'll have to rely on my reader's attention to context.

guess that grass is green, but it does not share content with your belief that broccoli is green. Thus, we may say that internalists (even dual content theorists) hold that, necessarily, if two thinkers are *internal duplicates*, their respective mental states all share content. Externalists deny this claim.

Some examples will help. It's the year 1750, well before the denizens of planet Earth discover that the clear, drinkable liquid filling their rivers and oceans (and so forth) consists of the chemical compound H_2O. But Oscar (who lives on Earth) believes – here's how Oscar would put it – "Water is H_2O." (Who knows why Oscar believes this? Perhaps it came to him in a dream.) Light years away from Earth, there is another planet – call it "Twin Earth" – that, with one exception, is an internal duplicate of Earth. The exception is that the clear, drinkable liquid filling the rivers and oceans (and so forth) of Twin Earth consists of a chemical compound other than H_2O. Call this chemical compound "XYZ." Twin Oscar is Oscar's Twin Earth duplicate.[4] And Twin Oscar believes – here's how Twin Oscar would put it – "Water is H_2O." If there is narrow content (i.e., if internalism is true), then Oscar's and Twin Oscar's respective belief states share content of this sort. On the other hand, if Oscar's and Twin Oscar's respective belief states do not share content, then it must be that their belief states have only wide content. And, in this case, we should say that, in general, there is only wide (mental) content (i.e., externalism is true).

Consider a related example. April is sitting on the edge of Lake Michigan, dangling her feet in the lake. She looks down and thinks to herself – here's how she'd put it – "This is water." Meanwhile, Twin April is sitting on the edge of Twin Lake Michigan, dangling her feet in the lake. She looks down and thinks to herself – here's how she'd put it – "This is water." If there is narrow mental content, April and Twin April's respective belief states share content of this sort. On the

[4] Well, there is the following difference between Oscar and Twin Oscar: Oscar's body is over 50% H_2O (by mass) whereas Twin Oscar's body is over 50% XYZ. I address this difference briefly in §2. But everywhere else, I ignore the difference, both as it applies to Oscar and Twin Oscar and as it applies, *mutatis mutandis*, to other (conscious) Earth-to-Twin-Earth internal duplicates.

other hand, if April and Twin April's respective belief states do not share content, then it must be that their belief states have only wide content. And, in this case, we should say that, in general, there is only wide (mental) content.

Examples like the two above have standardly been used to motivate externalism over internalism.[5] Arguably, two mental states cannot share content while differing in semantic value. But internalism is at odds with this claim. For Oscar's belief is true (since water is H_2O) whereas Twin Oscar's belief is false (since *twin* water – the clear, drinkable liquid filling the rivers and oceans of Twin Earth – is XYZ). Unlike Oscar's and Twin Oscar's respective beliefs, April's and Twin April's respective beliefs both seem to be true. But April's belief and Twin April's belief refer to different things – to water (H_2O) and to twin water (XYZ), respectively. It arguably follows from this that April's and Twin April's belief states do not share content. But, if there is narrow content, Oscar's and Twin Oscar's belief states share such content, as do April's and Twin April's. So, arguably, the two examples we've been discussing are both *counter*examples to internalism.

More generally, let us say that a mental representation M is *Twin-Earth-susceptible* just in case it includes, as a constituent, some concept (a) that serves as a rigid designator and (b) that picks out some entity whose essential characteristics cannot be known simply by reflection on the concept. And let us say that a *Twin Earth case* is any case in which two internal duplicates have corresponding Twin-Earth-susceptible mental states that – in light of some difference in the duplicates' respective environments – differ in semantic value. (So a Twin Earth case needn't

[5] These examples are variations on Hilary Putnam's (1973, 1975) Twin Earth scenario. Putnam initially appeals to these examples to argue for *semantic* externalism, the view that the meanings of linguistic entities are, in general, wide (1973, 1975). Colin McGinn (1977) is the first to observe (in print) that an analogous consideration may be made in favor of content externalism, while Tyler Burge (1979) is the first to offer a sustained defense of content externalism (or, for Burge, "anti-individualism") based on this consideration.

literally involve Twin Earth, *per se*.) The above two cases are Twin Earth cases.[6] And, if the above, externalist-friendly analyses of these two Twin Earth cases are accurate, then – for any Twin-Earth-susceptible mental representation M – we can construct a Twin Earth case in which two mental states, one of which is M, are had by two internal duplicates but nevertheless do not share content.

(This is what I have in mind when I say that, for the externalist, mental content is *in general* wide. Accordingly, when I say that the externalist and internalist disagree over whether there is narrow mental content, what I mean – more precisely – is that they disagree over whether *even* Twin-Earth-susceptible representations have such content. We should all agree that, for instance, the content of my belief that 2 + 2 = 4 is narrow.)

We have just seen what the overall motivation for externalism is supposed to be. Here are three standard – and powerful – considerations in favor of internalism. First, arguably, we need to appeal to narrow mental content to explain the causal roles that mental states play. For instance, we ought to expect Oscar and Twin Oscar to be alike with respect to what their mental states cause them to do, and – arguably – to accommodate this we must insist that there is a kind of content that Oscar's and Twin Oscar's respective mental states have in common. Second, we all have privileged access to the contents of our own respective mental states, and – arguably – the externalist cannot accommodate this fact. For instance, arguably, if the contents of Oscar's and Twin Oscar's respective mental states are wide, then Oscar and Twin Oscar do not have privileged access to their mental states since there is nothing available by introspection to Oscar that is not available to Twin Oscar, and *vice versa*. Third, arguably, externalists cannot account for the fact

[6] Or so I'll take for granted, up until §7. The alternative would be to hold that – at least assuming that there really is a Twin Earth – what is essential to water is not its underlying chemical structure but, instead, its surface structure (i.e., its being clear, drinkable liquid, etc.). This seems to be what Crane (1991) thinks.

that we are, by and large, rational cognitive agents. For instance, arguably, if April is rational (which, let us stipulate, she is) and if the content of her mental states is wide, then *simply* from her belief that – here's how she'd put it – "this is water" she ought to be able rationally to infer that – here's how she'd put it – "this is H_2O"; but, of course, she can't do this.[7]

2 – A Complication

Among the respondents to the 2009 Philpapers survey question 'Mental content: internalism or externalism?', 51.1% of "target faculty" accepted or leaned towards externalism, 20.0% accepted or leaned towards internalism, and 28.9% accepted or leaned towards neither (e.g., 5.7% identified themselves as being "insufficiently familiar with the issue") (2014: 495).[8] This gives us fairly strong inductive evidence for thinking that a majority of professional philosophers in the analytic tradition think that the question 'Is there narrow mental content?' has a straightforward answer. And my characterization of this question in the preceding section also suggests that it can be answered in a straightforward way.

But, on reflection, I think this is mistaken. There is no straightforward way to answer the question 'Is there narrow mental content?'. To begin, the question is ambiguous. For recall from §1 that different philosophers use 'content' in different ways, even within the context of the mental. True, they generally use 'content' to pick out *some* representational – and, as such, intentional –

[7] For classic presentations of these three considerations in favor of internalism, see, respectively, Fodor (1987, 1991), Loar (1988) and Boghossian (1989), and Kripke (1979). (Kripke does not endorse internalism, nor does he present the consideration from rationality as a consideration in favor of internalism, *per se*. Rather, Kripke introduces a puzzle for accommodating the rationality of cognitive agents, and internalism is an essential part of a well-recognized solution to this puzzle. See, e.g., Chalmers 2002a: 608, 628.)

[8] Target faculty are identified, for the purposes of this survey, as "regular faculty members in 99 leading departments of philosophy" (2014: 468). By design, no specification was given in the question as to what exactly was meant by 'externalism' and 'internalism' (other than that the question had to do with mental content), including whether dual content views count as versions of internalism or as alternative views (2014: 470).

feature of mental states, which feature plays a role in determining mental states' semantic values. But philosophers differ on *which* feature of this sort they use 'content' to identify. So 'content' admits of multiple disambiguations. As a result 'Is there narrow content?' admits of multiple disambiguations as well. Moreover – and this is the real issue – it is at the very least reasonable to answer different disambiguations of this question in different ways.[9]

Here is an example, one that centers on two radically different uses of the term 'content'. Michael Tye is a paradigmatic externalist (2000: Ch. 3, 2009: Ch. 8). Given Tye's use of 'content', this should not surprise us. Brian Loar, meanwhile, is a paradigmatic internalist (1988, 2003). Given Loar's use of 'content' (at least in the context we'll be discussing), this should not surprise us either.

Tye characterizes his use of 'content' as follows:

[9] There is an additional source of ambiguity: the term 'narrow'. Again, a representational feature of a mental state is narrow just in case it is determined wholly by what is internal to the thinker. Traditionally, what is *internal* to the thinker is simply what is "in the head" or "within the skin," i.e., within some physical boundary of the thinker (or some part of the thinker). But, as Katalin Farkas (2002) has observed, this conception of what is internal – and, consequently, of what is narrow – will not do. Recall, for instance, that the Twin Earth scenarios are supposed to provide us with paradigmatic examples of cases that, at least according to the externalist, motivate externalism. And they serve as such paradigmatic examples only insofar as Oscar and Twin Oscar (April and Twin April, etc.) are internal duplicates. But Oscar and Twin Oscar differ with respect to what is within their physical boundaries (as do April and Twin April). In particular, Oscar is over 50% H_2O whereas Twin Oscar is over 50% XYZ. (Farkas provides a related example centering on meningitis. See her 2002: 190 - 191.) So we need to appeal to a different understanding of 'internal' – and, consequently, of 'narrow' – if we're to adequately capture what is at stake in the debate between internalists and externalists. And it is not enough just to say, as I have said above, that an internal feature of a thinker is a qualitative feature that is intrinsic to the thinker. For this just pushes the question back to what the relevant sense of 'intrinsic to the thinker' should be. Brie Gertler has argued that "[a]ny way of explicating 'intrinsic to the thinker' will clash with the usual taxonomy of leading externalist and internalist views, or construe these positions as involving claims that are standardly regarded as orthogonal to them – and, in some cases, explicitly rejected by their most prominent exponents" (2012: 51). I agree with Gertler. However, for our purposes, we can just stipulate that Oscar and Twin Oscar, April and Twin April, etc., are paradigmatic examples of intrinsic, qualitative (i.e., internal) duplicates. In doing so, we run the risk of talking past other philosophers only insofar as the arguments we are exploring turn on differences in *exactly* how we understand what counts as being intrinsic to a thinker. But none of the arguments we'll be exploring turn on these differences.

> [T]he term 'content' is a term of art, and there is no one correct way to use it. As I use the term, the content of a thought is what is thought. It is expressed in the 'that'-clause and it is either true or false (or neither true nor false, on three-valued views). (2007: 610, fn. 14)

For Tye it is just a matter of stipulation that the content of a thought has a truth value, independent of a context. So, as Tye uses the term 'content', the possibility of two respective mental states' sharing content yet differing in truth value is just ruled out by fiat. This commits Tye to externalism given that narrow content can be shared by mental states with different truth values.[10]

Meanwhile, Loar characterizes his use of 'content' (or at least *one of* his uses of 'content') as follows:

> The mental or psychological content of a thought is a matter of how it conceives things; and that is what we hope to grasp, at least approximately, when we try to understand another person. We want to know not merely what her thoughts represent as it were impersonally, but also how they represent things *to* her. (2003: 229)

So, for Loar, two mental states are alike in content just in case they conceive things in the same way, i.e., just in case things are represented to the respective thinkers of these thoughts in the same way. Intuitively, even if *what* is being represented to a thinker may be in part determined by factors

[10] Not all philosophers who characterize content in terms of what is thought (or what is represented, what is conceived, etc.) are self-proclaimed externalists. For instance, Katalin Farkas is a self-proclaimed internalist (2008a, 2008b), and according to Farkas "the content of a thought or a judgment is what we *grasp* when we think the thought, or reflect upon a belief" (2008a: 158). But simply to characterize content as what is thought or what we grasp (or what is represented, what is conceived, etc.) is not yet to give a clear enough sense of what content is for us to determine whether there is narrow content.

external to the thinker, *how* something is represented to a thinker is fully determined by a thinker's internal features.[11] This suggests that, as Loar uses 'content', mental content is narrow.[12]

We have just seen an example of the way in which different uses of the term 'content' can reasonably lead one to different answers to the question 'Is there narrow content?'. The practical upshot of this is that any genuinely illuminating investigation of the question 'Is there narrow content?' must begin by our getting clear on exactly what notion of content we wish to employ. This is the task of the next section.

3 – Representational Content

Curtis Brown characterizes mental content as follows:

> *Mental* content simply means the content of a mental state such as a thought, a belief, a desire, a fear, an intention, or a wish. *Content* is a deliberately vague term; it is a rough synonym of another vague term, 'meaning'. A state with content is a state that *represents* some part or aspect of the world; its content is the way it represents the world as being. For example, consider my belief that water is a liquid at room temperature. The content of this belief is what it says about the world, namely that a certain substance, water, has a certain property, being a liquid, under specified conditions, namely being at room temperature. (2016: §1)

And Jesper Kallestrup, in a similar vein, writes the following:

[11] To put the point differently: how a thinker conceives of something is, intuitively, fully determined by a thinker's internal features. Loar says something similar. According to Loar, "[i]t is not unnatural to suppose that conceivings are in the head" (2003: 229). Again, this conception of what is internal as what is literally within certain physical boundaries is problematic. (See fn. 9.) But Loar's basic idea, I take it, is that it is natural to think of how one conceives of something as internal to the thinker, on whatever uses of 'internal' capture the intuitive sense in which, e.g., Oscar and Twin Oscar are internal duplicates.

[12] Interestingly, Tye has a way of resisting the move from Loar's notion of content to internalism, at least given what Loar himself has in mind by *how* a thought conceives things. For Loar (2003), the way a thought conceives things is constituted by its phenomenal character. So, for Loar, mental content supervenes on phenomenal character. But the supervenience of mental content on phenomenal character implies that mental content is narrow only given the assumption that phenomenal character is narrow. And Tye is happy to allow that phenomenal character is wide. In fact, Tye allows this precisely because he accepts a strong version of intentionalism, according to which the phenomenal character of a mental state reduces to its content (in Tye's sense of 'content', as characterized above). See Tye (2000: Ch. 3 and 2009: Ch. 8).

> To see that Jane's jumper is red and to believe that apples are wholesome are both mental states. These two states are also representational. To say that a mental state is representational is to say that it serves the function of being about something in the world, or that it takes the world to be a certain way. Perceptions and beliefs are representational states. To be in such states is to represent the world as being a certain way. The way a state represents the world as being is its representational content. (2011: 1)

Finally, Gilbert Harman says the following about the content of *perceptual* experience:

> Our experience of the world has content – that is, it represents things as being in a certain way. In particular, perceptual experience represents a perceiver as in a particular environment, for example, as facing a tree with brown bark and green leaves fluttering in a slight breeze. (34)

Brown, Kallestrup, and Harman are all talking about mental representations of some kind or other, and each takes it that to characterize such a state in terms of the particular content it has is to identify the way in which that state represents things (or the world) as being (or as being in). This is what Kallestrup refers to as a state's *representational* content.

Let us say, then, that for a certain representational state to have a certain representational content is, simply, for it to *represent things as being a certain way*. Accordingly, let us say that two representational states have the same representational content just in case the way they represent things as being is the same. (Dual content theorists should say that two representational states share representational content of *a certain sort* just in case, for some sort of way of representing things as being, the way they represent things as being is the same. I'll largely ignore this complication except for in §6, when I explicitly discuss Chalmers's version of dual content theory.)

Given how I am characterizing what it is for a representational state to have (and, accordingly, for two representational states to share) a certain representational content, it seems

natural to characterize representational content itself as the way a state represents things as being. I shall often speak this way, but we need to be careful. Say that two mental states M and M* have the same representational content, in the sense that the way these mental states represent things as being is the same. This does not mean that there literally is this particular *entity* – the content, which is a way of representing things as being – that M and M* share. All it means is that they are the same with respect to how they represent things as being. So we can make good sense of a state's having a particular content – and of two states' sharing that content – without reifying content.[13] So, while it will often be helpful (just to avoid undue verbosity) to speak of representational content itself as being the way things are represented as being, we should keep in mind that this is *only* a helpful way of speaking. (Of course, one might have philosophical reasons for thinking that there really are such entities as representational contents. My point is just that there being such entities does not follow, *as a matter of definition*, from the fact that thoughts and other mental states have and share representational contents.)

One way to interpret the question 'Is there narrow mental content?' is to interpret it asking, *Do mental states have narrow representational content?* This is how I shall interpret the question. Interpreting the question this way allows us to make sense of it as being both substantial and difficult, in no small part because the competing considerations for internalism and externalism,

[13] This is true even given the dual content theory (with respect to representational content), but here things get a little trickier. Say that two mental states M** and M*** have the same narrow representational content (hence, share content) yet have distinct wide representational contents. This does not mean that there literally are these three entities – M**'s wide representational content, M***'s wide representational content, and the narrow content shared by M** and M***. Rather, it means that M** and M*** differ with respect to those wide features they have which constitute their representing things as being certain ways yet, at the same time, are the same with respect to the narrow features they have which constitute their representing things as being certain ways.

presented in §1, can all be interpreted as respective competing considerations for claiming – and for denying – that mental states have narrow representational content.[14]

Begin with the standard considerations in favor of internalism, now interpreted as considerations in favor of the view that there is narrow *representational* mental content.

First, arguably, we need to appeal to narrow representational content to explain mental states' causal roles. Intuitively, we should expect two thinkers' total mental states to effect the same behaviors in them only insofar as these total mental states are the same with respect to the ways in which they represent things as being (setting aside those cases in which differences cancel each other out). And, intuitively, to accommodate this we must say that representational mental content is narrow since, e.g., Oscar and Twin Oscar behave in exactly the same way.

Second, each of us has privileged access to the representational contents of our mental states, and – arguably – we need to appeal to narrow representational mental content to accommodate this fact. Each of us is in a unique position to tell, by introspection alone, the ways in which our various mental states represent things as being. But, arguably, if the ways our mental states represent things as being were determined even in part by factors external to us, we would not be in this unique position. For instance, there seems to be nothing available by introspection to Oscar that is not available by introspection to Twin Oscar, and *vice versa*. So, arguably, if the

[14] One may worry that, in interpreting the question 'Is there narrow mental content?' in this way, we eliminate – as part of the debate over this question – the debate over how to individuate concepts in terms of their content. For, while concepts in some sense *represent* things, a given concept does not all by itself represent things *as being a certain way*. I do not share this worry. To begin, it seems to me that – while concepts are constituents of representational mental states (at least in the minimal sense that we employ concepts in having representational mental states) – they are not themselves mental states. But set this aside. The particular way in which a given representational mental state represents things as being is wholly determined by the particular representational features of the concepts that are its constituents, together with the way in which these concepts are arranged. Let 'conceptual content' pick out precisely these representational features of concepts. While I am focusing here on the question of how to individuate mental states in terms of representational content, the answer to this question entails – and is entailed by – the answer to the related question of how to individuate concepts in terms of what I am calling "conceptual content."

ways that Oscar's and Twin Oscar's mental states represent things as being were determined even in part by factors external to them, then neither Oscar nor Twin Oscar would be in a unique position to tell, by introspection alone, the way in which his mental state represents things as being.

Third, arguably, we need to posit narrow representational content to accommodate the fact that we are, by and large, rational cognitive agents. For instance, arguably, if April is rational (which, let us stipulate, she is) and if the ways that her mental states represent things as being are in part determined by factors external to her, then *simply* from her belief that – here's how she'd put it – "this is water," she ought to be able rationally to infer that – here's how she'd put it – "this is H_2O." But, of course, she cannot do this.

So the standard considerations in favor of internalism may be interpreted as considerations in favor of the view that there is narrow *representational* mental content.

The standard consideration in favor of *externalism*, again, centers on the observation that internalism is at odds with the plausible claim that two mental states cannot be alike in content while differing in semantic value. This consideration is especially compelling if we treat externalism and internalism as having to do with representational content.

In the following section, I offer a detailed account of this consideration in favor of externalism, with one caveat: I focus on this consideration as it involves representational *thought* content and *truth* value. I do this, first, for simplicity's sake and, second, because this narrow focus more precisely targets the versions of internalism I will be discussing in subsequent sections.

4 – Truth and (Representational) Content

Content Fixes Truth (CFT) is the following claim:

> **CFT** Necessarily, if two thoughts share representational content (of any sort), then they have the same truth value.

CFT is a disambiguation (and clarification) of the common assertion, in favor of externalism, that thoughts cannot share content while differing in truth value. (See footnote 5.)

CFT is straightforwardly in tension with internalism about representational content. If there is narrow representational content, then there are at least some instances (or possible instances) in which thoughts share content of this sort yet differ in truth value. For instance, if there is narrow representational content, then Oscar's and Twin Oscar's respective "water" beliefs – each of which its thinker would express with the sentence 'Water is H_2O' – share that content even though Oscar's belief is true whereas Twin Oscar's belief is false.

Given this tension between CFT and internalism with respect to representational content, arguments for CFT are, indirectly, arguments against internalism of this sort. In what follows I present – though do not explicitly endorse – two related arguments for CFT.

The first argument for CFT I'll call "the Argument from How Things Are."[15] Let T_1 and T_2 be two arbitrary thoughts. T_1 represents things as being a certain way, and T_2 also represents things as being a certain way. T_1 is true if and only if the way it represents things as being is the way things really are. Likewise, T_2 is true if and only if the way *it* represents things as being is the way things really are. If T_1 and T_2 have the same representational content, then the way T_1 represents things as being is the same as the way T_2 represents things as being. It follows that, if T_1 and T_2 have the same representational content, then they also have the same truth conditions: in particular, each thought is true if and only if things really are the way T_1/T_2 represents them as being. So, if T_1 and T_2 have the same representational content, then they have the same truth value. Now recall that T_1 and T_2 are two arbitrary thoughts, and note that nothing in the argument up to

[15] Jesper Kallestrup (2012: 70 – 71) offers a helpful overview of this argument.

this point relies on anything contingent. We may conclude, then, that – necessarily – if two thoughts have the same representational content then they also have the same truth value.

The second argument for CFT I'll call "the Argument from How Things Could Be."

One ardent defender of this argument is Robert Stalnaker (1981, 1984: Ch. 1, 1989, 1990/1999: 197 - 198, 2003). In the introduction to a 2003 collection of his essays on metaphysics, Stalnaker writes the following:

> The basic assumption that informs all of my attempts to come to terms with metaphysical questions is that what it is to represent the world – to say how things are – is to locate the world in a space of possibilities. One understands what someone else is saying by understanding how that person is distinguishing between the possibilities, as one takes those possibilities to be. (2003: 8)

To put Stalnaker's basic assumption a different way: to represent things as being a certain way is locate the way things are within a space of possibilities, i.e., a space of ways things could be. It is to identify – among the various ways things could be – which of these is a live option for how things actually are. For Stalnaker, this picture of representation applies both to speech and to the mental. Consider, for instance, my belief that Sandy Koufax is left-handed. Among the various ways things could be, some of them are consistent with Koufax's being left-handed, and some are not. My belief that Koufax is left-handed represents things as being a certain way, and one way to characterize the way my belief represents things as being is to say that my belief locates how things are (i.e., locates reality) as being among those ways things could be that are consistent with Koufax's being left-handed. Or so Stalnaker would say.[16]

So Stalnaker endorses the following:

[16] One standard objection to this way of characterizing representation and representational content is that it does not seem fine grained enough. I discuss this objection in Appendix A when I consider one way of improving upon the Argument from How Things Could Be.

> **Representationalism.** Thoughts represent things as being certain ways. And we can identify the way a thought represents things as being by the space of possibilities within which the thought locates reality. To put the point differently: we can identify the representational content of a thought by identifying which possibilities the thought is consistent with and which possibilities the thought rules out.[17]

It is easy to see why, at least *prima facie*, representationalism implies CFT. Assume representationalism, and consider any two thoughts with the same representational content. These two thoughts locate reality within the same space of possibilities. If this space of possibilities includes the way things actually are, then both thoughts are true.[18] Otherwise, neither thought is true. Since nothing just argued relies on anything contingent, we can conclude that, necessarily, if two thoughts have the same representational content, then they have the same truth value.

Or so the argument goes. In fact, things aren't quite so simple. As we'll in §5, there are different ways of filling out the details of the representationalist picture, and not every way of filling out the details results in a version of representationalism according to which CFT is true. To see how *Stalnaker* is led to endorse CFT, we need to look more closely at the version of representationalism he endorses.

First, the space of possibilities within which a given thought locates reality is constituted by *possible worlds*. A possible world is a maximal way things could be, objectively speaking, where 'could' here is meant to indicate metaphysical possibility. Among the possible worlds, there is the maximal way things *are* (the possible world that is *actual*) as well as all of the maximal ways things merely *could have been*. And possible worlds are maximal in the sense that each is a way

[17] A word of caution: the term 'representationalism' is used in other contexts to pick out different views (e.g., the view about experience that, in fn. 12, I refer to as "intentionalism"). My use of the term follows, roughly, Jackson's (2003b) use of the term. (See §5, fn. 22.)

[18] Here and going forward, 'the way things actually are' should be given a *de dicto* reading and, consequently, should *not* be treated as a rigid designator. (This also applies *mutatis mutandis* to other expressions involving 'actual' and its derivatives, unless I note otherwise.)

reality as a whole could be, objectively speaking, down to the smallest detail. So, for any possible world W and for any way things could be (maximal or otherwise), either W involves things' being that way or W involves things' *not* being that way. To say the former is to say that things are that way in W, whereas to say the latter is to say that, in W, things are not that way. So, for instance, it could have been the case that Clinton is president. For any possible world W, either W involves Clinton's being president – which is to say that, in W, Clinton is president – or W involves Clinton's not being president – which is to say that, in W, it's not the case that Clinton is president.[19]

Second, and consequently, Stalnaker conceives of a thought's representational content as a function from possible worlds to truth values. Consider a possible world W, along with a thought T. T represents things as being a certain way. To say that T is true *at* W is to say that, *in* W, things are that way, whereas to say that T is false at W is to say that, in W, things are not that way. Consider, for instance, my politically uninformed neighbor's belief that Clinton is president. For any possible world W, my neighbor's belief is true at W just in case, in W, Clinton is president, whereas my neighbor's belief is false at W just in case, in W, it's not the case that Clinton is president. For Stalnaker, a thought T's representational content is the function *f* from possible worlds to truth values such that, for any possible world W, *f* maps W to the value *true* if and only if T is true at W (and, otherwise, maps W to the value *false*). So, for instance, on Stalnaker's account, the representational content of my neighbor's belief that Clinton is president is the

[19] We could make the notion of a possible world more precise by saying something about what sort of *entity* a possible world is. But different philosophers (including Stalnaker 2011) use 'possible world' to pick out (or at least to attempt to pick out) entities of different sorts, and for our present purposes we shouldn't get bogged down in the competing ontologies of possible worlds. I shall take it, then, that the above, metaphorical picture of what a possible world is will be enough for us to make sense of the Argument from How Things Could Be.

function f_N such that, for any possible world W, f_N(W) = *true* if and only if my neighbor's belief is true at W, i.e., if and only if Clinton is president *in* W (and, otherwise, f_N(W) = *false*).[20]

It's worth noting that we don't need to read anything ontologically deep into this account of representational content. Stalnaker's intention is just to individuate thoughts in terms of representational content by appeal to functions from possible worlds to truth values, as characterized above. The key, for Stalnaker, is that, necessarily, for any thoughts T* and T**, T* and T** have the same representational content just in case, for any possible world W, T* is true at W if and only if T** is true at W. It is in this sense that a thought's representational content may be identified with the space of possibilities within which it locates reality. The relevant space of possibilities, again, will be a collection of possible worlds. And for a thought to locate reality within this space is for it to be true at the possible worlds that collectively constitute it.

With all this in mind, it is easy to see why, for Stalnaker, thoughts cannot differ in truth value without differing in representational content. For say that Stalnaker's version of representationalism is correct. Then, necessarily, for any two thoughts T* and T**, T* and T** have the same representational content if and only if, for any possible world W, T* is true at W if and only if T** is true at W. *A fortiori*, it's necessarily the case that, for any two thoughts T* and T**, T* and T** have the same representational content if and only if, for whichever possible world is actual, either both T* and T** are true at this possible world or neither is.[21] And,

[20] See Stalnaker (1981: 134, 1984: 2, and 1989: 300). (Stalnaker calls these functions "propositions," but I want to reserve the use of the term 'proposition' for a more involved discussion, beginning in Chapter Two.) In his 1976 (79 – 80), Stalnaker identifies a thought's content with the *set* of possible worlds at which the thought is true. But, as will be obvious from the next paragraph, nothing substantial rides on the difference here, at least given our assumption of bivalence.

[21] As implied in fn. 18, 'whichever possible world is actual', 'the possible world that is actual', and related expressions should all be given *de dicto* readings and, consequently, should not be treated as rigid designators. But there is one exception: the expression 'the actual world' should be given a *de re* reading and, consequently, should be treated as a rigid designator.

necessarily, for any thought T, T is true (i.e., T is true *simpliciter*) if and only if T is true at whichever possible world is actual. So, necessarily, for any two thoughts T* and T**, T* and T** have the same representational content only if they have the same truth value.

So, to summarize, the Argument from How Things Could Be amounts to the following. First, thoughts have representational content, and for a given thought T to have such content – i.e., for T to represent things as being a certain way – is for T to locate reality within a certain space of possibilities (representationalism). Second, this characterization of representational content should be cashed out, more precisely, in terms of possible worlds. In particular, we should identify a thought T's representational content with the function *f* such that, for any possible world W, *f*(W) = *true* if and only T is true at W (and *f*(W) = *false* otherwise). Third, it follows that CFT is true.

5 – Jackson on Representational Content

Frank Jackson begins his essay "Narrow Content and Representation – or Twin Earth Revisited" with the following:

> Intentional states represent. Belief represents how we take things to be; desire represents how we would like things to be; and so on. To represent is to make a division among possibilities; it is to divide the possibilities into those which are consistent with how things are being represented to be and those which are not. I will call the possibilities consistent with how some intentional sate represents things to be, its *content*. (2003a: 55)

In other words (and to focus on *thought* content in particular): thoughts represent things as being certain ways, and we can identify the way a thought represents things as being – that is to say, we can identify the *representational content* of the thought – by identifying which possibilities the thought is consistent with and which possibilities the thought rules out.

This picture of thought content should sound familiar. It is the same picture of thought content I've attributed, above, to Stalnaker, under the heading 'representationalism'.[22] However, while Stalnaker believes that representational content is wide, Jackson believes that it is narrow.[23]

In presenting Jackson's view, I'll follow Jackson in focusing on representational *belief* content, in particular. But, as Jackson himself notes (2003a: 55), what he has to say about the contents of beliefs may be extended to the contents of other mental states as well. Also, since Jackson often uses simply the term 'content' for what I am calling *representational content*, I'll often drop the modifier 'representational', though it should be taken to be implicit.

Jackson's view that (representational) belief content is narrow stems from the importance he puts on *egocentric* belief, which he characterizes as follows:

> The belief that there is an apple immediately above one's head is an example of an egocentric belief. Egocentric belief is ubiquitous. When we believe that it is raining, we typically believe that it is raining near where we *ourselves* are. Beliefs about the arrival of planes are typically beliefs about when the tarmac in front of *us* will contain a plane. Perceptual beliefs are typically beliefs about how we ourselves stand with respect to objects of various shapes, locations, colors and so on. (2003a: 58)

[22] Jackson himself, in his 2003b, uses the term 'representationalism' to pick out a very similar view (which he endorses), one that combines representationalism, as I use the term, with a fairly weak variety of physicalism regarding intentionality. Jackson writes: "By representationalist views of belief, I mean views that say that the cases are alike in that in both cases we have a parcel of matter that represents that things are thus and so by virtue of being in one of a certain range of possible states, and that we should understand their representational contents in terms of mappings from the possible states of these parcels of matter into various ways things might be" (2003b: 99). (Jackson also endorses "representationalism" in the sense of *intentionalism* – see, e.g., Jackson 2006. But this is a different – albeit related – view. See fn. 12.)

[23] Jackson is sometimes characterized as endorsing the view that there is both wide mental content and narrow mental content. This is misleading. For Jackson, there is a common interpretation of 'content' on which the sentence 'Mental content is wide' is true, but this is not the interpretation of 'content' according to which 'content' picks out *representational* content. For Jackson, *representational* mental content is always narrow. (See also McGinn 1982 and Loar 1988. Like Jackson, McGinn and Loar both hold that the term 'content' can be used to pick out a narrow feature of our mental states and can also be used to pick out a wide feature of our mental states. But Jackson, McGinn, and Loar are not dual content theorists, in the sense of 'dual content theory' I have laid out in §1.)

By 'egocentric belief' Jackson means *first-personal* belief or, in other words, *de se* belief. To have such a belief is to have a belief about oneself, wherein one thinks about oneself *as* oneself. That is, to use David Lewis's terminology, it is to *self-ascribe* a property (1979: 520). As Jackson observes, egocentric belief is *essentially* egocentric (2003a: 58). Think of John Perry's famous supermarket story (1979: 3). Perry is pushing a shopping cart around a supermarket, following a trail of sugar, in search of the hapless shopper spilling the sugar. He thinks to himself – here's how he'd put it – "The sugar spiller is making a mess." Unbeknownst to him, *he* is spilling the sugar. So he is attributing to himself the feature of making a mess. But it is only later, when he realizes that he *himself* is making a mess, that he *self*-attributes the feature of making a mess. So Perry, in coming to *self*-attribute this feature, has gained knowledge, knowledge that he could not have had simply in virtue of his having attributed the feature of making a mess to some individual, which individual – as it turns out – is Perry.[24]

Most philosophers acknowledge that egocentric belief is essentially egocentric.[25] But not everyone who acknowledges this irreducibility thinks that there is egocentric *content*, i.e., that there is a proprietary sort of representational content that is egocentric. But Jackson thinks that there is, and it is because of this – along with the ubiquity Jackson attributes to egocentric belief – that he takes it that representational belief content is narrow. Jackson writes the following:

> …[V]ery often what is represented is how things are *in relation to* that which is doing the representing. When I believe that there is an apple on my head, I believe something about how the region immediately above *my* head is. When T(win) Jackson, a duplicate from the skin (or brain) in of me, believes that there is an apple

[24] To be clear, this is not to say that Perry's initial belief that the sugar-spiller is making a mess is *not* a *de se* belief. Indeed, plausibly, for Jackson, it is. Plausibly, for Jackson, when Perry attributes the feature of making a mess to the sugar spiller, he is self-attributing a different feature – e.g., the feature of *following the trail of a sugar spiller, who is making a mess*. In self-attributing this feature, Perry is not thereby self-attributing the feature of *making a mess*. He only self-attributes the feature of *making a mess* when he comes to realize the he *himself* is the sugar spiller.

[25] Stalnaker (1981) seems to be an exception. Herman Cappelen and Josh Dever (2013) are also exceptions.

> on his head, he believes something about how the region immediately above *his* head is. In consequence, my belief and his differ in truth conditions: one is true if and only if one of the regions is the relevant, apple-containing way. Likewise, my thoughts and his will differ in reference. But this does not mean that our belief contents differ. We believe alike concerning how things are in the regions immediately above our respective heads – how things would have to be above our two heads for our beliefs to be true is the same for each of us. (2003a: 57 – 58)

According to Jackson, the representational content of his belief and his Twin Earth duplicate's belief is egocentric and, as such, is the same. Each belief is one in which the thinker is representing *himself* as having an apple on his head. And, as such, each belief locates reality within the same space of possibilities – namely, the space constituted by those ways things could be that involve an apple's being on the head of the first-personal subject.[26]

And what Jackson thinks of this particular case he thinks of any case that is at the very least a *candidate* example of a case involving wide content – that is to say, is a genuine Twin Earth case. Consider, for instance, the case involving Oscar's and Twin Oscar's respective "water" beliefs, which each would express with the sentence 'Water is H_2O'. I have been assuming (and I will continue to assume, for the sake of illustration) that this case really is a Twin Earth case. But, for Jackson, it is an open question as to whether a case of this sort is a Twin Earth case (in my sense of 'Twin Earth case', as made explicit in §1). If it is *not*, then Oscar's and Twin Oscar's beliefs do not differ in truth value (or in semantic value of any other sort) and, consequently, are not even candidates for beliefs having wide content. If it *is* a Twin Earth case, then Oscar's and Twin Oscar's beliefs differ in truth value but still have the same representational content. And this

[26] Of course, Jackson would not *express* the content of his belief in the same way that he would express the content of Twin Jackson's belief. Jackson would use 'I' to express the content of his own belief while using something else – e.g., 'he himself' – to express the content of Twin Jackson's belief. But this is just to observe that, often, to report the belief state of another successfully, we must shift the perspective from which that belief state is characterized from that of the believer to that of ourselves. (See Jackson 2003a: 65 as well as Jackson 2003b: 102. See also Chalmers 2002a: 622 – 624 (§8).)

is explained by the fact that the content is egocentric: Oscar's belief is about the watery stuff around *him* (Oscar), and Twin Oscar's belief is about the watery stuff around *him* (Twin Oscar).[27]

Notice that the above-quoted passage doesn't just reveal how Jackson's views on egocentric belief motivate his internalism; it also reveals the role that these views play in Jackson's response to the Argument from How Things Are. Essential to the Argument from How Things Are is the claim that thoughts with the same representational content cannot differ in truth conditions. This is supposed to follow from the claim that, for any thought T, T is true if and only if the way T represents things as being is the way things are (along with the claim – which is just true by definition – that two thoughts have the same representational content if and only if the way they represent things as being is the same). But, for Jackson, this doesn't follow. Indeed, for Jackson, two thoughts can have the same representational content yet still differ in truth conditions, provided that that content is egocentric.

Jackson's views on egocentric content also underlie his response to the Argument from How Things Could Be.

To see how, we need to get clear on the notion of a *centered world*. Recall that a possible world is a maximal way things could be, objectively speaking (where 'could' is meant to indicate metaphysical possibility). A centered world is a maximal way things could be *for* a certain individual, *at* a particular time. Centered worlds are maximal, then, in the sense that each is a way reality as a whole could be, down to the smallest detail, *for* a certain individual, *at* a certain time. A centered world can be modeled as an ordered pair {W, *C*}, where W is a possible world and *C* is a center. Here, *C* is itself an ordered pair – in particular, an ordered pair {x, t}, where x is a possible individual and t is a time. For any possible world W, possible individual x, and time t,

[27] See Jackson (2003a: 62 – 65). And, for a related argument from Jackson in favor of narrow representational belief content, see Jackson (2003b: 103 – 104).

there is the centered world {W, {x, t}}, whose center is {x, t}. And, for any individual y and time t*, this centered world {W, {x, t}} is actual *for* y *at* t* just in case W is actual, x = y, and t = t*. So consider, for some possible world W*, the centered world {W*, {NKR, 9 p.m.}}. (I shall often ignore dates just to keep things simple.) This world is actual *for me now* just in case W* is the actual world, I am NKR, and it is currently 9 p.m. As it happens, I *am* NKR, and it *is* currently 9 p.m. So the centered world {W*, {NKR, 9 p.m.}} is actual for me now if and only if W* is the one maximal way things could be, objectively speaking, that is currently the way things actually are.[28]

Now return to the Argument from How Things Could Be. In this argument, CFT is inferred from a particular way of fleshing out representationalism. On this way of fleshing out representationalism, the representational content of a thought T is the function *f* from possible worlds to truth values such that, for any possible world W, *f*(W) = *true* if and only if T is true at W (and *f*(W) = *false* otherwise). Jackson's response to the Argument from How Things Could Be, as I'll construe it, is twofold. First, not all belief contents can be equated with functions from possible worlds to truth values, for some belief contents – in particular, egocentric belief contents – must be characterized in terms of *centered* worlds. Second, egocentric belief contents, so depicted, can be shared – and, in some cases, are shared – by beliefs that differ in truth value.

So, to begin: for Jackson, egocentric contents are *centered* contents. For Jackson, centered contents are (or at least may be modeled as) sets of centered worlds. In particular, for any belief B whose content is centered, that content is identical to the set of centered worlds such that, for any

[28] The notion of a centered world goes back to Quine (1969), for whom centered worlds are possible worlds centered on a space-time point. Lewis (1979), meanwhile, characterizes centered worlds as "pairs of a world and designated inhabitants thereof" (532). (For Lewis the inclusion of a time is redundant since Lewis is a B-theorist and a perdurantist who, as such, analyzes an individual thinker x's t-located mental states as the mental states of x's t-located temporal part. A-theorists about time should also find the inclusion of a time – or of a space-time point – redundant since, for A-theorists, what time it is *now* is an objective matter.)

centered world W, B is true at W's center, at the time at W's center. To put it more carefully: for any belief B, B's content – if it is centered – is the set of centered worlds at which B is true, where – for any belief B, possible world W, possible individual x, and time t – B is true at the centered world {W, {x, t}} if and only if, at W, B is true for x at t.

Jackson's depiction of egocentric content as centered is a natural upshot of his commitment to representationalism. For assume that there really is egocentric content, had by egocentric beliefs. Then we cannot adequately characterize the space of possibilities within which an egocentric belief locates reality wholly in terms of different ways things could be objectively. Instead, we must characterize such a space of possibilities in terms of different ways things could be for a subject, at a time. Consequently, we cannot equate the contents of egocentric beliefs with sets of possible worlds (or with functions from possible worlds to truth values). The natural alternative is to equate the contents of egocentric beliefs with sets of *centered* worlds (or, for that matter, with functions from centered worlds to truth values). For instance, consider Jackson's belief that – as Jackson would put it – "There is an apple on my head." If this belief is essentially egocentric, then its content cannot be characterized as, e.g., the set of possible worlds at which Jackson has an apple on his head (or, for Stalnaker, the corresponding function from possible worlds to truth values). The natural alternative is to characterize the content of Jackson's belief as a set of centered worlds: in particular, the set $\boldsymbol{S}_{\text{J}}$ of centered worlds such that – for any possible world W, any possible individual x, and any time t – the centered world {W, {x, t}} is a member of $\boldsymbol{S}_{\text{J}}$ if and only if W is a maximal way things could be that includes, among other things, x's having an apple on his or her (or their) head at t.[29]

[29] For Jackson's argument that egocentric contents are centered, see Jackson (2003a: 58 and 2004: 259). As far as I can tell, Jackson's decision to depict centered contents as *sets* of centered worlds as opposed to *functions* from centered worlds to truth values is not based on anything philosophically substantial. Again, I take it that Jackson, like Stalnaker, is not attempting to make an *ontological* claim about the nature of

If the contents of egocentric beliefs are themselves egocentric and – as such – centered, then these contents can be shared – and, in some cases, are shared – by beliefs with differing truth values. Jackson writes the following:

> The content of Jackson's and TJackson's beliefs that there is an apple [on] their respective heads is the same; it is the same set of centered worlds, the worlds with apple-[on] centers. But, as the center Jackson is actually at differs from the center TJackson is actually at, the conditions under which their beliefs are true differ. In the same way, the references of Jackson's and TJackson's beliefs may differ consistently with their having the same content in cases where we are dealing with centered content. This means that when differences in reference and truth are due solely to differences in centers, they are consistent with sameness of content. When we are dealing with centered content, truth and reference *simpliciter* are not truth and reference at the actual world; they are truth and reference at the actual world and the actual center, that is, at the world and center of the believer (or speaker). (2003a: 58)

So, to put the point a little more precisely: Jackson's and Twin Jackson's respective apple-on-head beliefs have, as content, the set $\boldsymbol{S}_{\text{J}}$ of centered worlds such that – for any possible world W, any possible individual x, and any time t – {W, {x, t}} is a member of $\boldsymbol{S}_{\text{J}}$ if and only if W is a maximal way things could be that includes, among other things, x's having an apple on his or her (or their) head at t. Now, for the sake of demonstration, say that Jackson and Twin Jackson both have their respective apple-on-head beliefs at 3 p.m., and say that at 3 p.m. Jackson really does have an apple on his head whereas Twin Jackson does not. For any belief B, had at a time t by a thinker x, B is true (i.e., is true *simpliciter*) if and only if, for whichever possible world W is actual, B is true at {W, {x, t}}. Let "W_{α}" be the actual world, i.e., the possible world that is actual. Then Jackson's 3 p.m. belief that there is (currently) an apple on his head is true just in case it is true at {W_{α}, {Jackson, 3 p.m.}}, whereas Twin Jackson's 3 p.m. belief that there is (currently) an apple on *his*

representational content. Jackson's point, rather, is that we can individuate the content of egocentric beliefs by appeal to sets of centered worlds, though (I assume) functions from centered worlds to truth values would do just as well.

head is true just in case it is true at {W_α, {Twin Jackson, 3 p.m.}}. So Jackson's 3 p.m. belief that there is (currently) an apple on his head is true just in case, in W_α, Jackson has an apple on his head at 3 p.m., whereas *Twin* Jackson's 3 pm. belief that there is (currently) an apple on *his* head is true just in case, in W_α, *Twin* Jackson has an apple on his head at 3 p.m. In W_α, at 3 p.m., Jackson has an apple on his head, but Twin Jackson does not. So Jackson's belief is true whereas Twin Jackson's is not, in spite of the fact that (on Jackson's version of representationalism) these two beliefs have the same content.

In this section I have discussed, primarily, Frank Jackson's conception of representational *belief* content as narrow. But, again, what I have said here of Jackson's conception of representational belief content is meant to apply – as Jackson himself notes (2003: 55a) – to representational *mental* content, in general.

So, to summarize: in spite of – or perhaps even *because* of – his commitment to representationalism, Jackson holds that mental representational content is narrow; and underlying Jackson's motivation for this view – as well as his responses both to the Argument from How Things Are and to the Argument from How Things Could Be – is Jackson's belief that there is genuinely *egocentric* content and that such content is ubiquitous in our mental lives.

6 – Chalmers on Subjunctive Content and Epistemic Content

David Chalmers offers something of a compromise between Stalnaker's and Jackson's respective accounts of mental content, though as we'll see there are ways in which Chalmers's account extends beyond both of these. According to Chalmers, a thought has (at least) two contents. It has a subjunctive content, which is identical to its subjunctive intension and is, in general, wide. And

it has an epistemic content, which is identical to its epistemic intension and is narrow. Chalmers, then, is a dual content theorist.[30]

While Chalmers has developed his dual content theory over the span of two decades and several publications (e.g., over his 1996, 2002a, 2002b, 2003, 2006, and 2011a), I focus here on Chalmers's 2002a. Here, Chalmers offers an account of the contents of *thoughts* – where by 'thoughts' Chalmers has in mind *token* mental attitudes with truth values (609) – and of the *concepts* that constitute such thoughts. I shall focus on what Chalmers has to say about the contents of thoughts. (See footnote 14 for why.) And for this section I shall follow Chalmers in using 'thoughts' to pick out thought *tokens*, in particular. I shall also follow Chalmers in simply using the term 'content', without any attempt to disambiguate this term; however, in the process of presenting Chalmers's characterization of subjunctive and epistemic thought contents, I shall demonstrate that such contents count as *representational* contents and that, beyond this, they conform to the representationalist picture of content that I have attributed to Stalnaker and to Jackson.

6.1 – Subjunctive Content as Wide Content

Up until now I've been skating over a distinction that is critical to Chalmers's characterization of thought content.

[30] He is also a *two dimensionalist* since, as we'll see shortly, he characterizes his dual content theory by appeal to two dimensional semantics. Two dimensional semantics (i.e., two dimensionalism) is just a formal framework for capturing the semantic features of representational states according to two separate dimensions – each which may be characterized as a mapping from possibilities (e.g., possible worlds), contexts, or some combination thereof to semantic values. Other versions of two dimensionalism are offered, for instance, by Robert Stalnaker (1978), Gareth Evans (1979), Martin Davies and Lloyd Humberstone (1980), David Kaplan (1989), and Frank Jackson (1998). For helpful overviews of the similarities and differences among the varieties of two dimensionalism, see Chalmers (2006), Brogaard (2012), and Schroeter (2012). (Note that not all varieties of two dimensionalism are meant to capture dual content theories of content or of meaning.)

To consider a possible world W as *counterfactual* is to grant that things are the way they actually are but then to acknowledge that things could have been as they are in W. So to say that things are thus and so in a possible world W, considered as counterfactual, is to say that – had W been actual – things would have been thus and so. Consequently, a thought T is true at a possible world W, considered as counterfactual, if and only if whichever way T represents things as being *would* have been the way things are, *had* W been actual.[31] To say that a thought T is true at a possible world W, considered as counterfactual, is – in Chalmers's terminology – to say that W *satisfies* T.

For Chalmers, to consider a possible world W as *actual* is to think of the way things are *qualitatively* in W as a way things may actually be, i.e., as a way things might be, for all we can know *a priori*. Accordingly, to say that things are thus and so in a possible world W, considered as actual, is to say, roughly, that – assuming that reality is qualitatively just as it is in W – things are thus and so. More precisely, for Chalmers, it is to say that – from the hypothesis that reality is qualitatively just like it is in W – it ought rationally to be concluded (which, for Chalmers, amounts to its being *a priori* knowable) that things are thus and so. Consequently, a thought T is true at a possible world W, considered as actual, if and only if – for whichever way T represents things being – it ought rationally to be concluded that things are that way *from the hypothesis* that reality

[31] Of course, among the possible worlds, there is the actual world, W_α. We can consider W_α as counterfactual, along with the other possible worlds. It's just that, in this case, it turns out (whether or not we know it) that what would have been the case, "had" W been actual, is actually the case. This way of speaking sounds a little odd. To make sense of it we need to stipulate – against convention – that expressions in the subjunctive mood *needn't* imply that something is *not* the case, and we also need to stretch our use of the term 'counterfactual' accordingly. (Having just acknowledged this complication and explained how we can deal with it, I will – for ease of exposition – ignore it in what follows.)

is qualitatively just like it is in W. To say that a thought T is true at a possible world W, considered as actual, is – in Chalmers's terminology – to say that W *verifies* T.[32]

When we evaluate a thought T with respect to a possible world, how we evaluate that thought can depend on whether we are considering that possible world as counterfactual or as actual. Take, for instance, the possible world that differs from the actual world precisely in that the clear, drinkable liquid in the rivers and oceans is composed of XYZ, not H_2O. Call this world the "XYZ world." Now say that I think to myself – here's how I'd put it – "Water is XYZ." To determine whether my thought is true at the XYZ world, considered as *counterfactual*, we first grant that things are the way they actually are, and then we ask, *Had the XYZ world been actual, would water have been XYZ*? Given how things actually are, water is H_2O, not XYZ. And, given that water is actually H_2O, water is *essentially* H_2O.[33] So, given how things actually are, water would not have been XYZ, even if the XYZ world had been actual. So my thought that water is XYZ is false at the XYZ world (and at any other world, for that matter), considered as counterfactual. On the other hand, this same thought is true at the XYZ world, considered as actual. For suppose that things actually are qualitatively just like they are in the XYZ world. In this case, the clear, drinkable liquid in the rivers and oceans – that is, the stuff we call "water" – is XYZ. So, from the hypothesis that things are qualitatively just like they are in the XYZ world, it ought rationally to be concluded that water is XYZ. So my thought that water is XYZ is true at the XYZ world, considered as actual.

[32] See Chalmers (2002a: 609 – 610) for this account of the *considering as counterfactual - considering as actual* distinction. As Chalmers notes, this account is based on – but not identical to – the account developed by Davies and Humberstone (1980). That account, in turn, is based (in part) on a distinction Evans (1979) makes between *superficial* and *deep* necessity.

[33] Here, for the sake of demonstration, I follow Kripke (1980) in taking it that the concept *water* is a natural kind concept that, as such, rigidly designates H_2O.

For Chalmers, the subjunctive content of a thought T is identical to T's subjunctive intension, the function *f* from possible worlds to truth values such that, for any possible world W, *f*(W) = *true* if and only if T is true at W *considered as counterfactual* (and *f*(W) = *false* otherwise).[34]

Note, then, that what Chalmers calls "subjunctive content" is precisely the representational content that Stalnaker attributes to thoughts. Chalmers simply makes explicit something that we ought, charitably, to take to be implicit in Stalnaker's account of content – namely, that a thought's truth value is to be evaluated with respect to a possible world considered as counterfactual. (We ought, charitably, to read Stalnaker this way since the default way of considering a possible world is as counterfactual, given that – at least setting aside the actual world – the possible worlds just are the maximal ways things could have been, objectively speaking.)

So characterized, subjunctive content both is representational content and, beyond this, conforms to the representationalist picture of mental content endorsed by Stalnaker and Jackson. When we ask what the subjunctive content of a thought T is, the answer tells us, for every possible world W, whether the way T represents things as being conforms to how things would have been, had W been actual. So subjunctive content is representational content of a certain sort, and for a thought T to have representational content of this sort is for it to locate reality within a certain space of possibilities – in particular, one that consists of possible worlds. Chalmers (2006: 103) calls the totality of possible worlds – i.e., the *overall* space of possibilities within which a thought, in virtue of its subjunctive content, may locate reality – the *subjunctive* space of possibilities.

Subjunctive content is, in general, wide. Say that Cordelia thinks to herself – here's how she'd put it – "There is water." And say that Cordelia's Twin Earth duplicate, Twin Cordelia, thinks to herself – here's how *she'd* put it – "There is water." The subjunctive intension of

[34] See Chalmers (2002a: 614 – 616 (§5)) for this characterization of subjunctive content.

Cordelia's belief is identical to the function f_C such that, for any possible world W, $f_C(W)$ = *true* if and only if Cordelia's belief is true at W, considered as counterfactual. Likewise, the subjunctive intension of Twin Cordelia's belief is identical to the function f_{TC} such that, for any possible world W, $f_{TC}(W)$= *true* if and only if Twin Cordelia's belief is true at W, considered as counterfactual. Let "the H_2O world" be an arbitrary possible world meeting the following criterion: had the H_2O world been actual, there would have been a liquid composed of H_2O but no liquid composed of XYZ. Cordelia's belief is true at the H_2O world, considered as counterfactual, since – if the H_2O world had been actual – there would have been water (H_2O). Meanwhile, Twin Cordelia's belief is false at the H_2O world, considered as counterfactual since, if the H_20 world had been actual, there would not have been *twin* water (XYZ). So, while the subjunctive intension of Cordelia's belief maps the H_2O world to the value *true*, the subjunctive intension of Twin Cordelia's belief maps the H_2O world to the value *false*. So Cordelia's and Twin Cordelia's respective beliefs have different subjunctive intensions. Consequently, they have different subjunctive contents. But, again, Cordelia and Twin Cordelia are internal duplicates. So the subjunctive contents of Cordelia and Twin Cordelia's beliefs are wide. And what goes for Cordelia's and Twin Cordelia's beliefs goes for any Twin-Earth-susceptible thought. So subjunctive content is, in general, wide.[35]

6.2 Epistemic Content as Narrow Content

So far, we've been talking about the way in which a thought might locate reality within a space of *metaphysical* possibilities. Chalmers characterizes *epistemic* possibility (and necessity) as follows:

> Let us say that a thought is *epistemically necessary* when it can be justified a priori: that is, when there is a possible reasoning process that conclusively justifies the thought with justification independent of experience. A thought is *epistemically possible* (in a broad sense, related to but distinct from the usual philosophical sense)

[35] See Chalmers (2002a: 616). Chalmers puts the point by saying that "subjunctive content is *often* wide" (616, my emphasis).

> when the thought cannot be ruled out by a priori reasoning: that is, when its negation is not epistemically necessary. Intuitively, this holds when the thought does not involve an a priori contradiction. More precisely, this holds when there is no possible reasoning process that can conclusively justify the thought's negation, with justification independent of experience. (2002a: 609-610)

So, for instance, my belief that there is no largest prime is epistemically necessary, whereas my uncle's belief that there is a largest prime is epistemically impossible. On the other hand, my belief that the clear, drinkable liquid in the rivers and oceans is H_2O, not XYZ, is epistemically possible but not epistemically necessary.

Chalmers's notion of epistemic content centers on his notion of a *scenario*, which he characterizes in terms of epistemic possibility. Intuitively, "[a] scenario can be thought of as a maximally specific epistemic possibility: one with all the details filled in" (2002a: 610). A bit more formally, a scenario is simply an epistemically possible *centered* world (2002a: 611).[36]

Just as we can consider possible worlds as actual, we can consider centered worlds – consequently, scenarios – as actual. For one to consider a centered world as actual is, more carefully, for one to consider that centered world as actual for herself (or himself, etc.) at the current time. Take the centered world {W, {x, t}}, for some possible world W, individual x, and time t. For me to consider {W, {x, t}} as actual is for me to consider – as a genuine epistemic possibility (i.e., as a way things may be, for all I can know *a priori*) – the hypothesis that I am x, that t is the current time, and that things are qualitatively just like they are in W. Accordingly, things are thus and so in a centered world W, considered as actual by (and *for*) me now, just in case I ought rationally to conclude that things are thus and so from the hypothesis that I am x, that

[36] Chalmers (2002a) defines a centered world as "a world marked with an individual and a time at its 'center'" (611), which is roughly how I characterize centered worlds in §5. One might insist that a centered world can be epistemically possible only from a perspective. But, for Chalmers, while it isn't wrong to say of a given centered world that it is possible from a certain perspective, it is unnecessary to say this. For what is *a priori* knowable is a wholly objective matter.

t is the current time, and that things are qualitatively just like they are in W. Consequently, for any thought T, T is true at the centered world {W, {x, t}}, considered as actual, if and only if, for whichever way T represents things as being, T's thinker ought rationally to conclude (at the time at which T occurs) that things really are that way from the hypothesis that she is x, that t is the current time, and that things are qualitatively just like they are in W. (So, for any thought T, had by a thinker x at a time t, T is true at a given centered world, considered as actual, just in case T is true at that centered world, considered as actual *for* x *at* time t.) To say that a thought T is true at this centered world {W, {x, t}}, considered as actual, is – in Chalmers's terminology – to say that {W, {x, t}} *verifies* T.[37]

For Chalmers, the epistemic content of a thought T is identical to T's epistemic intension, the function *f* from scenarios to truth values such that, for any scenario S, *f*(S) = *true* if and only if T is true at S, considered as actual (and *f*(S) = *false* otherwise).[38] Or, to put the point in Chalmers's terms: the epistemic content of a thought T is identical to T's epistemic intension, the function *f* from scenarios to truth values such that, for any scenario S, *f*(S) = *true* if and only if S *verifies* T (and *f*(S) = *false* otherwise).

So characterized, epistemic content both is a kind of representational content and, beyond this, conforms to the representationalist picture of mental content we have been discussing. Consider a thought T, had by a thinker x at a time t. When we ask what the epistemic content of T is, the answer tells us, for every scenario S, whether the way T represents things as being is part of how things are for x at t, on the supposition that W is actual for x at t. So epistemic content is representational content of a certain sort, and for a thought T to have representational content of this sort is for it to locate reality within a certain *epistemic* space of possibilities.

[37] See Chalmers (2002a: 609 – 611).
[38] See Chalmers (2002a: 611).

Chalmers puts the point in a slightly different way:

> [T]he space of scenarios [that verify my beliefs, at the current time] constitutes my *epistemic space*: the space of specific epistemic possibilities that are open to me a priori. If I had no empirical beliefs, all of epistemic space would be open to me. As I acquire empirical beliefs, my epistemic space is narrowed down. Any given belief will typically *divide* epistemic space into those epistemic possibilities that it endorses and those that it excludes. (610)

So the scenarios – as maximal epistemic possibilities – constitute epistemic space *overall*. And to identify a thought in terms of its epistemic content is to identify a certain portion of the overall epistemic space of possibilities – in particular, the portion constituted by those scenarios that verify that thought. Moreover, for Chalmers, one's *own* epistemic space, given the beliefs one has at a given time, may be identified with precisely those scenarios that *verify* one's beliefs at that time.

Epistemic content is always narrow. Consider, for instance, Oscar's and Twin Oscar's respective "water" beliefs. Again, these are the beliefs that – here's how each would put it – "water is H_2O." Now take any scenario S*, identical to an epistemically possible centered world {W*, {x*, t*}}, for some possible world W*, possible individual x*, and time t*. S* verifies Oscar's belief if and only if Oscar ought rationally to conclude – from the hypothesis that he is x*, that t* is the current time, and that reality is qualitatively just like it is in W* – that (currently) the clear, drinkable liquid in the rivers and oceans around him is H_2O. Likewise, S* verifies *Twin* Oscar's belief if and only if *Twin* Oscar ought rationally to conclude – from the hypothesis that *he* is x*, that t* is the current time, and that reality is qualitatively just like it is in W* – that (currently) the clear, drinkable liquid in the rivers and oceans around *him* is H_2O. In both cases, Oscar and Twin Oscar ought rationally to draw these respective conclusions from these respective hypotheses if and only if *anyone* ought rationally to conclude, from the hypothesis that things are qualitatively just like they are in W*, that (at t*) the clear, drinkable liquid in the rivers and oceans around x*

is H_2O. So S* verifies Oscar's belief if and only if it verifies Twin Oscar's belief. So, to generalize (since S* is just any scenario whatsoever): the scenarios that verify Oscar's belief are precisely the scenarios that verify Twin Oscar's belief. So these two beliefs have the same epistemic intension. Consequently, they have the same epistemic content.

Roughly the same story can be told for any Twin Earth case, involving any pair of Twin-Earth-susceptible thought tokens. In any such case, internal duplicates have the same epistemic content. So epistemic content is always narrow.

Since epistemic content is always narrow, it should come as no surprise that Chalmers's reasons for holding that thoughts have epistemic content are among the reasons for holding that thoughts have narrow content. For one, according to Chalmers, "epistemic content reflects the rational relations between thoughts" in that the inferences it is rational for an individual to make, given her current belief state, are determined by the epistemic contents of her beliefs (2002a: 618). (Indeed, for Chalmers, this connection to rationality is guaranteed by the precise way in which he conceives of what it is for a world to verify a thought.) In addition, for Chalmers, it is a thought's epistemic content – not its subjunctive content – that accounts for the causal relations between one's thoughts and one's behavior (2002a: 620).[39]

Note further that epistemic content is often – though not always – egocentric. For instance, the epistemic content of Oscar's and Twin Oscar's "water" beliefs is egocentric. For, as we just saw, each belief represents the clear, drinkable liquid in the rivers and oceans around *its subject* as being H_2O, whether that subject is Oscar or, instead, is Twin Oscar. (Recall from §5 that Jackson is agnostic about whether such 'water" thoughts are really egocentric. Meanwhile, Chalmers

[39] In Chapter Two, §2, I discuss a related reason that Chalmers has for positing epistemic content: epistemic content is needed to account for the "*cognitive* relations between thoughts" (2002a: 619).

2002a: 611 takes it that they are.) On the other hand, the epistemic content of my belief that 2 + 2 = 4 is not egocentric, but it is narrow.

So some epistemic content is egocentric, and some isn't. All epistemic content is narrow.

6.3 Truth Conditions and CFT

Recall that Stalnaker accepts – and Jackson rejects – the following:

> **Content Fixes Truth (CFT)** Necessarily, if two thoughts share representational content (of any sort), then they have the same truth value.

While Chalmers would accept something *related* to CFT, he rejects CFT itself. Again, Chalmers holds that thoughts have subjunctive content, and – as my discussion of Stalnaker's argument in §4 reveals – thoughts cannot share subjunctive content while differing in truth value. But of course for Chalmers subjunctive content is not the only kind of representational content that thoughts have. They also have epistemic content. And it is possible for two thoughts to share epistemic content while differing in truth value. For instance, as we saw in §6.2, Oscar's "water" belief has the same epistemic content as Twin Oscar's "water" belief. But Oscar's "water" belief is true whereas Twin Oscar's "water" belief is false.[40] So, for Chalmers, CFT is false since there is a kind of content – epistemic content – that thoughts can share while differing in truth value.

[40] Moreover, we can demonstrate this difference in truth value by appeal to the depiction of these beliefs' epistemic intension (i.e., epistemic content) presented, on behalf of Chalmers, above. For any thought T, had by a thinker x at a time t, T is true (i.e., is true *simpliciter*) iff it is verified by the scenario that is actual for x at t. (For non-indexical thoughts, the 'for x at t' is unnecessary but does no harm.) Let's assume that Oscar's and Twin Oscar's beliefs both occur at 1 p.m. So Oscar's belief is true iff it is verified by the scenario that is actual for Oscar at 1 p.m. This is the scenario that is identical to the epistemically possible centered world $\{W_\alpha, \{\text{Oscar, 1 p.m.}\}\}$, where W_α is the actual world. This scenario verifies Oscar's belief iff, at 1 p.m., Oscar ought rationally to conclude – from the hypothesis that this scenario is (currently) actual for him – that the clear, drinkable liquid in the rivers and oceans around him is H_2O. Oscar *should* rationally conclude this from the hypothesis in question at 1 p.m. (and at any other time, for that matter). So the scenario that is actual for Oscar at the time of his belief verifies Oscar's belief. So Oscar's belief is true. Now consider *Twin* Oscar's belief. Twin Oscar's belief is true iff it is verified by the scenario that is actual for Twin Oscar at 1 p.m. This is the scenario that is identical to the epistemically possible centered world

Since Chalmers rejects CFT, he owes us an explanation of where the two arguments for CFT we've been discussing go wrong.

It is straightforward where, for Chalmers, the Argument from How Things Could Be goes wrong. Again, essential to this argument is an inference to CFT from the version of representationalism endorsed by Stalnaker. As we saw in §6.1, this version of representationalism amounts to the claim that representational thought content is subjunctive content. For Chalmers, however, this claim is false. Chalmers holds that thoughts have both subjunctive content and epistemic content, and subjunctive content and epistemic content both qualify as representational content.

Turn now to Argument from How Things Are. Essential to this argument is an inference to CFT from the claim that thoughts cannot have the same (representational) content while differing in truth conditions. As we saw in §5, Jackson's way of resisting this argument is to say that thoughts *can* have the same (representational) content while differing in truth conditions.

For Chalmers, things are a little more complicated. Consider the following passage from Chalmers:

> It is widely held that because narrow content is internal, it lacks the sort of relation to the external world that is required to qualify as *content*. For example, many have thought that narrow content is not the sort of thing that can be true or false, as the Twin Earth cases show us that truth-conditions are not determined internally.
>
> I think that these problems are illusory, and that there is a robust and natural notion of narrow content such that narrow content has truth-conditions of its own. (2002a: 608)

$\{W_\alpha, \{\text{Twin Oscar, 1 p.m.}\}\}$. This scenario verifies Twin Oscar's belief iff, at 1 p.m., Twin Oscar ought rationally to conclude – from the hypothesis that this scenario is (currently) actual for *him* – that the clear, drinkable liquid in the rivers and oceans around *him* is H_2O. Oscar *should not* rationally conclude this from the hypothesis in question at 1 p.m. (nor at any other time, for that matter). So the scenario that is actual for Twin Oscar at 1 p.m. fails to verify Twin Oscar's belief. So Twin Oscar's belief is false.

It will no doubt be obvious to the reader, at this point, what these truth conditions are:

> The epistemic content of a thought delivers conditions that one's actual centered world must satisfy in order for one's thought to be true. We might think of these as a thought's *epistemic* truth conditions, as opposed to a thought's *subjunctive* truth-conditions, which govern truth across counterfactual worlds. (2002a: 618)

In short, for Chalmers, a thought has epistemic truth conditions, which are given by its epistemic intension, and has subjunctive truth conditions, which are given by its subjunctive intension.

This suggests the following explanation of where, for Chalmers, the Argument from How Things Are goes wrong. First, since for Chalmers there are truth conditions of different sorts, the assertion that thoughts cannot share representational content while having different truth conditions needs to be clarified. For this assertion to be interpreted in a way that makes it true, we must interpret it as the claim that thoughts cannot share representational content *of any sort* while having different truth conditions of the corresponding sort. But, so interpreted, this assertion implies CFT only on the supposition that – for truth conditions *of any sort* – thoughts cannot share such truth conditions while differing in truth value. For Chalmers, this supposition is false. Consider, for instance, Oscar's and Twin Oscar's "water" beliefs. These beliefs have the same epistemic content/intension, so they have the same epistemic truth conditions. But they differ in truth value. So thoughts (or at least some thoughts) can have the same epistemic truth conditions while differing in truth value. So, for Chalmers, the Argument from How Things Are fails.

7 – Internalism and CFT

There is a difference of scope between the views discussed over the past two sections. While Jackson and Chalmers both have a lot to say about "the nature of mental and linguistic content," broadly construed, I've focused on what each has to say about the nature of *mental*

representational content, in particular. What Jackson has to say here he often puts in terms of representational *belief* content, though he takes it to apply, more broadly, to representational mental content in general. On the other hand, Chalmers focuses on the contents of thoughts, where a thought is a mental representation with a truth value.[41]

I mention this difference so that I can set it aside. Going forward, for simplicity's sake, I'm going to take for granted that what goes for representational *thought* content goes for representational *mental* content in general, at least with regards to how that content should be individuated. Likewise, I'm going to take for granted that – insofar as a given philosopher X has a view on how to individuate the representational contents of thoughts – we have (defeasible) justification for attributing to philosopher X the analogous view of representational mental content, in general. But those readers who find what I'm taking for granted problematic can read me as just making claims about representational thought content (and about what others have to say about such content).

The above-mentioned difference in scope is not the only difference between Jackson's and Chalmers's respective accounts of representational mental content. Most obviously, Jackson is a strong internalist (with respect to *representational* mental content) whereas Chalmers is a dual content theorist. Jackson and Chalmers also conceive of narrow mental content somewhat differently. For instance, Jackson is happy simply to say that representational mental content – which, for Jackson, is narrow – consists of the sets of centered worlds at which thoughts are true, period. Chalmers, meanwhile, is explicit that content of the sort he takes to be narrow consists of functions from *epistemically possible* centered worlds (i.e., scenarios) to thoughts' truth values at

[41] Again, Chalmers's 2002a is about the contents of thoughts *and of concepts*, but – as noted in §3, fn. 14 – only content of the former sort plausibly counts as representational content.

those centered worlds, *considered as actual*.[42] (Recall also from §6.3 that Chalmers differs from Jackson with regards to where, for him, the Argument from How Things Are goes wrong. But, as I'll discuss in Appendix A, this difference does not stem from any philosophically substantial difference in Jackson's and Chalmers's accounts of representational mental content).

Even still, there is much in common between Jackson and Chalmers. Both posit representational mental content of a sort (a) that is narrow, (b) that can, as such, be shared by thoughts (or by other representational mental states) with different truth values (or with different semantic values of some other sort), and (c) that conforms to representationalism. And both, in positing representational mental content of this sort, provide us with versions of internalism that, first, rule out CFT and, second, come equipped with the resources to resist the two arguments for CFT we've considered in this chapter.[43]

---

[42] On the other hand, the way Jackson conceives of the (representational) contents of sentences is nearly identical to the way Chalmers conceives of the (representational) contents of sentences, which is roughly analogous to how Chalmers conceives of (representational) thought content. Jackson holds that a sentence (at least provided that it includes indexicals, proper names, or natural kind terms) has two contents. One of these is identical to the sentence's *C*-intension and is, as such, wide; and the other is identical to the sentence's *A*-intension and is, as such, narrow. Jackson characterizes such intensions as follows:

> The *C*-intension of '*a* is *F*' is the set of [possible] worlds where the object *a* is *F*. The *A*-intension of '*a* is *F*' is the set of [possible] worlds *w* (or centered worlds) such that 'The actual so and so is *F*' is true at *w* (or at <*c*, *w*>) under the supposition that *w* (or <*c*, *w*>) is actual; that is, the set of worlds (or centered worlds) where 'The so and so is *F*' is true. It is the *A*-intension and not the *C*-intension that gives the content of '*a* is *F*' in the sense of the content of the belief about how things are which we give voice to using the sentence '*a* is *F*'. (2003a: 68)

(Note the connection here between the content of a belief and the A-intension of a sentence used to express that belief. To be honest, I am not sure why, in light of this, Jackson does not just say that the representational contents of beliefs – and of mental representations, more generally – *are* A-intensions. But it seems important for Jackson that they are not. For more on this, see Jackson 2001: 662 and 2003a: 70, endnote 22, which endnote appends the above-quoted passage.)

[43] Even externalists about mental content should can draw lessons from Jackson and Chalmers regarding what these two arguments from CFT turn on. In Appendix A, I appeal to these lessons in order to develop clearer and more easily defensible versions of these arguments. While these arguments do not play a central role in my own eventual defense of CFT, they should be of independent interest to externalists and internalists about representational mental content.

For what follows, it will be important to be absolutely clear that the internalist *must* follow Jackson and Chalmers in rejecting CFT. That is, it will be important to be absolutely clear that internalism (about representational mental content) is false if CFT is true.

To begin, if CFT is true and there is at least one case in which internal duplicates have thoughts that differ in truth value, then internalism (about representational mental content) is false. For internalism is true just in case thoughts – along with any other mental states with representational content – have *narrow* representational content. And, if so, then – by definition – the thoughts of internal duplicates share such content. So, if internalism is true, then CFT is false provided that there are (or at least could be) cases in which internal duplicates have thoughts that differ in truth value.

And there certainly are such cases (either in the actual world or, at least, in nearby possible worlds), each of which is a Twin Earth case of some variety. Now, to be fair, there is room for disagreement about *how many* of these cases there are. For instance, while many philosophers will be happy to say that Oscar's and Twin Oscar's "water" beliefs differ in truth value, some philosophers – e.g., Tim Crane (1991) – will deny that they differ in truth value. (Recall also from §5 that Jackson is agnostic about whether such "water" beliefs differ in truth value.) Nevertheless, there are other cases involving internal duplicates where there is just no question that they have corresponding thoughts with different truth values. A good example of this is the case of Jackson's and Twin Jackson's respective apple-on-head beliefs. (See §5.) And here's a related example: I am currently thinking to myself – here's how I'll put it – "I am over 50% filled with H_2O"; meanwhile, my Twin Earth duplicate is thinking to himself – here's how he will put it – "I am over

50% filled with H_2O." (Note that the thoughts involved in Jackson's example and the thoughts involved in my example are all obviously egocentric.)[44]

So, if CFT is true, then internalism (about representational mental content) is false. And, if so, then externalism (about representational mental content) is true.

Conclusion

I've done three things in this chapter. First, I've proposed that we interpret the question 'Is there narrow mental content?' as the question, *Do mental states have narrow representational content?,* and I've tried to show what is fruitful about interpreting the question this way. Second, I've presented both Jackson's and Chalmers's versions of internalism (a) as versions of internalism about *representational* mental content, in particular, and (b) as providing the internalist with plausible ways of rejecting CFT (along with two standard arguments for CFT). Third, I've shown that internalists about representational mental content must follow Jackson and Chalmers in rejecting CFT.

In the following chapter, I shall – among other things – present and defend an additional argument for CFT and, consequently, for externalism. I shall seek to show that this argument is sound, at least given the truth of a plausible claim about how it is that beliefs – and thoughts, more generally – inherit their representational contents.

[44] Crane (1991) holds that the following two claims are consistent: (a) psychological states cannot differ in type unless the internal properties of their thinkers differ; (b) psychological states cannot be of the same type while differing in truth value. (An implication of this is that, for Crane, there are no Twin-Earth-susceptible thoughts – hence, no genuine Twin Earth cases.) My argument here shows that Crane is mistaken, at least assuming – which I think is reasonable to assume – that Crane is individuating psychological states in terms of their representational contents. (Crane actually recants his 1991 position in his 2001: 166, endnote 30. And in doing so he cites McCulloch 1992 as having demonstrated why he is mistaken.)

Going forward, it will be helpful for me to take some liberties for the sake of simplicity. I'll use 'content' to mean, more specifically, *representational mental content*, unless I specify otherwise. Likewise, I'll use 'internalism' and 'externalism' to mean, respectively, internalism and externalism about representational mental content. (I'll use 'content', 'internalism', and 'externalism' in this way even when characterizing the views of other philosophers, unless I am providing a direct quotation or unless I specify otherwise.[45]) In addition, as noted in §7, I'll take for granted that what goes for representational *thought* content goes for representational *mental* content in general, at least with regards to how that content should be individuated. And I'll treat beliefs as paradigmatic instances of thoughts.

[45] Of course, I will need to specify otherwise when addressing philosophers whose views about content cannot reasonably be interpreted as views about representational content.

Two

Narrow Content and Propositions

I hope to have shown – over the previous chapter – just how compelling a view internalism about representational mental content really is. *Prima facie,* we need to appeal to narrow representational content (a) to explain the causal roles that our mental states play, (b) to accommodate the fact that we have privileged access to the contents of our mental states, and (c) to countenance the rational relations between our mental states. Moreover, the standard objection to internalism is not obviously damning. The standard objection is that internalism is inconsistent with what I am calling "Content Fixes Truth (CFT)" – the claim that thoughts cannot share representational content while differing in truth value. As we saw in Chapter One, both Jackson's and Chalmers's versions of internalism have the resources to resist the standard arguments for CFT.

Nevertheless, given a plausible – and common – assumption about how thoughts derive their contents, CFT is true and, consequently, internalism is false. Roughly, the assumption is that thoughts derive their contents from *propositions*. More precisely, the assumption is the following: necessarily, if two thoughts share representational content (of some sort), then they inherit that content from a single proposition. My primary goal in this chapter is to argue from this assumption to the truth of CFT and, consequently, to the falsity of internalism. A secondary goal is to draw out a further consequence of this assumption – namely, that there are private propositions.

1 – The Traditional Doctrine of Belief

Gottlob Frege, in "On Sense and Reference," writes the following:

> By a thought I understand not the subjective performance of thinking but its objective content, which is capable of being the common property of several thinkers. (1892: 214, fn. 5)

And, in "The Thought: A Logical Inquiry," Frege writes:

> Without wishing to give a definition, I call a thought something for which the question of truth arises. So I ascribe what is false to a thought just as much as what is true. (1956: 292)

Frege is positing a certain class of entities, which he calls "thoughts" ("Gedanken") and which nowadays philosophers often call "propositions." These entities are not mental states, so – in my sense of the term 'thought' – these entities are not thoughts. That is, they are not mental states with truth values. Instead, they are these mental states' "objective contents" ("objektiven Inhalte"). And, like these mental states, they have truth values. And they have their truth values *absolutely*, rather than merely relative to contexts.[1]

For another way of putting what is ostensibly the same view, consider John Perry's presentation of what he calls "a traditional way of thinking of belief," and which Perry explicitly attributes to Frege (1979: 5-6). (Perry himself doesn't endorse this view in its entirety.) Perry characterizes the view in terms of three tenets, *two* of which he characterizes as follows:

> The first [tenet] is that belief is a relation between a subject and an object, the latter being denoted, in a canonical belief report, by a that-clause. So "Carter beliefs [sic] that Atlanta is the capital of Georgia" reports that a certain relation, *believing*, obtains between Carter and a certain object – at least in a suitably wide sense of object – *that Atlanta is the capital of Georgia*. These objects are called *propositions.* (6)
>
> The second [tenet] is that [propositions] have a truth-value in an absolute sense, as opposed to merely being true for a person or at a time. (6)[2]

[1] Frege's use of the term 'objective' (well, 'objektiven') in the (1892) passage suggests this latter position. Perry (1977: 476) reads Frege (1956) as endorsing this position in characterizing propositions (a.k.a., Frege's "thoughts") as the things "for which the question of truth arises." More telling is that Frege accepts so-called "private propositions," and without the acceptance of this position the acceptance of private propositions is unmotivated. (See my §6.1, as well as Perry 1977.)

[2] The third tenet is that, necessarily, for any proposition P and any proposition Q, P's having the same truth value as Q is *necessary* but not *sufficient* for P's being identical to Q, where different versions of this third

Robert Stalnaker, in 1981, calls this view a "received doctrine" about belief (129). And he summarizes it as consisting of the following two tenets: first, that "[b]elief is a relation between an animate subject and an abstract object which we will call, without prejudging its nature, a proposition" and, second, that "[p]ropositions have truth values, and their truth values do not vary with time, place, or person" (130).

And here's the way of cashing out the view that I prefer:

(I) Necessarily, if two thoughts share (representational) content (of some sort), then they inherit that content from a single proposition.

(II) Necessarily, propositions (if they exist) have truth values, and they have their truth values absolutely.

It'll make things easier down the road to say a little up front about what exactly (I) and (II) amount to. To begin, while there is no single, universally-accepted *definition* of the term 'proposition', philosophers (in the analytic tradition) generally agree that propositions (i) are the entities toward which our thought tokens are *attitudes* and (ii) are bearers of truth and falsity.[3] There's no consensus on what exactly this first feature amounts to, except to say that it is at least *part* of what the first tenet of the Fregean view of belief is meant to capture. (Of course, as we've just seen, there are different ways of fleshing out this first tenet.) So let's just stipulate that propositions – if they exist – are entities of the sort that meet criteria (i) and (ii), and let's leave some wiggle room regarding how exactly to cash this out. Building anything more into what is meant by 'proposition' would, in the present context, be question begging.

---

tenet involve different specifications of what, in addition to truth value, is relevant to the individuation of propositions (1979: 6).

[3] See, e.g., John Perry (1979: 5 - 6), Robert Stalnaker (1981: 129 - 130), Timothy Williamson (2002: 235), Trenton Merricks (2015: 21), and Matthew McGrath (2014). (David Lewis is an exception. See my fn. 10.)

With this understanding of 'proposition' in mind, consider claim (I). To say that an entity x inherits content from an entity y is just to say that this entity x inherits y's content, where:

> x inherits y's (representational) content $=_{df}$ x has the (representational) content that y has, and x has this content in virtue of (***i***) y's having it and (***ii***) x's standing in a certain relation to y.

So (I) amounts to the claim that, for any two thoughts T* and T** with the same (representational) content (of a certain sort), there is a unique proposition P that has precisely that content, and T* and T** both inherit that content from P. [4] This implies that, for any set S of thoughts that share a (representational) content (of a certain sort), there is a proposition P, also with that particular content, such that each member x of S has that particular content in virtue of (***i****) P's having that content and (***ii****) T's standing in a certain relation to P. Notice that, based on the above definition, (I) does not imply that there is a *single* relation Φ such that, for any thought T, regardless of its content, T stands in Φ to a certain proposition. But, in what follows, to keep things simple, I'll assume that there is a single relation of this sort. For lack of a better phrase, I'll call this relation the "being an attitude towards" relation, with the qualification that there are related, alternative uses of this phrase that are just as legitimate. (Note that this is consistent with there being other relations in virtue of which representations besides thoughts – sentences, perhaps – inherit the contents of propositions. See §5.1.)

[4] So I take the expression 'x has the (representational) content that y has' to convey that x has whatever (representational) content y has. I do not take it to convey (or to rule out) that the *only* (representational) content that x has is what y has. So it is perfectly coherent to say, for instance, that an entity x inherits (representational) content of some sort (e.g., subjunctive content) from an entity y while inheriting (representational) content of some *other* sort (e.g., epistemic content) from some entity z other than y. For this reason, there is no *obvious* inconsistency between (I) and dual content theory. To avoid wordiness, I will often ignore the possibility that a single thought token may have representational contents of different sorts, but I will not rely on such an assumption in the arguments I defend in this chapter.

An example may help. My (token) belief that snow is white, your (token) belief that snow is white, and my epistemically idiosyncratic uncle's (token) guess that snow is white all have the same content. In particular, they all represent snow's being white. Given (I), these thoughts are all attitudes toward the proposition that represents snow's being white. Let's call this proposition "the proposition <that snow is white>."[5] It is in virtue of our three thoughts' all being attitudes towards <that snow is white> – along with this proposition's representing snow as being white – that these thoughts all represent snow as being white. And it is in this sense that these thoughts may all be said to inherit the content of the proposition <that snow is white>.

Turn now to claim (II), the claim that – necessarily – propositions (if they exist) have truth values, and have their truth values absolutely.

A representation R is *relativized*, let's say, just in case it has truth values relative to – and *only* relative to – contexts of evaluation. For our purposes, we can treat contexts of evaluation as subject-time ordered pairs, relative to which the truth values of a representation may be evaluated – where to evaluate a representation R relative to a subject-time ordered pair {x, t} is simply to determine whether R is true or false *for* x *at* t. Plausibly, then, the class of contexts of evaluation is identical to the class of contexts of *belief*, where a context of belief, let's say, is an individual-time ordered pair, for some (possible) individual x and time t at which a belief may occur. Nevertheless, *to be* a context of evaluation is not *to be* a context of belief. To be a context of evaluation is to be an ordered pair (of a certain sort) relative to which a representation has a truth value; to be a context of belief (or the context of *a* belief) is to be an ordered pair (of a certain sort) at which a belief token may occur.[6]

---

[5] Here (and throughout this dissertation) I follow the tradition of naming a given proposition by the 'that'-clause that most perspicuously captures the way it represents things as being.

[6] My use of 'relativized', as in 'relativized representation', is based on John Perry's use of 'relativized proposition' (1979: 12 – 15). Here, Perry also distinguishes between the context in which a belief occurs

A representation R has a certain truth value *absolutely* just in case R has a certain truth value and it has that truth value independently of any context of evaluation. This is not to say that such a representation cannot also have truth values relative to contexts of evaluation. For instance, my belief that there is no set of all sets is, in a trivial sense, true at, e.g., {Michelle Obama, noon}. But it is also true *simpliciter*. My belief, as such, has its truth value absolutely.

So, by definition, if a given representation has its truth value absolutely, it is not relativized. With this in mind, here is another way to put (II): necessarily, there are no relativized propositions, nor are there any propositions that simply lack truth values.

Going forward, by 'the traditional doctrine of belief' I shall mean, in particular, the conjunction of the two tenets (I) and (II). But I want to acknowledge that there are alternative renditions of Frege's doctrine – for instance, Perry's rendition, as presented above – that are just as worthy of this title.[7]

---

and the context relative to which we might evaluate a belief (or, more generally, a representation). It's worth distinguishing *contexts* of evaluation from *circumstances* of evaluations. For Kaplan, (possible) circumstances of evaluation are "actual and counterfactual situations with respect to which it is appropriate to ask for the extensions of a given possible well-formed expression" (1989: 502). More broadly, let's say, a circumstance of evaluation is just any entity (a) relative to which we might evaluate a given representation (or a given constituent of a given representation, such as a word or concept) as having a given extension and (b) that includes (or depicts) a way things could be. For instance, both possible worlds and centered worlds are circumstances of evaluation, but contexts of evaluation are not circumstances of evaluation since contexts of evaluation do not meet criterion (b).

[7] On the other hand, Stalnaker's rendition, it seems to me, fails to capture the spirit of Frege's doctrine of belief. My complaint is with the second tenet of Stalnaker's "received doctrine." Again, this is the claim that propositions have truth values and that their truth values do not vary with time, place, or person. Given the A-theory of time, the second tenet of Stalnaker's "received doctrine" comes apart both from (II) and from Perry's second Fregean tenet. For instance, given the A-theory, it is an objective, subject-independent fact that, while Obama *was* the president of the U.S, he is not *currently* the president of the U.S. So the A-theorist should say that the proposition <that Obama is currently president of the U.S.> is *absolutely* false, and not just false relative to certain contexts of evaluation. (The A-theorist should also say that this proposition *was* absolutely true.) This analysis is wholly consistent with (II) and with Perry's second Fregean tenet, but it is inconsistent with the second tenet of Stalnaker's "received doctrine." To my mind, this difference is a reason to think that Stalnaker's second tenet fails to capture the spirit of the second tenet of Frege's doctrine of belief.

2 – CFT and the Traditional Doctrine of Belief

Recall, from Chapter One, Jackson's and Twin Jackson's respective apple-on-head beliefs. These are the beliefs that each would express by uttering the sentence 'There is an apple on my head'. Each of these beliefs occurs at 3 p.m. (or so we're stipulating), at which time Jackson has an apple on his head but Twin Jackson does not. So Jackson's belief is true whereas Twin Jackson's belief is false. According to Jackson (2003a: 57 – 58), however, these beliefs have the same (representational) content: each represents its first-personal subject as having an apple on his head.

If Jackson's and Twin Jackson's beliefs really do have the same content, then their respective truth values are not fully determined simply by that content, along with what is objectively the case. Rather, these beliefs' respective truth values are determined by that content, along with what is objectively the case *and* along with the respective contexts in which these beliefs occur. For instance, Jackson's belief is true since (i) it represents its first-personal subject as having an apple on his head, (ii) it occurs at 3 p.m. and is had by Jackson, and (iii) at 3 p.m. Jackson has an apple on his head. Meanwhile, Twin Jackson's belief is false since (i) it represents its first-personal subject as having an apple on his head, (ii*) it occurs at 3 p.m. and is had by Twin Jackson, and (iii*) at 3 p.m. Twin Jackson does not have an apple on his head.

Now assume that (I) is true. That is, assume that, necessarily, if two thoughts share (representational) content (of some sort), then they inherit that content from a single proposition. And assume, also, that Jackson's and Twin Jackson's apple-on-head beliefs share content. Then Jackson's and Twin Jackson's beliefs are attitudes toward the same proposition – call it "the proposition <that there is an apple on my head>" – which proposition represents its first-personal subject as having an apple on his (or her, or their) head. The truth value of this proposition cannot be fully determined simply by this proposition's content, along with what is objectively the case.

A context is needed. But this proposition has no context – at least, it has no *single* context at which it occurs (or is entertained, or whatever). It does not have a unique thinker (or a thinker at all), nor does it occur at a particular time. So the proposition <that there is an apple on my head> does not have a particular truth value *absolutely*. It merely is true relative to some contexts of evaluation – e.g., to {Jackson, 3 p.m.} – and false relative to other contexts of evaluation – e.g., to {*Twin* Jackson, 3 p.m.}. It is, as such, a relativized proposition.

If (I) is true, then, necessarily, if two thoughts share content yet differ in truth value, then these thoughts are attitudes towards a particular relativized proposition, the same in both cases. We've just seen an example of this: if (I) is true, then – necessarily – if Jackson's and Twin Jackson's apple-on-head beliefs share content yet differ in truth value, then these two beliefs are both attitudes towards the proposition <that there is an apple on my head>.

I shall now demonstrate the more general point. Assume that (I) is true. And consider, for some possible world W, two thoughts that – in W – share content even though one and only one of these thoughts is true.[8] Let 'T_T' be a name for whichever of these thoughts is true, and let 'T_F' be a name for whichever of these thoughts is false. (If you think that a representation can fail to be true without being false, replace 'is false', throughout this paragraph with 'is not true'.) Given (I), there is a proposition P such that T_T and T_F are both attitudes toward P and, in light of this, both inherit P's content. Neither T_T nor T_F nor P has a truth value that can be fully determined by its content (which is the same for all three representations), along with what is the case. In addition, for any of these three representations to have a truth value, a unique context of belief (or, more generally, a unique context at which the representation occurs, is entertained, etc.) is needed. But

[8] Throughout this chapter, unless I say otherwise, when I speak of something's being the case in a given possible or centered world, I have in mind its being the case in that world, considered as *counterfactual*. (See Chapter One, §6.) I make this stipulation just to keep things simple. And, where relevant, I demonstrate, in footnotes, that my argument does not turn on this stipulation.

the proposition P has no such context. So P has no truth value. That is, it has no truth value independently of a context of evaluation. It only has truth values relative to contexts of evaluation. For instance, the proposition P is true relative to whichever context of evaluation is identical to the context in which T_T occurs. Meanwhile, P is false relative to whichever context of evaluation is identical to the context in which T_F occurs. Since the proposition P only has truth values relative to certain contexts of evaluation, P is a relativized proposition. So, to generalize: for any possible world W, if two thoughts have the same content yet differ in truth value, those thoughts are attitudes towards a relativized proposition (the same for both thoughts).

With all this in mind, consider the following argument:

(1) Necessarily, if two thoughts share (representational) content (of some sort) but only one of them is true, then those two thoughts are attitudes towards a relativized proposition (the same for both thoughts).

(2) Necessarily, there are no relativized propositions.

Therefore,

(3) CFT: Necessarily, if two thoughts share (representational) content (of some sort), then either both thoughts are true or neither is. (from 1 & 2)

Call this argument "the Argument for CFT from the Traditional Doctrine of Belief" – for short, "the Argument from Tradition."

The Argument from Tradition is sound provided that (I) and (II) are both true. To begin, since necessitation is closed under entailment, (1) and (2) entail (3). So the Argument from Tradition is valid. And, above, we saw that, if (I) is true, (1) is true. Now recall from §1 that (II) is equivalent to the claim that, necessarily, there are no relativized propositions, nor are there any propositions that simply lack truth values. So (II) entails (2). So the Argument from Tradition is

sound – and its conclusion, CFT, is true – provided that both tenets of the traditional doctrine of belief, as I am characterizing this doctrine, are true.

3 – The Case for Relativized Propositions

Recall from Chapter One that, if CFT is true, internalism is false. So the internalist – along with anyone else who denies CFT – must resist the Argument from Tradition. And since the Argument from Tradition is sound if both doctrines of the traditional doctrine of belief are true, there are precisely two options for resisting the argument. One option is to deny (I). The other option is to accept (I) but then to insist that some thoughts – namely, those thoughts whose contents can be shared by thoughts with differing truth values – are attitudes towards relativized propositions.[9] In insisting that there are thoughts of this sort, one is thereby denying (II).

In this and the next two sections, I consider the case for relativized propositions, given (I). I begin, in this section, by presenting the case for such propositions. I construe the case for relativized propositions as consisting of two steps. The first step is to make relativized propositions plausible by presenting a picture of what propositions of this sort might look like (§3.1). The second step is to extoll the virtues of relativized propositions (§3.2).

3.1 –Different Conceptions of Relativized Propositions

Again, a relativized proposition is a proposition that has truth values relative to – and *only* relative to – contexts of evaluation. Any plausible, informative conception of relativized propositions will tell us, among other things, how to determine, for any relativized proposition P, which contexts of

[9] Technically, one need only hold that it's possible for there to be such thoughts (and, consequently, that it's possible for there to be relativized propositions). But anyone who thinks that this is possible should also think that it is actual. So, for simplicity's sake, I'll characterize this second strategy for resisting the Argument from Tradition as centering on the claim that there really are thoughts that are attitudes toward relativized propositions.

evaluation are the ones relative to which P is true, and which contexts of evaluation are the ones relative to which P is false. The most straightforward way to develop such a conception of relativized propositions is to characterize relativized propositions partly in terms of centered worlds – formal devices that include, among other things, contexts of evaluation for propositions (beliefs, etc.). (See Chapter One, §5, for a more detailed characterization of centered worlds.)

David Lewis (1979: 531 – 533), for instance, characterizes relativized propositions as *sets* of centered worlds (though see footnote 10 for a qualification). More specifically, for Lewis, a relativized proposition P is the set of centered worlds at which P is true. Since this proposition P is a set of *centered* worlds, there is a fact of the matter about what its members are. In particular, for any possible world W, possible individual x, and time t, P is true (false) at the centered world {W, {x, t}} if and only if, at W, P is true (false) relative to the context of evaluation {x, t}. Note that this gives us a prescription for determining whether a relativized proposition P is *actually* true (false) relative to a subject-time ordered pair. In particular, for any individual x and time t, a relativized proposition P is actually true (false) relative to {x, t} if and only if the centered world $\{W_\alpha, \{x, t\}\}$ is (is not) a member of P. (And, W_α is the actual world.) For instance, the proposition <that there is an apple on my head> both is true at $\{W_\alpha$, {Jackson, 3 p.m.} and has $\{W_\alpha$, {Jackson, 3 p.m.} as a member; meanwhile, this proposition is false at $\{W_\alpha$, {*Twin* Jackson, 3 p.m.}}, and it does *not* have $\{W_\alpha$, {*Twin* Jackson, 3 p.m.}} as a member.[10]

[10] Quine (1969: 15-22) is the first to propose that we conceive of relativized propositions as sets of centered worlds, but Quine himself quickly abandons this proposal. Meanwhile, Lewis (1979) uses the *term* 'proposition' to characterize sets of worlds (including sets of centered worlds), but strictly speaking Lewis would not say that there are propositions (relativized or otherwise) in *my* sense of 'proposition'. For Lewis (1979) not only denies the first tenet of the traditional view of belief. He also denies – or at least would deny – that there are entities of a sort that both (i) are *the* entities toward which our thought tokens are attitudes and (ii) are bearers of truth and falsity. (Rather, for Lewis, our thought tokens, at least fundamentally, are attitudes wherein we ascribe – and, typically, *self*-ascribe – *properties*, where for Lewis all sets of worlds are properties but not all properties are sets of worlds.)

A related characterization of relativized propositions takes these propositions to be *functions* from centered worlds to truth values. More specifically, on this account, a relativized proposition $\underline{P}$ is a function *f* such that – for any possible world W, possible individual x, and time t – *f* maps the centered world {W, {x, t}} to the value *true* (*false*) if and only if, at W, $\underline{P}$ is true (false) at {x, t}. Accordingly, on this account, a relativized proposition $\underline{P}$ is *actually* true (false) relative to a context of evaluation {x, t} just in case $\underline{P}$ maps $\{W_{\alpha}, \{x, t\}\}$ to the value *true* (*false*). David Chalmers (1996: 63 – 65) adopts a particular version of this view. For Chalmers (1996: 63 – 65), a thought is an attitude toward (though not exclusively toward) a "centered proposition," which is what Chalmers (2002a) identifies with the thought's epistemic intension.[11]

On both of the above centered-worlds-based characterizations of relativized propositions, relativized propositions are identical to mathematical entities of a certain sort. But one needn't take it that relativized propositions literally are mathematical entities to find the appeal to centered worlds useful. For instance, Andy Egan (2006: 104-111) holds that relativized propositions either are sets of centered worlds or at least can be *modeled* as sets of centered worlds. In other words, for Egan, whether or not relativized propositions literally are sets of centered worlds, we can individuate relativized propositions in terms of their truth conditions, as represented by sets of centered worlds. More specifically, for any relativized proposition $\underline{P}$, there is one (and only one) set S of centered worlds such that – for any centered world $\underline{W}$ – $\underline{P}$ is true at $\underline{W}$ if and only if $\underline{W}$ is a member of S. And, for any set S of centered worlds, there is one (and only one) relativized proposition $\underline{P}$ such that – for any centered world $\underline{W}$ – $\underline{P}$ is true at $\underline{W}$ if and only if $\underline{W}$ is a member of S.[12] (A close cousin of this account – or perhaps just a different way of putting the same account

[11] So, for Chalmers, a centered proposition maps epistemically possible centered worlds to its truth values at those worlds, considered as *actual*. (See Chapter One, §6.)

[12] In adopting any one of the candidate accounts of relativized propositions I've just presented, one is thereby committed to representationalism. (See Chapter One, §4.) But one could avoid this commitment by

– individuates relativized propositions in terms of their truth conditions, as represented by *functions* from centered worlds to truth values.)

3.2 – Why Believe in Relativized Propositions?

We've already seen one consideration in favor of relativized propositions: their existence follows from the conjunction of internalism and the first tenet of the traditional doctrine of belief. For internalists must deny CFT and, consequently, must either deny the first tenet of the traditional doctrine of belief or, instead, posit relativized propositions.

So internalists (e.g., Chalmers 1996, 2002a, & 2002b) have a special reason to accept relativized propositions. But philosophers posit relativized propositions for all sorts of reasons: e.g., to accommodate differences in taste (Kölbel 2002 and Einheuser 2008 & 2012), to make sense of epistemic modal assertions (Egan 2006), to explain knowledge attributions (MacFarlane 2005), to countenance differences in conceptual schemes (Lynch 1998), and to make sense of moral discourse (Capps, Lynch, & Massey 2009).[13]

Finally, and ironically, perhaps the most compelling consideration in favor of relativized propositions comes from, though is not endorsed by, Frege (1892, 1956). Frege individuates propositions by their sense, or cognitive value. Given Frege's commitment to the two traditional

---

accepting, instead of the final account of relativized propositions I've presented, the following weaker conception of relativized propositions: for any relativized proposition P, there is one (and only one) set S of centered worlds such that – for any centered world W – P is true at W if and only if W is a member of S. On this weaker conception of relativized propositions, such propositions are no longer individuated completely in terms of sets of centered worlds. But this is the right result for those who deny representationalism.

[13] As Max Kölbel observes, Prior (1967) and Kaplan (1989) both "have allowed *tensed* propositions, i.e., propositions whose truth-value varies with time" (2008: 2). Kölbel infers from this observation that Prior and Kaplan both reject the "Fregean assumption that propositions have their (actual) truth-values absolutely" (2008: 2). But see fn. 7 for why Kölbel's inference fails to acknowledge the spirit of Frege's assumption and for why, as I define 'relativized proposition', a proposition's varying with time is not sufficient for its being relativized.

doctrines of belief, a consequence of this way of individuating propositions is that some beliefs – for instance, egocentric beliefs – are attitudes toward *private* propositions, where a private proposition is a proposition that only one subject can entertain. For many philosophers, this consequence is unacceptable. And one way to avoid this consequence is to posit the existence of relativized propositions, and in doing so to deny the second tenet of the traditional doctrine of belief. (Both David Chalmers 1996: 60-65, 2002a, & 2002b and Andy Egan 2006: 104-111, for instance, posit relativized propositions in part because it allows them to avoid private propositions. Chalmers 2011b, on the other hand, posits private propositions.)

In §6 and §7, I'll discuss the connection between Fregean sense, private propositions, and relativized propositions in great detail. And in doing so I'll have more to say about what exactly a private a proposition is supposed to be and about why many philosophers find belief in such propositions unacceptable.[14]

4 – Claim (I) Rules out Relativized Propositions as Sets of Centered Worlds (and as Functions from Centered Worlds to Truth Values)

I have just presented – but have not endorsed – the case (or at least *a* case) for relativized propositions. And in doing so I've introduced some ways of conceiving of relativized propositions, two of which are the following: (i) relativized propositions are sets of centered worlds, and (ii) relativized propositions are functions from centered worlds to truth values.

[14] There is at least one (*prima facie*) reason to posit relativized propositions that I have not yet cited. John MacFarlane (2003) posits relativized propositions (or at least proposes an account that plausibly entails the existence of relativized propositions) in order to allow for the *openness of the future*. I'll discuss MacFarlane's view in a bit more detail in §5, at which point his view may seem to give rise to a plausible objection to my overall argument in this chapter. I shall show, however, that the objection in question is in fact a superficial one, and it could be avoided by adding some (philosophically harmless) complications to the Argument from Tradition.

It turns out that both of these ways of conceiving of relativized propositions are inconsistent with (I). That is unfortunate, for relativized propositions were supposed to provide the internalist with a way resisting the Argument from Tradition without having to deny (I).

We can begin to see why there is this inconsistency by noting that (I) entails the following:

> (I-Cor.) Necessarily, the fundamental bearers of (representational) content are propositions.

To say that an entity x is a *fundamental* bearer of content is to say that x has content but does not inherit that content from anything else. (I-Cor.) says two things: first, that there exist certain entities that are fundamental bearers of content; second, that these entities are propositions. For instance, if (I-Cor.) is true, then there exists a proposition representing snow as being white, but not in virtue of something else's representing snow as being white. Call this proposition "the proposition <that snow is white>." If (I-Cor.) is true, then the proposition <that snow is white> exists and is a fundamental bearer of content.

To begin to see that (I) entails (I-Cor.), assume, for *reductio*, that there are representations of some kind or other (thought tokens, propositions, speech acts, etc.) that are not fundamental content bearers and that do not inherit their contents (either directly or indirectly) from entities that are fundamental content bearers. For any such representation R, one of two things is the case:

> **Option One.** R is part of a *non*-terminating chain of inheritance, in the sense that R inherits its content from some other entity x, which inherits its content from some other entity y, and so on *ad infinitum* (where each member of the inheritance series is distinct)
>
> **Option Two.** R is part of a circle of inheritance, in the sense that R inherits its content from some other entity x, which inherits its content from some other entity y, and so on, leading up to some entity z's inheriting its content from R (which inherits its content from x…).

But there is no such chain of representations, nor is there any such circle of representations. For say that there is. Within such a chain or circle, the content of each member of the chain or circle is explained in terms of the content of some other member. But consider the chain or circle as a whole. There is no explanation whatsoever for why it involves the content it involves.[15] Of course, explanation has to stop somewhere. But it surely does not stop with such a chain or circle. Positing such a chain or circle would be entirely ontologically profligate. For positing such a chain or circle would not help us to explain why it is that thoughts, propositions, speech acts, or representations of any other sort have the contents they have. So we should deny that there is any such chain or circle (let alone a multitude of such chains or circles). So we should deny what we assumed for *reductio*. We should conclude, then, that there are no representations that neither are fundamental bearers of content nor inherit their contents (either directly or indirectly) from fundamental bearers of content. And, since we haven't relied on anything contingent in reaching this conclusion, we should further conclude that:

> (4) Necessarily, for any representation R, either R is a fundamental bearer of (representational) content or R inherits content (either directly or indirectly) from a fundamental bearer of (representational) content.[16]

Now assume, for conditional proof, that (I) is true. So thoughts inherit their contents from propositions. It follows that thoughts are not fundamental content bearers. But, if thoughts are not fundamental content bearers, then there are no plausible candidates left for being the fundamental content bearers other than propositions. So either there are no fundamental content bearers, or propositions are the fundamental content bearers. If there are no fundamental content bearers, then

[15] Cf. Rowe (2007: 24 – 30) on a version of the Cosmological Argument.
[16] My argument for (4) expounds upon an argument by Merricks (2015: 209) for roughly the same conclusion.

thoughts inherit their contents but do not (even indirectly) inherit their contents from fundamental content bearers. Given (4), this is impossible. It follows that propositions are the fundamental bearers of (representational) content. So (I-Cor.) is true. So, to discharge our assumption for conditional proof: if (I) is true, then (I-Cor.) is true. And nothing in our argument for this conclusion relies on anything contingent. So we may conclude that, necessarily, if (I) is true, then (I-Cor.) is true.

So (I) entails (I-Cor.).

Here's why this matters. Necessarily, if relativized propositions either are sets of centered worlds or are functions from centered worlds to truth values, then relativized propositions have whatever (representational) contents they have at least partly in virtue of how we interpret them. If so, relativized propositions are not fundamental content bearers. So the conception of relativized propositions as either sets of centered worlds or functions from centered worlds to truth values is inconsistent with (I-Cor.) and, consequently, with (I).

My argument for the requirement that relativized propositions, so conceived, be interpreted is a repurposing of a related argument versions of which are defended by Joseph Moore (1999), by Michael Jubien (2001: 48 – 49), and by Trenton Merricks. Moore, Jubien, and Merricks all argue that, necessarily, if propositions are sets of *possible* worlds or functions from *possible* worlds to truth values, then – lest propositions simply have the contents they have in virtue of how we interpret them – it is objectionably arbitrary which proposition represents what. But Moore, Jubien, and Merricks all reject the conception of propositions as having the contents they have in virtue of how they are interpreted. Consequently, Moore, Jubien, and Merricks reject the conception of propositions as sets of possible worlds or as functions from possible worlds to truth values.[17]

[17] And cf. Paul Benacerraf (1965) for an argument targeting the view that *numbers* are sets of a certain sort.

A slightly weaker version of the Moore, etc., argument leaves open the question of whether propositions may be interpreted and concludes, simply, that: necessarily, if propositions either are sets of centered worlds or are functions from centered worlds to truth values, then propositions have whatever contents they have at least partly in virtue of how we interpret them. I endorse this argument. And since I endorse this argument I also endorse the analogous argument that applies *mutatis mutandis* to relativized propositions, in particular, conceived of as sets of *centered* worlds or functions from *centered* worlds to truth values.

I present this argument – in two parts – in what follows.

Begin with the assumption that relativized propositions are sets of centered worlds. In §3.1, when presenting the view that relativized propositions are sets of centered worlds, I focused on the more specific view that relativized propositions are the sets of centered worlds at which they are true. But now consider, again, the relativized proposition representing *the first-personal subject* as having an apple on his head. Again, I have labeled this proposition "the proposition <that there is an apple on my head>." On our assumption that relativized propositions are sets of centered worlds, this proposition is a set – namely, whichever set of centered worlds represents the first-personal subject as having an apple on his head. On the specific version of the sets-of-centered-worlds view on which I focused in §3.1, this is the set Σ_J such that – for any possible world W, possible individual x, and time t – the centered world {W, {x, t}} is a member of Σ_J if and only if, in W, x has an apple on his head at t. But, we might reasonably ask, why is it that *this* set of centered worlds is the one that represents the first-personal subject as having an apple on his head? Why isn't it, instead, the set Δ_J such that – for any possible world W, possible individual x, and time t – the centered world {W, {x, t}} is a member of Δ_J if and only if, in W, x *does not* have an apple on his at t? It seems the difference here is entirely arbitrary, unless of course we simply

interpret the set Σ_J as the set that represents the first-personal subject as having an apple on his head. So, if the set of centered worlds representing the first-personal subject as having an apple on his head really is Σ_J (as opposed to, say, Δ_J), then this set has the content it has at least partly in virtue of its being interpreted as having that content. And note that this is also the case, for the same reason, if we take the set of centered worlds representing the first-personal subject as having an apple on his head to be some other set of centered worlds (e.g., the set Δ_J). So the set of centered worlds representing the first-personal subject as having an apple on his head has its content at least partly in virtue of our interpreting it in a certain way.[18] So the *relativized proposition* representing the first-personal subject as having an apple on his head has its content at least partly in virtue of our interpreting it in a certain way. And notice that we can tell the same story for any other relativized proposition, taken to be a set of centered worlds. So, to discharge our assumption: if relativized propositions are sets of centered worlds, then relativized propositions have the particular contents they have at least partly in virtue of how we interpret them.

An analogous argument may be given for the conclusion that, if relativized propositions are functions from centered worlds to truth values, then relativized propositions have the contents they have at least partly in virtue of their being interpreted in certain ways. Assume that relativized propositions are functions from centered worlds to truth values. Which function from centered worlds to truth values represents the first-personal subject as having an apple on his head? Is it the function f_J such that – for any individual x, time t, and possible world W – f_J maps {W, {x, t}} to the value *true* if and only if, in W, x has an apple on his head at t (and, otherwise, f_J maps {W, {x,

[18] Here we need to be careful. The expression 'the proposition <that there is an apple on my head>' is not short-hand for 'whichever proposition represents the first-personal subject's having an apple on his head'. Rather, it is a rigid designator for the entity that, *as a matter of fact*, is the proposition representing the first-personal subject as having an apple on his head (at least assuming that there is such a proposition).

t}} to the value *false*)? Or is it instead the function g_J such that – for any individual x, time t, and possible world W – g_J maps {W, {x, t}} to the value *false* if and only if, in W, x has an apple on his head at t (and, otherwise, g_J maps {W, {x, t}} to the value *true*)? Or is it some other function from centered worlds to truth values? It seems that there can be no non-arbitrary answer to this question, unless of course we simply *interpret* one of these functions (e.g., the function f_J) as representing the first-personal subject as having an apple on his head. But, then, this function has the content it has at least partly in virtue of how we interpret it. So the *relativized proposition* <that there is an apple on my head> has the content it has at least partly in virtue of how we interpret it. And the same general story may be told for any other relativized proposition, considered to be a function from centered worlds to truth values. So, to discharge our assumption: if relativized propositions are functions from centered worlds to truth values, then relativized propositions have the particular contents they have at least partly in virtue of how we interpret them.

The arguments presented over the above two paragraphs rely on nothing contingent. We may therefore conclude the following: necessarily, if relativized propositions either are sets of centered worlds or are functions from centered worlds to truth values, they have the particular contents they have at least partly in virtue of how we interpret them.

Necessarily, if a relativized proposition has the content it has partly in virtue of how we interpret it, then it is not a fundamental content bearer. For, necessarily, if a relativized proposition has the content it has partly in virtue of how we interpret it, then it inherits its content from some other representational state – in particular, the representational state (or one of the representational states) that is involved in the act of interpreting that proposition.

It follows that, necessarily, if (I) is true, then relativized propositions (even if they exist) are neither sets of centered worlds nor functions from centered worlds to truth values. Again, (I)

is the claim that, necessarily, if two thoughts share (representational) content (of some sort), then they inherit that content from a single proposition. And, as we saw above, (I) entails (I-Cor.), the claim that, necessarily, the fundamental bearers of (representational) content are propositions. But, again, necessarily, if a relativized proposition has the content it has partly in virtue of how it is interpreted, then it is not a fundamental bearer of content. So, necessarily, if relativized propositions either are sets of centered worlds or are functions from centered worlds to truth values, then they are not fundamental bearers of content. So (I-Cor.) entails that relativized propositions (if they exist at all) are neither sets of centered worlds nor functions from centered worlds to truth values. So (I) entails the same thing. In other words: necessarily, if (I) is true, then relativized propositions (even if they exist) neither are sets of centered worlds nor are functions from centered worlds to truth values.[19]

5 – Claim (I) Rules out Relativized Propositions and, Consequently, Internalism

(I) entails that, if there are relativized propositions at all, they are neither sets of centered worlds nor functions from possible worlds to truth values. More generally, (I) entails that, if there are relativized propositions at all, they are the fundamental bearers of content and, as such, do not need to be interpreted in any way to have the contents they have.

But this is not all that (I) entails. (I) entails that there are no relativized propositions, period. I argue for this entailment in what follows. I begin, in §5.1, by discussing a challenge faced by those who deny the possibility of relativized propositions, and then by proposing an argument

[19] Recall from §3.1 that David Lewis (1979) conceives of relativized "propositions" – in his sense of the term 'proposition' – as sets of centered worlds. (Indeed, for Lewis, 'proposition' really just means "set of worlds." See fn. 10.) But Lewis does not hold that "propositions" – in his sense of 'proposition' – are the fundamental bearers of content and, as such, are the entities from which thoughts inherit their contents. In fact, Lewis seems happy to allow that "propositions" represent what they do (*if* they represent at all) at least partly in virtue of how we interpret them. (See, once again, fn. 10. See also Lewis 1986: 144.)

based on a response to this challenge. The conclusion of this argument is that relativized propositions are impossible. The argument is sound provided that (I) is true.

5.1 – The Argument against Relativized Propositions

Just as a context of belief is a context in which a belief may occur, a context of utterance is a context in which a sentence (or some other linguistic entity) may be uttered.

Independent of a context of utterance, the sentence 'My pants are on fire' is neither true nor false. It is, however, true in certain contexts of utterance and false in others. Or at least it can be. Say, for instance, that at 1 p.m. today Darren and Dori both utter 'My pants are on fire', with assertoric force, even though at that time only Darren's pants are on fire. In this case, 'My pants are on fire' is true as uttered by Darren but false as uttered by Dori.

One *prima facie* plausible way to explain the above sensitivity of truth value to context of utterance is to say that the sentence 'My pants are on fire' is neither absolutely true nor absolutely false but, instead, is true relative to certain contexts of evaluation and false relative to others. In particular, one may say, the sentence 'My pants are on fire' is true relative to the context of evaluation {Darren, 1 p.m.} but false relative to the context of evaluation {Dori, 1 p.m.}.

This explanation suggests the following challenge to those who deny the possibility of relativized propositions: provide a principled explanation for how there can be relativized representations of certain sorts – e.g., the sentence 'My pants are on fire' – even though there cannot be relativized *propositions*.

One way to respond to this challenge would be to deny the presupposition on which the challenge is based, i.e., to reject the above explanation for the fact that the sentence 'My pants are on fire' has a truth value only in certain contexts of utterance. There is no such entity – one might argue – as the sentence 'My pants are on fire', which entity is true in certain contexts of utterance

and false in others. Rather, certain utterances of the sentence 'My pants are on fire' are true, whereas others are false. And what makes these different utterances both utterances of the sentence 'My pants are on fire' is not that there is some entity – the sentence 'My pants are on fire' – that these utterances have in common. Rather, these utterances have in common a certain phonic structure. Moreover – and in part because they have in common a certain phonic structure – they play a common role in a particular linguistic practice. (They are, in this sense, *tokens* of the 'My pants are on fire' sentence type.)

For what it's worth, I think that the above way of responding to the challenge is the right one. But set that aside. Let's just grant, for the sake of argument, that there is such an entity as the sentence (i.e., the sentence *type*) 'My pants are on fire', which sentence is relativized – true relative to the context of evaluation {Darren, 1 p.m.} and false relative to the context of evaluation {Dori, 1 p.m.}. In this case, those who deny the possibility of relativized propositions cannot simply dismiss the above-presented challenge; they must meet it.

And they *can* meet it, *at least provided that they accept (I)*. Those who deny the possibility of relativized propositions and who accept (I) can meet the above challenge by appeal to the following distinction:

> **Distinction.** It may be possible for there to be certain *non*-fundamental bearers of content that are relativized. But it is impossible for there to be a *fundamental* bearer of content that is relativized.

Recall that

> (I) Necessarily, if two thoughts share (representational) content (of some sort), then they inherit that content from a single proposition.

entails

(I-Cor.) Necessarily, the fundamental bearers of (representational) content are propositions.

So, for the proponent of (I), **Distinction** explains why propositions – as the fundamental bearers of content – cannot be relativized even if other bearers of representational content can be. Moreover, the proponent of (I) who denies the possibility of relativized propositions can offer a principled explanation for why **Distinction** is true. Consider the following:

> **Relativized Content.** Necessarily, for any relativized representation R and context of evaluation {x, t} relative to which R is true (false), there is a proposition P such that: (i), *relative to* {x, t}, R inherits the representational content of P, (ii) P is true (false), and (iii) it is in virtue of (i) and (ii) that R is true (false) relative to {x, t}.

For instance, the sentence 'My pants are on fire' is true relative to {Darren, 1 p.m.} since, relative to {Darren, 1 p.m.}, it expresses – and, consequently, inherits the representational content of – the true proposition <that Darren's pants are on fire at 1 p.m.>. Meanwhile, 'My pants are on fire' is false relative to {Dori, 1 p.m.} since, relative to {Dori, 1 p.m.}, it expresses – and, consequently, inherits the representational content of – the false proposition <that Dori's pants are on fire at 1 p.m.>. Or so one who endorses Relativized Content should say.

So Relativized Content provides a plausible explanation for how non-fundamental bearers of content can be relativized (at least given the assumption, which I made for the sake of argument, that there can be certain non-fundamental bearers of content that are relativized – e.g., the sentence 'My pants are on fire'). But the explanation cannot be extended to fundamental bearers of content. For essential to the explanation is that relativized bearers of content inherit different contents relative to different contexts of evaluations. But, necessarily, fundamental bearers of content do

not inherit different contents relative to different contexts of evaluation. For, necessarily, fundamental bearers of content do not *inherit* content at all.

This explanation suggests the following argument for the impossibility of relativized propositions:

(5) Necessarily, if propositions exist and are the fundamental bearers of content, then each proposition has the representational content it has independently of a context of evaluation.

(6) Necessarily, if a proposition P has the representational content it has independently of a context of evaluation, then P has the truth value it has independently of a context of evaluation, i.e., P is not relativized.

Therefore,

(7) Necessarily, if a proposition P is a fundamental bearer of content, then it is not relativized. (from 9 & 10)

(8) I-Cor.: Necessarily, propositions are the fundamental bearers of representational content.

Therefore,

(9) Necessarily, there are no relativized propositions. (from 11 & 12)

Call this argument the Argument against Relativized Propositions.

The Argument against Relativized Propositions is sound provided that (I) is true. To begin, (5) and (6) entail (7) since necessitation is closed under entailment, and (7) and (8) entail (9) for the same reason. So the argument is valid. And (8) is just (I-Cor.), which – again – is entailed by (I). In §5.2 I'll argue for (5), and in §5.3 I'll argue for (6).

5.2 – Premise (5)

Assume, for conditional proof, that there are representations of some sort that have their representational contents only relative to contexts of evaluation. The only plausible explanation

for this is that any such representation inherits its representational content from different fundamental bearers of content, relative to different contexts of evaluation. (E.g., relative to {Darren, 1 p.m.}, 'My pants are on fire' inherits the content of the proposition <that Darren's pants are on fire at 1 p.m.>, and relative to {Dori, 1 p.m.} it inherits the content of the proposition <that Dori's pants are on fire at 1 p.m.>.) This explanation entails that, if a representation R has its content only relative to a context of evaluation, R is not a fundamental bearer of content. So, discharging my assumption for conditional proof, I conclude the following: if a representation R is a fundamental bearer of content, R has the content it has independently of a context of evaluation. Note that my argument for this conclusion does not rely on anything contingent. So I conclude the following:

(5-R) Necessarily, the fundamental bearers of content – whatever they are – have their representational contents independently of contexts of evaluation.

(5) follows straightforwardly from (5-R). So, I conclude, (5) is true.

But say that, for whatever reason, you are suspicious of my argument for (5-R). Even still, you should follow me in endorsing (5). Here is why.

First, the point of positing propositions – at least for those who take them to be the fundamental bearers of content – is to appeal to propositions as ways of individuating the representational contents of thoughts and of other representational states. But this requires that we identify each proposition as having a unique representational content. And in doing so we commit ourselves to the view that a proposition has the representational content it has independently of a context of evaluation.

Second, assume, for *reductio*, that – possibly – (i) there are propositions that have the contents they have only relative to contexts of evaluation and (ii) propositions are the fundamental bearers of content. Now recall from §4 that,

> (4) Necessarily, for any representation R, either R is a fundamental bearer of (representational) content or R inherits content (either directly or indirectly) from a fundamental bearer of (representational) content.

So, necessarily, if propositions are the fundamental bearers of content, then thoughts inherit their contents from propositions, from which it follows that – if two thoughts share content – they inherit that content from a single proposition. Recall also that, by definition, thoughts inherit contents from propositions (in part) by being attitudes towards those propositions. So it follows from our assumption for *reductio* that – possibly, (i) there are propositions that have the contents they have only relative to contexts of evaluation and (ii*) propositions are the things from which thoughts inherit their contents, such that two thoughts that share content inherit that content from a single proposition. But, if a proposition has its content only relative to certain contexts of evaluation, then we can't make sense of that proposition as having a single content that different thought tokens – had in different contexts – inherit. So we should deny the compossibility of (i) and (ii*). But, then, we should deny what we assumed for *reductio*. So we should endorse (5).

It's worth noting, as an aside, that my argument in the above paragraph also shows that, just as (I) entails (I-Cor.), (I-Cor.) entails (I). Necessarily, if propositions are the fundamental bearers of content, then thoughts inherit their contents from propositions, such that – if two thoughts share content – they inherit that content from a single proposition. It follows that, necessarily, if propositions are the fundamental bearers of content, then – necessarily – if two

thoughts share content then they inherit that content from a single proposition.[20] In other words: if (I-Cor.), then (I). Furthermore, since we haven't relied on anything contingent to see this, we may conclude that (I-Cor.) entails (I).

5.3 – Premise (6)

Let 'R*' be an arbitrary name for any representation that has its representational content independently of a context of evaluation. In this case,

> (A) R* represents things as being a certain way, and which way R* represents things as being is independent of a context of evaluation.

But

> (B) For any way a representation could represent things as being, either – independent of any context of evaluation – things are that way or – independent of any context of evaluation – things are not that way.

Therefore,

> (C) Either – independent of any context of evaluation – things are the way R* represents them as being or – independent of any context of evaluation – things are not the way R* represents them as being.

But, of course,

> (D) If – independent of any context of evaluation – things are the way R* represents them as being, then – independent of any context of evaluation – R* is true;

and

[20] This is just an application of modal principle K: $\Box(A \rightarrow B) \rightarrow (\Box A \rightarrow \Box B)$, for any values of A and B.

(E) If – independent of any context of evaluation – things are not the way R* represents them as being, then – independent of any context of evaluation – R* is false.

And it follows from (C), (D), and (E) that

(F) Either – independent of any context of evaluation – R* is true or – independent of any context of evaluation – R* is false.

So, given (F),

(G) R* has its truth value independently of a context of evaluation, i.e., R* is not relativized.

Finally, since 'R*' is an arbitrary name for any representation that has its representational content independently of a context of evaluation, (G) generalizes to the following:

(H) For any representation whatsoever, if that representation has its content independently of a context of evaluation, it also has its truth value independently of a context of evaluation, i.e., it is not relativized.

Since my argument for (H) relies on nothing contingent, (H) is necessarily true. So, necessarily, if a representation has the representational content it has independently of a context of evaluation, it has the truth value it has independently of a context of evaluation as well. *A fortiori*,

(6) Necessarily, if a proposition P has the representational content it has independently of a context of evaluation, then P has the truth value it has independently of a context of evaluation, i.e., P is not relativized.

I can think of one plausible objection to my argument for (6), which objection should be especially attractive to internalists who accept (I). The objection runs as follows:

> **Objection from Internalism.** There are ways certain representations represent things as being such that things are that way relative to certain contexts of evaluation but are not that way relative to other contexts of evaluation. Consider, for instance, the proposition <that there is an apple on my head>, i.e., the proposition toward which Jackson's apple-on-head belief and Twin Jackson's apple-on-head belief are both attitudes. This proposition represents things as being a certain way. But there is no context-independent fact about whether things are that way. Rather, things are that way relative to some contexts of evaluation – e.g., relative to {Jackson, 3 p.m.} – and things are *not* that way relative to other contexts of evaluation – e.g., relative to {Twin Jackson, 3 p.m.}.

The objection from internalism purports to identify a counterexample to (B). Here, the counterexample (or supposed counterexample) is the proposition <that there is an apple on my head>. But any proposition will do provided that (and *only* provided that) it meets the following criterion: there are (or at least could be) multiple egocentric thought tokens, had by different thinkers, all of which are attitudes toward that proposition.[21] As we saw in §2, internalists who accept (I) should say that there are such propositions.

This objection is exactly the objection that the internalist who accepts (I) should make – and must make – against my argument for (6). But the objection fails. We can see that it fails by asking ourselves, "How does the proposition <that there is an apple on my head> represent things as being?"

[21] Say that, at 1 p.m., I think to myself – here's how I'd put it – "It is currently raining" And say that, at 2 p.m., I think to myself – here's how I'd put it – "It is currently raining." And imagine that it is raining at 1 p.m. but not at 2 p.m. Some philosophers sympathetic to the above objective may also wish to claim that my rain-related belief tokens are both attitudes towards the proposition <that it is currently raining> and that this proposition is also a counterexample to (B). But say that the A-theory of time is true. Then there is a fact of the matter – independent of context – about whether it is currently raining. Say, instead, that the B-theory is true. To my mind, the most plausible account of tensed belief consistent with the B-theory says that tensed beliefs are, most fundamentally, self-ascriptions of temporal properties, made by momentary stages. (See, e.g., Lewis 1979.) If so, then "my" 1 p.m. belief token and 2 p.m. belief token are themselves egocentric thought tokens, had (fundamentally) by different thinkers.

The only plausible answer to this question is that the proposition <that there is an apple on my head> represents a particular way things could be (in some sense of 'could'), namely the following:

The possibility of *the first-personal subject's* having an apple on his head.

But we need to be careful here. There is no x – nor could there by an x – such that x is *the first-personal subject*. So, for instance, the first-personal subject is not Jackson. Nor is it Twin Jackson. Nor is it you or I. Nor is it whichever individual happens to be the first-personal subject of the thought token in question, for there is no single thought token here.

So what *is* the first-personal subject? It is whatever we are all thinking about when we employ the *I*-concept. Of course, there is no x such that we are all thinking about x when we employ the *I*-concept. Nevertheless, there is, in some sense, a common subject of our first-personal thoughts – "something" that our first-personal thoughts all represent. And this is *the first-personal subject*. Or so one must say if one is to insist that the proposition <that there is an apple on my head> exists and represents a distinct possibility independently of any context of evaluation.

But characterizing the first-personal subject in this way undermines the objection from internalism. To demonstrate this, I'll report my own thought process in attempting to determine whether the first-personal subject has an apple on his head. When I employ the *I*-concept, I do not think about myself along with some ethereal first-personal subject. Of course, when I employ the *I*-concept, I do think about myself, and I do think about the first-personal subject. But for me to think about the first-personal subject *just is* for me to think about myself. And, when I ask whether this first-personal subject, so understood, has an apple on his head, I simply ask – here's how I'd put it – "Is there an apple on my head?" And I answer – here's how I'd put it – "I do not." I find,

then, that the answer to the question of whether the first-personal subject has an apple on his head is simply that he does not. It is just true *simpliciter* of this first-personal subject that he does not have an apple on his head. So it turns out that whether the first-personal subject has an apple on his head is not a relative matter at all. Independent of context, the first-personal subject does not have an apple on his head.

Or so I have just argued. But you don't have to take my word for it. Ask yourself whether the first-personal subject has an apple on his head. Most likely, you will get the same result: namely, that the first-personal subject does not have apple on his head. But maybe you're Frank Jackson and it's 3 p.m. In this case, the result you'll get will be that the first-personal subject does have an apple on his head. And this just makes things worse. For my answer is no better than yours. But my answer is the correct one. So yours must be correct as well. And now we have a contradiction.

I conclude that there is no plausible way to conceive of what the first-personal subject is that allows us to make sense of the proposition <that I have an apple on my head> as representing a way things could be, such that there is no fact of the matter – independent of a context of evaluation – whether things are that way.

I can think of three ways that one may attempt to resist my argument, but none of these ways of resisting my argument are successful:

> **First attempt.** Again, when I ask whether the first-personal subject has an apple on his head, I simply ask – here's how I'd put it – "Is there an apple on my head?" And I answer – here's how I'd put it – "I do not." I then infer from this answer that the first-personal subject does not have an apple on his head. But, one may insist, this inference is not valid. For the inference relies on the assumption that – independent of a context of evaluation – NKR is identical to the first-personal subject. And, one may insist, NKR is identical to the first-personal subject only relative to certain contexts of evaluation – e.g., {NKR, 3 p.m.} – and not to others – e.g., {Jackson, 3 p.m.}.

Why the first attempt fails. My inference does not rely on the assumption that – independent of a context of evaluation – NKR is identical to the first-personal subject. My inference relies only on the assumption that, when I ask myself whether there is an apple on my head, the subject of my question – i.e., whatever my *I*-concept is about – is the first-personal subject. And, insofar as it makes sense to speak of the first-personal subject as the common subject of all of our *I*-thoughts, then the subject of my question is the first-personal subject.[22]

Second attempt. The second attempt, then, is to offer an alternative to the characterization of the first-personal subject as the common subject of all of our *I*-thoughts. On this alternative, the first-personal subject of my *I*-thoughts is me, the first-personal subject of your *I*-thoughts is you, and so on. Of course, when I have an *I*-thought, I think about myself in the same way in which you think about yourself when you have an *I*-thought, in the same way in which Jackson thinks about himself when he has an *I*-thought, and so on. But to say that we think about things in the same way is not to say that we think about the same thing.

Why the second attempt fails. Say that this alternative way of conceiving of the first-personal subject is the right way of conceiving of the first-personal subject. Then there is no single possibility as the possibility of the first-personal subject's having an apple on his head. There is just the possibility of NKR's having an apple on his head, the possibility of your having an apple on your head, the possibility of Jackson's having an apple on his head, and so on.

Third attempt. Again, when I ask whether the first-personal subject has an apple on his head, I ask – here's how I'd put it – "Is there an apple on my head?" And I answer – here's how I'd put it – "I do not." But, one might insist, this is the wrong answer. Instead, I should say that there is an apple on my head with respect to certain contexts of evaluation but not with respect to others.

Why the third attempt fails. Say that my pants are on fire. And say also that the above-presented alternative analysis of how I should answer the question "Is there an apple on my head?" is correct. And say that I think to myself – here's how I'd put it – "Are my pants on fire?" What I should think, as a response to this question, is that my pants are on fire with respect to certain contexts of evaluation but not with respect to other contexts of evaluation. Consequently, I have no reason – at least no more reason than anyone else – to stop, drop, and roll. Premise: If my pants are on fire, then I have a special reason to stop, drop, and roll.[23]

[22] Of course, if my argument is successful, then one may run an analogous argument to show that there is no sense to be made of there being a common subject to our *I*-thoughts, which subject is NKR relative to certain contexts of evaluation, Jackson relative to other contexts of evaluation, and so forth.
[23] See Perry (1979: 4 - 5) for the connection between first-personal belief and reasons to act.

I conclude, then, that there is no way things could be represented by the proposition <that there is an apple on my head> such that things are that way relative to certain contexts of evaluation but not that way relative to other contexts of evaluation. The objection from internalism fails.

And the objection from internalism, again, is the one plausible way of resisting the above argument for (6).[24] So, I conclude:

> (6) Necessarily, if a proposition P has the representational content it has independently of a context of evaluation, then P has the truth value it has independently of a context of evaluation, i.e., P is not relativized.

---

[24] There may seem to be another plausible way of resisting (B) – namely, by appealing to the possibility of *objectively unsettled* representations. Consider, for instance, the proposition <that by the year 3000 Virginia's Eastern shore will be entirely underwater>. This proposition (if it exists) is a future contingent – a representation that purports to represent what will be the case (with respect to some contingent matter). Some philosophers – proponents of the so-called "open future" – will say that it is *objectively unsettled* whether Virginia's Eastern Shore will be entirely underwater by the year 3000. (Or perhaps they will prefer another example.) These philosophers might then infer that there is a way things could be in the future – which the proposition in question represents things as being – such that it is neither the case, independent of a context of evaluation, that things will be that way nor the case, independent of a context of evaluation, that things will not be that way. And, if this is right, then the proposition in question is a counterexample to (B). (Among those philosophers who accept or at least are sympathetic to belief in the open future, some – e.g., MacFarlane 2003 – do make this inference, whereas others – e.g., Barnes and Cameron 2009 & 2011, along with Cameron 2015 – do not.)

Here is a more general way to put the objection. First, there are objectively unsettled representations; second, at least some of these – perhaps the proposition <that Virginia's Eastern Shore will be entirely underwater by the 3000>, or perhaps the proposition <that Jay Z, who is 6'1", is tall>, or perhaps my *belief* that Jay Z, who is 6'1'', is tall – are counterexamples to (B).

For what it's worth, I am inclined to deny that there are any objectively unsettled representations that are genuine counterexamples to (B). But say that you think there are such counterexamples. Then replace (B) with the following:

> (B*) For any way a representation could represent things as being, *such that it is not objectively unsettled whether things are that way*, either – independent of any context of evaluation – things are that way or – independent of any context of evaluation – things are not that way.

The appeal to objective unsettledness does not threaten (B*) at all. Of course, strictly speaking it is (B) – and not (B*) that plays a crucial role in my defense of (6) and, accordingly, of the overall argument that I am defending, in this and the previous sections of §5, from (I) to the impossibility of relativized propositions. But we can reconceive of this argument by restricting it to *objectively settled* propositions, i.e., propositions that are not objectively unsettled (i.e., that represent things as being a certain way, such that it is *objectively settled* whether things are that way). And note that this argument is still a threat to such (supposed) propositions as <that I have an apple on my head>, for this proposition – if it exists – is an objectively settled relativized propositions.

5.4 – Internalism and the Argument against Relativized Propositions

Again, the Argument against Relativized Propositions is valid (§5.1). Moreover, its premise (5) is true (§5.2), and – as I've just argued – its premise (6) is true as well (§5.3). The remaining premise of the Argument against Relativized Propositions is

> (I-Cor.) Necessarily, the fundamental bearers of (representational) content are propositions.

And I-Cor., again, is entailed by (and entails)

> (I) Necessarily, if two thoughts share (representational) content (of some sort), then they inherit that content from a single proposition.

So, if (I) is true, then the Argument against Relativized Propositions is sound and its conclusion is true. So, if (I) is true, then

> (2/9) Necessarily, there are no relativized propositions.

Now recall, from §2, the Argument from Tradition:

> (1) Necessarily, if two thoughts share (representational) content (of some sort) but only one of them true, then those two thoughts are attitudes towards a relativized proposition (the same for both thoughts).
>
> (2) Necessarily, there are no relativized propositions.
>
> Therefore,
>
> (3) CFT: Necessarily, if two thoughts share (representational) content (of some sort), then either both thoughts are true or neither is. (from 1 & 2)

As we saw in §2, the Argument from Tradition is entailed by the traditional doctrine of belief, which again consists of (I) and the following:

(II) Necessarily, propositions (if they exist) have truth values, and they have their truth values absolutely.

As we saw in §2, (1) is entailed by (I), and (2) is entailed by (II).

Of course, we have now established that (2) is also entailed simply by (I). So (I) entails both (1) and (2). So, necessarily, if (I) is true, then the Argument from Tradition is sound and its conclusion, CFT, is true. And, again, it follows from CFT that internalism (about representational mental content) is false. So it follows from (I) that internalism is false.[25]

6 – Egocentric Beliefs as Attitudes toward Private Propositions

Let's say that a proposition P is *sharable* just in case the following is possible: some individual x entertains P and, possibly, some individual y – such that x is not identical to y – entertains P. (To entertain a proposition is to be the thinker of a thought that is an attitude toward that proposition.) For instance, on the traditional view of belief, your belief that snow is white and my belief that

[25] In fn. 24 I noted that, in order to sidestep a potential objection to (B), we can replace (B) with (B*), which differs from (B*) only in that it has to do specifically with ways a representation could represent things as being, such that it is not objectively unsettled whether things are that way. As I noted, in exchanging (B) for (B*) I would then need to restrict the scope of my overall argument for the impossibility of relativized propositions so that it applies only to objectively settled relativized propositions. It should be clear then, that, in light of this restriction in scope, the resultant argument establishes not (2/9) but, instead, the following:

(2/9*) Necessarily, there are no objectively settled relativized propositions.

But note that the scope of the Argument from Tradition may similarly be restricted, but that in restricting the argument in this way we are not thereby forced to reinterpret its conclusion. So we still get the same result: namely, that (I) entails the truth of CFT and, consequently, the falsity of internalism (about representational mental content). (Note also that the proposition <that I have an apple on my head> -- if it exists – is objectively settled. Moreover, if internalism is true and thoughts inherit their content from propositions, then the proposition <that I have an apple on my head> exists and is a relativized proposition.)

snow is white are both attitudes toward the proposition <that snow is white>. So you and I both entertain this proposition. So this proposition is sharable. (Indeed, it is *shared*.) Now think of some proposition that, as far as you know, nobody besides you is entertaining, ever has entertained, or ever will entertain. But say also that the proposition you have in mind could have been entertained by some individual other than you. This proposition, then, is sharable.

Recall from §3.2 that a *private* proposition is a proposition that only one subject can entertain. Another way to put this is to say that a private proposition is a proposition that is not sharable.[26]

If (I) is true, then there are private propositions. In §6.1 I present – though do not explicitly endorse – an argument for the claim that egocentric beliefs are attitudes toward private propositions. The argument I present goes back to Frege (1956), the original proponent of private propositions (though, as I discuss below, in footnote 30, Frege himself puts things somewhat differently form how I put them here). In §6.2 I identify what, it seems to me, are the two *prima facie* plausible strategies for resisting the Fregean argument, as I conceive of it. I argue that, necessarily, if (I) is true, these strategies both fail. I conclude that, necessarily, if (I) is true, egocentric beliefs are attitudes toward private propositions. Since there are egocentric beliefs, it follows that, if (I) is true, there are private propositions.

[26] As we'll see in what follows, my notion of a private proposition is related to Perry's notion of "propositions of limited accessibility" (1979: 16), Hanks's notion of "first-person propositions" (2013: 155), and Merricks's notion of a "just-for-me proposition" (2015: 31, fn. 19). But, as we'll also see in what follows, there are important differences among these notions as well. (See footnotes 40, 44, and 45.)

6.1 – The Fregean Argument for Private Propositions

Recall from Chapter One, §5, that egocentric belief (i.e., first-personal, or *de se*, belief) is *essentially* egocentric in that it grounds – or at least can ground – egocentric *knowledge*. Consider, for instance, John Perry's story of Rudolf Lingens:

> An amnesiac, Rudolf Lingens, is lost in the Stanford library. He reads a number of things in the library, including a biography of himself, and a detailed account of the library in which he is lost. He believes [any proposition simply about how things are, objectively] you think might help him.[27] He still won't know who he is, and where he is, no matter how much knowledge he piles up, until that moment when he is ready to say,
>
> > *This* place is aisle five, floor six, of Main Library, Stanford.
> > I am Rudolf Lingens. (1977: 492)

Now say that Rudolf Lingens does come to have the egocentric belief that he himself is Rudolf Lingens, i.e., that he *self*-attributes the property of *being Rudolf Lingens*. And say that he comes to believe this in the right sort of way, such that in coming to believe it he comes to know it. Then he has essentially egocentric knowledge. And what makes this knowledge essentially egocentric is that this knowledge is essentially constituted by an egocentric belief state of a certain sort – in particular, an egocentric belief state in which he self-attributes the property of *being Rudolf Lingens.* He could not have had this knowledge by, for instance, having a non-egocentric belief in which he simply attributes, to some individual x, the property of *being Rudolf Lingens*, even if it

[27] Two things are worth noting here. First, the bracketed insertion replaces Perry's phrase 'Fregean thought'. I take it that, in general, Perry means, by 'Fregean thought', *proposition*. But I take it that, in the present context, Perry is restricting his use of 'Fregean thought' to propositions that represent, simply, how things are objectively. Second, the locution 'believe a proposition' is potentially misleading. It is not really the proposition that is believed. Indeed, strictly speaking, there is no *thing* that is believed. (See Searle 1983.) For simplicity's sake, I'll follow Perry and others in occasionally employing this locution, but all I will mean in saying – for instance – that so and so believes a certain proposition is that so and so has a certain belief state that is an attitude toward that proposition.

turns out that Lingens is x. Nor could he have had this knowledge by self-attributing some other property, e.g., the property of *being John Perry*.

The argument for private propositions I have in mind rests on four claims. The first two of these are (I) and (II), i.e., the two tenets of the traditional doctrine of belief, as I conceive of this doctrine. The third claim is what I illustrated in the above paragraph, namely:

(10) Egocentric belief is essentially egocentric.

The fourth and final claim is the following:

(11) Necessarily, if egocentric belief is essentially egocentric, this is because egocentric beliefs have essentially egocentric *content*.

I'll have more to say in defense of (11) in §6.2. But for the moment I'll note simply that, if egocentric belief is essentially egocentric, then – by definition – what one knows in virtue of having a true, warranted, egocentric belief state cannot be known by one's having a true, warranted, non-egocentric belief state. The most straightforward explanation for this is that egocentric belief states have a kind of representational content that non-egocentric belief states do not (and cannot) have. That is to say: egocentric beliefs have essentially egocentric content.

With all this in mind, consider, as an example of egocentric belief, Lingens's belief that he himself is Rudolf Lingens. It follows from (10) and (11) that:

(12) Egocentric beliefs have essentially egocentric content.

So Lingens's egocentric belief that he is Rudolf Lingens has essentially egocentric content, i.e., content that no non-egocentric belief can have. It follows from (I) that Lingens's belief is an attitude toward a proposition P such that (i) it is impossible for a non-egocentric belief to be an

attitude toward P and (ii), necessarily, any belief state that is an attitude toward this proposition will be one in which the subject of that belief state *self-ascribes* the property of *being Rudolf Lingens*. [28] For ease of reference, let '$\underline{P_L}$' be a name for this proposition, whatever it is. Now, say that $\underline{P_L}$ is shareable. Then, in addition to the fact that Lingens has an egocentric belief that is an attitude toward $\underline{P_L}$, there is a possible world wherein some individual other than Lingens has a belief that is an attitude toward $\underline{P_L}$. Surely, then, there is a single possible world W in which both Lingens and this other individual entertain $\underline{P_L}$. (It would be wholly unprincipled to deny this.) Now, this belief, in this possible world W, must be egocentric, for again it is impossible for a non-egocentric belief to be an attitude toward $\underline{P_L}$. But, necessarily, any belief state that is an attitude toward $\underline{P_L}$ is one in which the subject of that belief state *self-ascribes* the property of *being Rudolf Lingens*. But any such belief state either is false or is true and is a belief state of Lingens himself. It follows, given (I), that – on the assumption that there really is a possible world W in which $\underline{P_L}$ is entertained both by Lingens and by somebody other than Lingens – $\underline{P_L}$ is a relativized proposition. Of course (II) tells us that relativized propositions are impossible. So $\underline{P_L}$ is not shareable. So $\underline{P_L}$ is a private proposition. So Lingens's egocentric belief that he is Rudolf Lingens is an attitude toward – and, as such, inherits the content of – a private proposition. Of course, Lingens's belief here is just meant to be one example of an egocentric belief.[29] So, to generalize:

---

[28] This implication may seem obvious, but – given the precise way in which I've articulated (I) – the move from (I) to the existence of such a proposition turns out to be somewhat complicated. Again, I present (I) as follows:

> (I) Necessarily, if two thoughts share (representational) content (of some sort), then they inherit that content from a single proposition.

If (I) is true, it is true because, necessarily, a thought has the particular content it has in virtue of being an attitude toward a proposition with that content. It thus follows from our assumption of (I) that, necessarily, thoughts are attitudes towards the same proposition P – and, thereby, inherit the content of that proposition P – if *and only if* they have the same (representational) content as each other.

[29] To be fair, my argument regarding Lingens's egocentric belief is made simpler by the fact that, necessarily, only Lingens can be right in self-ascribing the property of *being Rudolf Lingens*. But it would

egocentric beliefs are attitudes toward – and, as such, inherit the content of – private propositions. And note that my argument for this claim is based on nothing contingent. So:

(13) Necessarily, egocentric beliefs are attitudes toward private propositions.

(13) entails that, if there are egocentric beliefs, there are private propositions. And, of course, there are egocentric beliefs – e.g., I believe that I am NKR. So there are private propositions.

6.2 – Objections to the Argument for Private Propositions

Again, the argument for private propositions I have in mind goes back to Frege (1956), the original proponent of private propositions.[30] There are precisely two ways of resisting my argument that, it seems to me, are at least *prima facie* plausible (but see footnote 35 for a minor qualification). Both of these ways of resisting the above argument are versions of objections leveled against Frege's initial 1956 argument. Of course, I present the objections here as they apply to the version of the argument I have presented.

One *prima facie* plausible way to resist the argument from §6.1 is to insist that egocentric beliefs are attitudes toward relativized propositions. But recall from §5 that (I) entails that

---

be unacceptably *ad hoc* to insist that the only private propositions are those one person can be correct in endorsing.

[30] Frege's initial argument differs from the above argument in two main respects. First, whereas the version of the argument presented above focuses on propositions simply as the things toward which our thoughts are attitudes, Frege's argument has to do with propositions both as the things toward which our thoughts are attitudes and as the things that sentences express (or that we express by uttering sentences). (See §7.2 for how this difference leads Frege to incur additional commitments regarding the nature of these propositions.) Second, Frege (1956) does not speak explicitly of an egocentric belief's being *essentially* egocentric. But the moves I present above that appeal to this conception of egocentric belief are accomplished by Frege's identification of a thought's content with its *sense*. (In this way, the version of the argument for private propositions that I present is mediated by discussions of the so-called "problem of the essential indexical" to which Frege's argument gives rise. See Lewis (1979), Perry (1979), and Stalnaker (1981) for classic discussions of this problem. See also Perry (1977) for a related discussion (and critique) of Frege's argument.)

relativized propositions are impossible. So (I) entails that this first way of resisting the above-presented argument for private propositions fails.[31]

The other *prima facie* plausible way to resist the argument from §6.1 is what I'll call "the guise strategy." To explain this strategy, I'll first present it as it applies to a related issue. Imagine that, at t_1, Lois Lane believes – here's how she'd put it – "Superman is strong" – even though she does not believe – here's how she'd put it – "Clark Kent is strong." But say that, at t_2, Lois learns that Clark Kent *is* Superman. Consequently, Lois comes to believes – here's how she'd put it – "Clark Kent is strong." Notice that Lois's subsequent Clark Kent belief state constitutes a *discovery*; she comes to realize Clark Kent is strong only upon realizing that Clark Kent is identical to Superman. Followers of Frege will say that, to make sense of this discovery, we must insist that Lois's two beliefs – that is, her initial Superman belief and her subsequent Clark Kent belief – are attitudes toward distinct propositions. But many philosophers will resist this conclusion. For many philosophers will insist that Lois's two belief states are *singular* representations in that each is directly about something (as opposed to being about something, but only in virtue of representing certain properties). Indeed, these philosophers will say that Lois's two belief states are directly about the same thing – namely, Superman/Clark Kent. Consequently, these philosophers (or at least some of these philosophers) will insist that Lois's two belief states are attitudes toward the *same* proposition, the *singular* proposition <that Superman / Clark Kent is strong>. These philosophers then owe us an explanation of how to account for Lois's discovery.

[31] Strictly speaking, what I establish in §5 is just that (I) entails that there are no objectively settled relativized propositions. (See footnotes 24 and 25.) But the proposition toward which Lingens's egocentric belief is an attitude is an objectively settled proposition. It follows from (I) that it is not relativized. (Again, a proposition is objectively unsettled just in case it purports to represent a portion of reality that, as a matter of fact, is objectively unsettled, i.e., is unsettled with respect to how things are, objectively speaking.)

The standard way of accounting for a discovery such as this is to characterize different knowledge states in terms of different guises under which a proposition may be entertained. To begin to see how this works, consider the following passage from Nathan Salmon:

> The mode of acquaintance by which one is familiar with a particular object is part of the mode of apprehension by which one grasps a singular proposition involving that object. For example, if one is familiar with some individual by having read his or her writings, then the reading of these writings is also part of the means by which one is acquainted with a singular proposition about that individual – say, that he or she had an unhappy childhood. (1986: 108)

For Salmon, the guise by which one entertains a given proposition can be identified with the "mode of apprehension" by which one grasps that proposition. Moreover, at least when the proposition in question is singular, we can entertain that proposition under different guises (1986: 107 – 109).

Salmon's picture provides us with a straightforward way of understanding what is special about Lois's knowledge that Clark Kent is strong, even granting that Lois's Superman belief and her Clark Kent belief are attitudes toward the same proposition. Knowing something, on the picture Salmon is proposing, is constituted not just by one's endorsing a particular true proposition (along with one's being warranted in endorsing this proposition) but – in addition to this – endorsing this proposition under a certain guise. This explains why, when Lois comes to believe that Clark Kent is strong, her coming to believe this constitutes a discovery. All along, Lois endorses (or at least is in a position to endorse) the proposition <that Superman / Clark Kent is strong). Her discovery amounts to her coming to endorse this proposition under a new guise. In particular, Lois's believing, at t_1, that Superman is strong amounts to Lois's endorsing, at t_1, the proposition <that Superman / Clark Kent is strong> under one guise – call it the "Superman guise"; her believing, at t_2, that Clark Kent is strong amounts to her endorsing this same proposition at t_2, but now under a

different guise – call it the "Clark Kent" guise.[32] And it is Lois's coming to endorse this proposition under the latter guise that constitutes Lois's discovery.[33]

This strategy for explaining how someone can discover something without coming to endorse a new proposition can be used to resist a key move in the argument for private propositions presented in §3.1. Recall from §3.1 that, since Lingens's egocentric belief is essentially egocentric, it constitutes essentially egocentric knowledge – in this case, Lingens's knowledge that he is Rudolf Lingens, where this knowledge is essentially egocentric in that it can only be had by one's self-attributing the property of *being Rudolf Lingens*. It was then argued that, to accommodate this condition, we must insist that Lingens's egocentric belief is an attitude toward a proposition that must be entertained first-personally, i.e., a proposition toward which only egocentric beliefs (or thoughts, more generally) can be attitudes. But the proponent of the guise strategy can resist this move. First, the proponent of the guise strategy can insist that Lingens's egocentric belief is an attitude toward a proposition that can be entertained third-personally. For ease of illustration, let us take this to be the proposition <that Lingens is Lingens>, i.e., the proposition toward which my belief that Lingens is Lingens is an attitude. The proponent of the guise strategy can then explain the essentially egocentric nature of Lingens's knowledge that he is Rudolf Lingens by insisting that this knowledge can only be had by one's entertaining that proposition under a certain guise – call it the (or a) *first-personal* guise, i.e., the (or a) guise wherein one thinks of oneself *as* oneself. Note that this explains why only Lingens can know that he is Lingens: while many of us can

[32] But isn't the meaning of a proper name fixed by its referent? If you think this (as those who say that propositions can be entertained under guises are likely to think), then replace 'Superman' and 'Clark Kent' with the definite descriptions (or primary intensions, if you prefer) of your choice.

[33] Salmon endorses this general strategy for explaining how we can come to know something new without coming to endorse a new proposition (1986: Ch. 8). Other proponents of roughly this strategy include Perry (1979), Kaplan (1989), Soames (1989), and Merricks (2015). See Perry (1979: 16 – 20), Kaplan (1989: 533 – 534), and Merricks (2015: 59 – 60) for versions of the application to egocentric belief I characterize below.

entertain the proposition <that Lingens is Lingens>, only Lingens can entertain this proposition under the first-personal guise; and this is what it takes for one to know that he is Rudolf Lingens.

I find this guise strategy *prima facie* plausible. Nevertheless, I think it fails. To see that it fails, we should think about which premise of the argument from §3.1 the guise strategy targets. The move from the essentially egocentric nature of Lingens's belief to the essentially egocentric nature of the proposition toward which this belief is an attitude is entailed by two claims:

> (I) Necessarily, if two thoughts share (representational) content (of some sort), then they inherit that content from a single proposition.

and

> (11) Necessarily, if egocentric belief is essentially egocentric, this is because egocentric beliefs have essentially egocentric *content*.

Now, for what it's worth, if the guise strategy presented above is at odds with (I), then – even if this strategy is successful as a response to the argument for private propositions – it does nothing to undermine my argument. For I am arguing only that the argument for private propositions presented in §3.1 is sound provided that (I) is true.[34] But set that aside. For it seems to me that the appropriate target of the guise strategy is not (I) but (11). For the guise strategy attempts to show us that we can explain what makes egocentric belief *essentially* egocentric without having to posit essentially egocentric content. It attempts to show us this by telling us two things, which I'll put

[34] Recall from §1 that Perry's version of the first tenet of what Perry calls "a traditional way of thinking of belief" is that "belief is a relation between a subject and an object, the latter being denoted, in a canonical belief report, by a that-clause" (1979: 5-6). It is this tenet that Perry targets in his own version of the guise strategy (or, perhaps, proto-guise strategy). This suggests that, as Perry conceives of this tenet, it is stronger than my [I]. Indeed, it seems to me that what Perry identifies as the first tenet of "a traditional way of thinking of belief" is roughly equivalent to the conjunction of [I] and (11), at least on the assumption that what this first tenet says of belief applies, *mutatis mutandis*, to thoughts in general.

in terms of our specific example involving Lingens's egocentric belief that he is Lingens. First, the guise strategy tells us that we should explain the essentially egocentric nature of Lingens's *knowledge* state as follows:

> **Guise Explanation.** The state of Lingens's (or anyone else's) knowing that he *himself* is Rudolf Lingens is essentially constituted, in part, by Lingens's entertaining a certain proposition about himself under the (or a) first-personal guise, i.e., the (or a) guise under which one thinks of oneself *as* oneself.

Second, the guise strategy insists that the above explanation is a genuine alternative to the following explanation:

> **Content Explanation.** The state of Lingens's (or anyone else's) knowing that he *himself* is Rudolf Lingens is essentially constituted, in part, by his having a belief whose content is essentially egocentric.

But, I shall now argue, the Guise Explanation is not an alternative to the Content Explanation. For the Content Explanation follows from the Guise Explanation. And so, I shall conclude, the guise strategy fails.

To see that the Content Explanation follows from the Guise Explanation, begin by assuming the Guise Explanation for conditional proof. The guise explanation entails that Lingens's state of knowing that he himself is Rudolf Lingens is essentially constituted, in part, by Lingens's thinking of himself *as* himself. Now, in general, to think of something as such and such is to think of something as being a certain way. For me to think of Michael Jordan as the greatest basketball player of all time is for me to think of Jordan as being a certain way – namely, as being the greatest basketball player of all time. For me to think of the narrator of James Joyce's "Araby" as "a creature driven and derided by vanity" is for me to think of the narrator of Joyce's "Araby" as being a certain way – namely, as being a creature driven and derided by vanity. Likewise, for

Lingens (or anyone else, for that matter) to think of himself as himself is for Lingens to think of himself as being a certain way – namely, as being himself. It follows that for Lingens to think of himself as himself is for Lingens to have a belief that represents him as being a certain way, where the way Lingens is represented as being is distinctive of egocentric belief. And the way a belief (or any other mental state) represents something as being is part of what determines, overall, the way that that belief represents things as being, i.e., the belief's representational content. So, insofar as Lingens's knowledge that he himself is Lingens essentially involves Lingens's thinking of himself as himself, Lingens's knowledge that he is Lingens is essentially constituted by a belief whose content is *essentially* egocentric. It follows that the Content Explanation is true. So, to discharge our assumption: the Content Explanation follows from the Guise Explanation. Consequently, the guise strategy against the argument for private propositions fails.

Now perhaps one might object that, when Lingens thinks of himself *as* himself, he is not thereby thinking of himself as being a certain way but, instead, is simply thinking of himself from a certain perspective – namely, from the *first-personal perspective*. Call this response to my argument "the perspective response."

The proponent of the perspective response faces a dilemma. For we might ask, "When Lingens thinks of himself from the first-person perspective, is his thinking of himself from this perspective essential to his coming to know something?" If the proponent of the response we're considering answers *No*, then she cannot insist – as she must insist – that an essential part of Lingens's knowing that he is Lingens is Lingens's thinking of himself as himself. Say, instead, that the proponent of the response we're considering answers *Yes*. Then the proponent of this response owes us an explanation of how it could be that simply taking up a certain perspective can play a distinctive role in one's coming to know something. And it seems to me that there is only

one plausible answer to this question – namely, that in taking up a certain perspective one comes to recognize some different aspect of something or, to put it differently, one comes to recognize something as being a certain way. The implication of the dilemma is this: if Lingens's thinking about himself from the first-personal perspective is essential to Lingens's knowing that he is Lingens, then Lingens's thinking about himself from this perspective just is – or, at least, constitutively involves – Lingens's thinking of himself *as being* himself. So the perspective response fails as a way of resisting my move from the Guise Explanation to the Content Explanation.

Since the Content Explanation follows from the Guise Explanation, the appeal to guises fails as a way of resisting the argument from the essentially egocentric character of egocentric *belief* to the essentially egocentric character of egocentric belief *content*. Moreover, this appeal to guises was the best hope for resisting this argument. Indeed, it was the best hope for resisting the overall argument for private propositions, at least on the assumption that (I) is true.[35] I conclude, then, that (I) entails the following:

(13) Necessarily, egocentric beliefs are attitudes toward private propositions.

So, if (I) is true, then private propositions exist and are the entities toward which our egocentric beliefs are attitudes.

[35] Besides the appeal to relativized propositions (which, again, (I) rules out), there is one other way of resisting the argument from private propositions that is somewhat prominent in the literature. Stalnaker (1981) and Cappelen and Dever (2013) argue, in effect, that we should deny (10), i.e., the claim that egocentric belief is essentially egocentric. For my part, the truth of (10) seems to me undeniable, given my own introspective access to my egocentric thoughts. Of course, Stalnaker and Cappelen and Dever provide sophisticated alternatives for how we should understand what is special about egocentric thought. While these alternatives all seem to me to miss something important about egocentric belief, a longer, more complete work would engage with these alternatives.

7 – Modus Ponens or Modus Tollens?

We began with the two tenets of the traditional view of belief:

(I) Necessarily, if two thoughts share (representational) content (of some sort), then they inherit that content from a single proposition.

(II) Necessarily, propositions (if they exist) have truth values, and they have their truth values absolutely.

It is fairly uncontroversial that, since (II) straightforwardly entails that relativized propositions are impossible, (I) and (II) together entail that CFT is true and, consequently, that externalism (about representational mental content) is true.

I have argued for something stronger. Like (II), (I) entails – all by itself – the impossibility of relativized propositions. So (I) entails – all by itself – that CFT is true and, consequently, that externalism is true. (I) also entails – all by itself – that, necessarily, egocentric beliefs are attitudes toward private propositions.

Now, from all I've said so far, one might be led simply to deny (I). For, again, while CFT is *prima facie* plausible in its own right, both externalism and the view that there are private propositions (conceived of as the entities toward which our egocentric beliefs are attitudes) face serious objections.

But there are compelling considerations in favor of (I). First, it is standardly held that propositions are necessary to explain what it is that different mental states – e.g., my belief that John is tall, and your supposition that John is tall – have in common. For this appeal to propositions to do any explanatory work, propositions must be the fundamental bearers of content, in which case (I) must be true.[36] Second, some philosophers argue that, in order to accommodate the validity

[36] See Cargile (1979: 104 – 105).

of certain arguments (and the invalidity of others), we must say that propositions exist necessarily and essentially represent whatever they represent.[37] But this is plausible only given the truth of (I-Cor), the claim that, necessarily, the fundamental bearers of (representational) content are propositions. Consequently, it is plausible only if (I) is true. Third, and relatedly, it seems difficult to accommodate the truths of propositional logic without characterizing these truths, straightforwardly, as being about propositions – e.g., for any proposition P and Q, if P materially implies Q, then Q's negation materially implies P's negation. And, plausibly, for the appeal to propositions to accommodate such truths successfully, we need to hold that propositions exist and represent whatever they represent essentially (if not necessarily). And to accommodate this we need to accept (I-Cor) and, consequently, (I).

For (at least some of) these reasons, I endorse (I). So I endorse CFT. So I endorse externalism (and reject internalism), and I hold that our egocentric thoughts are attitudes toward private propositions.

I owe my reader a response to the worries faced by externalists and by proponents of private propositions, conceived of as the entities toward which our egocentric beliefs are attitudes. I confine myself here to sketches of responses to what I consider to be the three most pressing objections to private propositions, along with the three standard objections to externalism discussed in Chapter One.

7.1 – Distinguishing between Private Propositions

Of the three objections to private propositions I'll consider, two are explicitly articulated by opponents of private propositions, and one – I suspect – underlies the initial, intuitive resistance to private propositions that philosophers sometimes seem to have.

[37] See Cargile (1979: 104 – 105), Williamson (2002: 235 - 236), and Merricks (2015: Ch. 1).

I begin with the latter. We often appeal to 'that' clauses to identify propositions and to distinguish one proposition from another. Say, for instance, that Jack believes in – and Jill denies – the existence of private propositions. We can distinguish between the propositions toward which their beliefs are attitudes by characterizing these propositions, respectively, as follows: the proposition <that there are private propositions>, toward which Jack's belief is an attitude, and the proposition <that there are no private propositions>, toward which Jill's belief is an attitude. In distinguishing between propositions in this way, it seems crucial that we grasp each proposition. But, then, we cannot distinguish between private propositions in this way since, by definition, private propositions are unsharable. This may not by itself be a reason to think that there are no such things as private propositions. But, the objection goes, it may plausibly lead us to doubt the rationality of believing in them.[38]

My response to this objection is to show that we *can* use 'that' clauses, slightly modified, in order to distinguish between propositions. To do this, we needn't come up with 'that' clauses that allow us to grasp each proposition. To grasp a proposition is to entertain that proposition directly, i.e., to have a thought whose content is (or includes) the content of that proposition. To distinguish between propositions by appeal to 'that' clauses, we needn't do this. It is enough for us to use 'that' clauses to think of propositions indirectly. Consider, for instance, the following proposition:

<that I_{RL} am Rudolf Lingens>

[38] According to Peter Hanks, private propositions – at least of the sort endorsed by Frege – "are widely regarded as mysterious and metaphysically suspect" (2013: 155 – 156). What he has in mind by this seems more closely related to the objection to private propositions I consider in §7.3. But I suspect that philosophers do often find private propositions "mysterious," precisely because we cannot distinguish between them in the standard way.

Here, 'RL' stands for Rudolf Lingens. In subscripting 'RL' to a first-person pronoun (here, 'I') within the 'that' clause, we are identifying the proposition as being one in which Lingens is the object of first-personal reference. The proposition <that I_{RL} am Rudolf Lingens>, then, is the proposition whose content is – or would be – inherited by any thought constituted, simply, by the self-ascription of the property of *being Rudolf Lingens* to Lingens. Consider, instead, the following proposition:

<that I_{JP} am Rudolf Lingens>

Here, 'JP' stands for John Perry. In subscripting 'JP' to a first-person pronoun (here, 'I') within the 'that' clause used to label the proposition, we are identifying the proposition as being one in which Perry is the object of first-personal reference. The proposition <that I_{JP} am Rudolf Lingens>, then, is the proposition whose content is – or would be – inherited by any thought constituted by the self-ascription of the property of *being Rudolf Lingens* to Perry. Note, then, that the proposition <that I_{JP} am Rudolf Lingens> is false whereas <that I_{RL} am Rudolf Lingens> is true.[39]

[39] Recall from Chapter One, §4, that representationalists hold that a thought's representational content is wholly a product of which possibilities the thought is consistent with, along with which possibilities the thought rules out. (I) entails that representationalism is coextensive with the analogous claim about propositions – i.e., the claim that the representational content of a *proposition* is wholly a product of which possibilities the thought is consistent with, along with which possibilities the thought rules out. It follows that the representationalist who accepts (I) has available to her an additional way of distinguishing between private propositions, along with propositions more generally.

Begin with the notion of a thought T's subjunctive and epistemic intensions, as characterized by Chalmers (2002a, 2002b). (See Chapter One, §6.1, and §6.2.) Extend this notion, *mutatis mutandis*, to propositions. A proposition P's subjunctive intension, then, is the function *f* from possible worlds to truth values such that, for any possible world W, *f*(W) = *true* if and only if P is true at W *considered as counterfactual* (and *f*(W) = *false* otherwise). A proposition P's epistemic intension is the function *f** from scenarios to truth values such that, for any scenario S, *f**(S) = *true* if and only if P is true at S, considered as actual (and *f**(S) = *false* otherwise). It follows from (I) that, necessarily, for any thought T and proposition P toward which T is an attitude, T and P have the same subjunctive intension as well as the same epistemic intension. Now, let a thought or proposition's dual intension be the ordered pair {*f*, *f**} such that *f* is the thought or proposition's subjunctive intension while *f** is the thought or proposition's epistemic intension. If representationalism is true, then every non-private proposition has a unique subjunctive intension. Moreover, if representationalism is true, then every proposition – including every private

7.2 – Private Propositions as Incommunicable

Perhaps the most common objection to private propositions is that they are incommunicable. Peter Hanks puts the objection as follows:

> There is…a problem about incommunicability. If you cannot grasp my first-person mode of presentation then you cannot grasp the propositions I express using 'I'. This is unacceptable. Anyone who speaks English should be able to grasp the proposition I assert when I say 'I am on fire'.

By 'first-person mode of presentation' Hanks has in mind the *sense* of an egocentric belief, which for Frege is identical to the private proposition one entertains in having that belief. While Hanks himself believes in private propositions, he does not believe in private propositions conceived as Fregean senses.[40] Setting Frege scholarship aside, we can say that – for Hanks – the problem of incommunicability arises when we conceive of private propositions as the entities whose contents our egocentric beliefs inherit. So the objection Hanks has in mind applies to private propositions as I conceive of them – namely, as the entities toward which egocentric beliefs are attitudes.

The objection, so understood, amounts to the following argument:

(14) A subject can communicate successfully by uttering (with assertoric force) 'I am on fire' to a competent English speaker.

---

proposition – has a unique dual intension. So the representationalist who accepts (I) can distinguish between private propositions, in terms of their representational contents, by appeal to these propositions' dual intensions. (Note that proponents of (I) who deny representationalism are still in a position to distinguish between *classes* of private propositions in this way.)

[40] Hanks believes in *first-person propositions*, where a "first-person proposition is a proposition that only a single subject can assert or believe" (2013: 155). So characterized, first-person propositions are private propositions since, at least in general, if one cannot believe a proposition, one cannot entertain it in any other sense (consider it, doubt it, etc.). (Note that it does not follow that private propositions are first-person propositions. See fn. 44.) Hanks's project is "to show how to make sense of first-person propositions without relying on first-person senses or anything in the vicinity, such as individual essences or haecceities" (2013: 156).

(15) A successful 'I am on fire' communication consists in the following: first, the subject asserts the proposition that she entertains in self-ascribing the property of *being on fire*; second, the other speaker grasps that proposition, where to grasp that proposition is – at the very least – to entertain it.

(16) If egocentric beliefs are attitudes toward private propositions, then the proposition a subject entertains in self-ascribing the property of *being on fire* is a private proposition, and as such it is a proposition that nobody else can entertain.

Therefore,

(17) It is not the case that egocentric beliefs are attitudes toward private propositions.

This argument is valid, and (14) and (16) are beyond dispute. The problem with this argument, I'll argue, is with (15).

To clarify, in premise (15), by "successful 'I am on fire' communication", I mean an *event* consisting of the following: first, an 'I am on fire' assertion, i.e., an utterance of 'I am on fire', made with assertoric force; second, successful uptake, on behalf of a listener, of the speaker's intended message. (Note that it is not true by definition that an 'I am on fire' assertion involves the assertion of a proposition from one individual to another.)

With this in mind, consider the following (in)famous passage from Frege's "The Thought: A Logical Inquiry":

> Now everyone is presented to himself in a particular and primitive way, in which he is presented to no-one else. So, when Dr. Lauben thinks that he has been wounded, he will probably take as a basis this primitive way in which he is presented to himself. And only Dr. Lauben himself can grasp thoughts determined in this way. But now he may want to communicate with others. He cannot communicate a thought which he alone can grasp. Therefore, if he now says "I have been wounded", he must use the "I" in a sense which can be grasped by others, perhaps in the sense of "he who is speaking to you at this moment", by doing which he makes the associated conditions of his utterance serve for the expression of his thought. (1956: 298)

Recall that by 'thought' ('Gedanke') Frege means, roughly, what I mean by 'proposition'. So Frege is endorsing belief in propositions that are both private – i.e., unsharable – and *incommunicable*, where a proposition is incommunicable – let us say – just in case it cannot be asserted by one subject and grasped by another. Frege is also providing us with a way of resisting the Argument from Communication. The response, following Frege, is to deny (15) by insisting that – in a successful 'I am on fire' communication – the proposition that the speaker is asserting and that the listener is grasping is *not* the proposition that the speaker entertains (or would entertain) in self-ascribing the property of *being on fire*. Instead, the proposition in question is (something like) the proposition that the speaker entertains (or would entertain) in thinking – here's how he'd put it – "He who is speaking to you at this moment is on fire." Or so the Fregean response goes.

The Fregean response, it seems to me, is only half right. I follow the Fregean response in denying that – in a successful 'I am on fire' communication – there is a single proposition P such that (i) the speaker asserts P and the listener grasps P and (ii) P is the proposition that this speaker entertains (or would entertain) in self-ascribing the property of *being on fire*. But I deny the Frege-inspired alternative. That is, I deny that – in a successful 'I am on fire' communication – there is a single proposition P such that (i) the speaker asserts P and the listener grasps P and (ii*) P is the proposition that this speaker entertains (or would entertain) in thinking – here's how he'd put it – "He who is speaking to you at this moment is on fire." The problem with this proposed alternative is that the proposition in question is another private proposition. For second-person thoughts are really first-person thoughts in disguise. For instance, when I think to myself – here's how I'd put it – "He who is speaking to you at this moment is on fire," I am thinking – here's another way I could put it – "He who is speaking to the individual to whom I speaking is on fire." So my belief

is an egocentric belief. Consequently, my belief is an attitude toward a private proposition. This proposition, then, cannot be grasped by the person to whom I am speaking.

The problem here cannot be avoided by amending the Fregean response so that the proposition being asserted is a non-private proposition. For no non-private proposition will allow for successful communication in this case, at least not if egocentric beliefs really are attitudes toward private propositions. For a successful 'I am on fire' communication somehow involves a listener's receiving information that is not simply third-personal, that the listener could not learn, for instance, just in knowing everything that is objectively the case about the world. At the very least, this requires that the proposition being asserted be one that the speaker asserts – and the listener grasps – by appeal to indexical concepts. If my argument for private propositions above is sound, then such propositions are private.[41]

What this reveals is that the friend of private propositions should deny that – in a successful 'I am on fire' communication – there is *any* proposition P such that (i) the speaker asserts P and the listener grasps P. Rather, in a successful 'I am on fire' communication, there is a proposition P and a proposition P* such that (i*) the listener grasps P as a result of the speaker's 'I am on fire' assertion, and (ii**) the speaker's 'I am on fire' assertion is essentially preceded by the speaker's self-ascribing the property of *being on fire*, where P* is the proposition that the speaker entertains in self-ascribing this property. Premise (15), then, is false. Or so the friend of private propositions should say.

And so I say. To lay out the details of my proposal, I'll focus on the following example of a successful 'I am on fire' assertion:

[41] Or, even if they are not private, they are propositions of "limited accessibility." See §7.3. (My own view is that all such propositions really are private propositions.)

> **Fire.** Vi is on fire. Vi sees an individual – a stranger to her – who is nearby and who has a bucket of water. Vi utters, with assertoric force, 'I am on fire', and her utterance is directed at the stranger. There is successful uptake of Vi's message (whatever that is) on behalf of the stranger.

Implicit in the above scenario is that Vi believes – and is right to believe – that she is on fire. Her belief is an attitude toward the (true) private proposition <I_{Vi} am on fire>. And, as will soon become clear, it is essential to Vi's 'I am on fire' assertion that it is occasioned by Vi's having this particular belief.[42] Vi's 'I am on fire' assertion is partly constituted Vi's having and acting upon a certain intention – an intention for the stranger to grasp a certain proposition, where in this case the stranger's grasping the proposition constitutively involves his *believing* it. But the proposition Vi intends for the stranger to believe is not the proposition <that I_{Vi} am on fire>. Vi has no hope of getting the stranger to believe this proposition. Nor would it be helpful for Vi if the stranger did – *per impossibile* – believe this proposition. What Vi intends is for the stranger to attribute to her whatever property she self-attributes in believing the proposition that <that I_{Vi} am on fire>. Vi recognizes that her belief in the proposition that <that I_{Vi} am on fire> is an egocentric belief, one that amounts to her self-ascribing the property of *being on fire*. So, she concludes, what she must do is get the stranger to ascribe to her the property of *being on fire*. Vi also recognizes that she is (or is about to) address the speaker. So, she concludes, what she must do – here's how Vi would put it – "is to get the stranger to self-ascribe the property of *being addressed by an individual who is on fire*." One norm of speech is that, for any property *F-ness,* if a speaker wants a listener to self-ascribe the property of *being addressed by an individual who is exemplifies F-ness,* the speaker ought to utter the sentence 「I am F」. A correlative norm of speech is that, when a speaker utters a sentence of the form 「I am F」, the listener ought to self-ascribe the property of *being addressed*

[42] This is not quite right. What is essential is that this belief precedes the assertion in those cases in which the assertion is sincere.

by an individual who exemplifies F-ness. (These norms, of course, have developed in relation to one another.) In accordance with the former norm, Vi utters 'I am on fire'. In accordance with the latter norm, the stranger self-ascribes the property of *being addressed by an individual who is on fire.* In other words, for the individual x such that x is identical to the listener, the listener comes to believe the private proposition <that I_x am being addressed by an individual who is on fire>. So there is successful uptake of the speaker's message.[43]

Notice that this explanation for where (15) goes wrong has important implications for how we should understand assertion. It is common to think of an assertion, more specifically, as an assertion of a proposition from one individual to another. My proposal rejects this view of assertion. On my view, assertion amounts to a speaker's acting on certain norms of speech, along with a listener's acting on another norm of speech, where both norms are designed to get the listener to believe a certain proposition P. Moreover, the speaker intends for the listener to believe P precisely because the speaker wishes to convey, at the very least, part of what she believes. But what the speaker believes needn't be P. It need only be a proposition related to P in the appropriate way.[44]

[43] Likely, this oversimplifies things, by failing to acknowledge the way in which successful communication involves participants' mutual recognition of their intentions. But the account here can be complicated to accommodate this. (See Grice 1957.)

[44] This way of understanding assertion has important implications for the relationship between my notion of a private proposition and both Hanks's notion of a first-person proposition and Merricks's notion of a just-for-me proposition. Again, for Hanks, a first-person proposition "is a proposition that only a single subject can assert or believe" (2013: 155). Hanks's first-person propositions count as private propositions since they are unsharable. The question of whether private propositions count as first-person propositions, in Hanks's sense, depends on whether, by 'only a single subject', Hanks means *at most one single subject* or *exactly one single subject*. If he means the former, then private propositions are first-person propositions. If he means the latter, then private propositions are not first-persons and – if I am right – first-person propositions are impossible. For Merricks, meanwhile, "a just-for-me proposition is a proposition that, necessarily, exactly one person can express" (2015; 31, fn. 19). If to express a proposition is to express it *to another*, then propositions – or at least first-person propositions – cannot be expressed. If so, then – as Merricks himself claims – just-for-me propositions are impossible. But, then, private propositions are not just-for-me propositions.

7.3 – Private Propositions as Not Worth the Cost

In the "Problem of the Essential Indexical," John Perry objects to what he calls "propositions of limited accessibility," which – for our purposes – we can think of as propositions that can only be entertained at one time, or at one place, or by one person.[45] While Perry's notion of a proposition of limited accessibility is related to my notion of a private proposition, these notions are also importantly different. Nevertheless, Perry's objection to propositions of limited accessibility can be modified to apply to private propositions. So modified, the objection runs as follows:

> Such a theory of [private propositions] seems acceptable, even attractive, to some philosophers. Its acceptability or attractiveness will depend on other parts of one's metaphysics; if one finds plausible reasons elsewhere for believing in a universe that has, in addition to our common world, myriads of private perspectives, the idea of [private propositions] will fit right in. I have no knockdown argument against such propositions, or the metaphysical schemes that find room for them. But I believe only in a common actual world. And I do not think the phenomenon of essential indexicality [and essential *egocentricity*, in particular] forces me to abandon this view. (1979: 16)[46]

Perry's objection – as I understand it – is a common one, which amounts to the following:

> First, (a) belief in private propositions comes at a high price. Second, the value of belief in private propositions is not all that great. For, to begin, (b) there is no good reason to believe in private propositions other than the fact that belief in such propositions provides us with a way to accommodate the phenomenon of essential egocentricity, and (c) there are plausible alternative ways of accommodating this phenomenon. So belief in private propositions is just not worth the cost.[47]

---

[45] Perry defines propositions of limited accessibility as "propositions which can only be expressed in special circumstances" (1979: 16). See fn. 44 for why, strictly speaking, there can be no such propositions, even though there can be private propositions.

[46] Egocentric beliefs are paradigm cases of indexical beliefs – beliefs that are constituted, in part, by the exercise of indexical concepts (e.g., *I*, *here*, and *now*). Just as the egocentric beliefs are essentially egocentric, indexical belief – more generally – is essentially indexical.

[47] Perry also raises roughly this objection in his (1977: 490). Additional proponents of this objection include David Kaplan (1989: 533 – 534), Colin McGinn (1983: Ch. 5), and Neil Feit (2008: 13).

I shall respond to parts (a), (b), and (c) of this objection in reverse order.

Response to (c). The whole point of §6 was to argue that accepting (I) commits us to private propositions. Now, there are ways to resist this move, but I have shown – or at least I take myself to have shown – that the most plausible ways to resist the move fail. I conclude, then, that – since there are very good reasons to accept (I) – there are very good reasons to accept private propositions.

Response to (b). Note that my response to (c) is really enough to resist the charge that the value of belief in private propositions is not all that great. But there's more to be said here. One of the most pressing objections to externalism is best met by the acceptance of private propositions. Perhaps at this point in the chapter it is already obvious that this is the case. Nevertheless, I shall demonstrate this explicitly in §7.4, when I respond to some standard objections to externalism.

Response to (a). Why take it that belief in private propositions comes at such a high price? Read straightforwardly, the above passage by Perry does present us with such a reason, but one that – it seems to me – is deeply confused. Again, according to Perry, belief in private propositions commits us to a "metaphysics" according to which the "universe…has, in addition to our common world, myriads of private perspectives" (16). Read straightforwardly, Perry's claim here is that belief in private propositions commits us to the view that reality itself is inherently subjective in nature – that it somehow consists of multiple universes each of which has an irreducible perspective built into it. Now, I agree that commitment to such a picture of reality would be costly, to say the least. But the believer in private propositions needn't endorse such a picture. The believer in private propositions need only accept that the way each of us represents reality as being is itself irreducibly perspectival.

But perhaps there is a less metaphysically loaded way of reading Perry's charge. Perhaps Perry's charge just amounts to this: belief in private propositions commits one to the view that the way each of us represents reality as being is itself irreducibly perspectival; and, at least pre-theoretically, this idea just seems implausible.

My response to this charge centers on Frege's claim, discussed in §7.2, that "everyone is presented to himself in a particular and primitive way, in which he is presented to no-one else" (1956: 298). Frege's statement, it seems to me, is really just a more poetic way of articulating the view that the way each of us represents reality as being is irreducibly perspectival. Consequently, we might interpret the charge of pre-theoretic implausibility that this view faces as a charge against Frege's statement.[48] So understood, the charge seems misguided. For, to begin, anyone who believes that egocentric beliefs are essentially egocentric should believe that everyone is presented to himself in a *particular* way, in which he is presented to nobody else. According to Frege and myself, this particular way in which each of us is presented to ourselves is primitive in the sense that it is constituted by our having egocentric thoughts that are attitudes toward private propositions. According to, e.g., Perry, this particular way in which each of us is presented to ourselves reduces to something like the following: my having an egocentric belief is constituted by my entertaining a non-private proposition P, in a distinctively first-personal way, where only I am able to entertain P in a first-personal way (even though others can entertain P in different ways, and even though others can entertain other propositions in first-personal ways). Now, I have argued against this reductionist strategy in §6.2. But, whether or not you were persuaded by my arguments in §6.2, you ought at least to agree with the following: it is not *pre-theoretically* obvious that the reductivist picture is the right one whereas the primitivist picture is the wrong one. Indeed, the

[48] Indeed, Peter Hanks (2013: 155) characterizes the objection I am considering at present as an objection to Frege's statement.

debate between the reductivist and the primitivist is just not the sort of thing that it would be reasonable for one to have a pre-theoretic stance on. So the charge of pre-theoretic implausibility we are considering fails.

Now, perhaps there is more to be said in defense of the claim that accepting private propositions comes at a high cost. But let me now respond to this claim in a way that will serve as an overall response not just to the objection I have been considering in the present section but, even more broadly, to resistance to private propositions in general. I have shown – and shall continue to show, in §7.4 – that we have very good reason to accept private propositions. In §7.1 and §7.2 I considered a couple of reasons – or at least candidate reasons – to reject private propositions. I argued against these reasons. Now, the claim that acceptance of private propositions comes at a high cost does not itself count as a reason against private propositions unless it is supported by some independent reason to resist private propositions. So the objection being considered in the present section does not give us any additional reason to doubt the existence of private propositions. So we are left with good reason to accept private propositions and no good reason to deny them. So, absent any additional reasons to deny private propositions, we ought to accept them.

7.4 – Objections to Externalism

It will be helpful to recall, from Chapter One (§1 & §3), the exact way I characterize internalism and externalism. To begin, *narrow* mental content – if there is such a thing – is mental content that is fully determined by a thinker's *internal* features, where a feature counts as internal to a thinker just in case it is a *qualitative* feature and is *intrinsic* to the thinker. *Wide* mental content – if there is such a thing – is mental content that is not narrow, i.e., mental content that is not fully determined by what is internal to a thinker. *Internalism* (about representational mental content) is the view that

there is narrow mental content. *Externalism* (about representational mental content) is the view that mental content is, in general, wide. (Again, unless I specify otherwise, by 'content' I mean *representational* content.)

The first objection I'll discuss has to do with our behavior. Intuitively, we should expect two thinkers' total mental states to effect the same behaviors only insofar as these total mental states are the same with respect to the ways in which they represent things as being. And, intuitively, to accommodate this we must say that representational mental content is narrow since, e.g., Oscar and Twin Oscar behave in exactly the same way.

The problem with this objection is that it does not take into account the full variety of mental states that make up a thinker's total mental states. The total mental state of a thinker – at least when that thinker is engaged in behavior – includes not just beliefs but desires and intentions. Say that, at time t, both Oscar and Twin Oscar perform a certain behavior that each would characterize as "drinking a glass of water." We might explain the way in which Oscar's mental states give rise to his behavior as follows: Oscar has the desire to drink a glass of water, and Oscar believes (truly, let's say) that the glass in front of him is filled with water; so Oscar forms the intention to drink the water from this glass, and he then acts upon this intention. Likewise, we might explain the way in which Twin Oscar's mental states give rise to his behavior as follows: Twin Oscar has the desire to drink a glass of twin water, and Oscar believes (truly, let's say) that the glass in front of him is filled with twin water; so Oscar forms the intention to drink the twin water from this glass, and he then acts upon this intention. Are Oscar and Twin Oscar the same in terms of their behavior? How we answer this question will depend on how we individuate behavior. If we individuate behavior widely, then the objection never gets off the ground. If we individuate behavior narrowly, then Oscar and Twin Oscar are the same in terms of their behavior, in which

case it is false that two thinkers' total mental states effect the same behaviors in them only insofar as these total mental states are the same with respect to the ways in which they represent things as being. Either way, the objection fails.

The second objection to externalism I'll discuss has to do with *privileged access*. Each of us has privileged access to the representational contents of our mental states, in the sense that each of us is in a unique position to tell, by introspection alone, the ways in which our various mental states represent things as being. But, arguably, if the ways our mental states represent things as being were determined even in part by factors external to us, we would not be in this unique position. For instance, there seems to be nothing available by introspection to Oscar that is not available by introspection to Twin Oscar, and *vice versa*. So, arguably, if the ways that Oscar's and Twin Oscar's mental states represent things as being were determined even in part by factors external to them, then neither Oscar nor Twin Oscar would be in a unique position to tell, by introspection alone, the way in which his mental state represents things as being.

For the moment, I have nothing new to say in response to this worry. But there is a standard response to the worry, which I endorse. Imagine that Oscar asks himself – here's how he'd put it – "Am I thinking about water?" Even given externalism, Oscar is in a position to answer this question. For Oscar is aware that he is thinking about a certain kind of substance – whatever substance, as a matter of fact, is wet and drinkable and fills the oceans and lakes around him. He thinks to himself – employing a demonstrative concept – "I am thinking about *that substance*." In thinking this to himself, Oscar is picking out a certain substance – namely, water – and he is ascribing to that substance the property of being a thing that he is thinking about. So Oscar is in a position to realize that he is thinking about water. And note that only Oscar is in a position to employ this particular, sure-fire method in coming to realize that he is thinking about water. So

there is a sort of privileged access to his mental states that Oscar has. And the standard externalist response – which I endorse – is to insist that this sort of privileged access is all we have reason to think we have.

The third and final objection to externalism I'll consider says that, arguably, we need to reject externalism to accommodate the fact that we are, by and large, rational cognitive agents. Recall that Oscar lives at a time before scientists have discovered the chemical make-up of water. With this in mind, imagine that Oscar thinks to himself – here's how he'd put it – "Water is wet." Furthermore, let's take for granted that Oscar is rational. Even still, Oscar is surely not in a position to infer that – here's how he'd put it – "H_2O is wet." But, the objection goes, the externalist is committed to saying that Oscar is in a position to make this inference. For, the objection goes, the externalist must say that Oscar's "water" thought is no different in content from the corresponding "H_2O" thought he might have.

In response to this objection, I deny that the externalist must say that Oscar's "water" thought is no different in content from the corresponding "H_2O" thought he might have. Now, to be fair, in Chapter One, in my initial illustration of the distinction between externalists and internalists, I did often treat thoughts of these sort as the same in content, given externalism. But I did this just to avoid complexity. And I also did this because, as a matter of fact, it is *standard* for the externalist to treat thoughts of this sort as being the same in content. But the externalist *needn't* treat thoughts of this sort as being the same in content. Indeed, I deny that thoughts of this sort have the same content. For instance, I deny that Oscar's thought that – here's how he'd put it – "water is wet" has the same content as the corresponding "H_2O" thought would have.

I owe my reader an externalist-friendly account of what this difference is. Begin with Oscar's belief that – here's how he'd put it – "water is wet." Oscar's belief is an egocentric belief

of a certain sort. In particular, it is an attitude toward (something like) the following private proposition:

<that the clear, drinkable liquid in the rivers and lakes in my_O environment is wet>,

where 'O' picks out Oscar. On the other hand, say that Oscar were to go on to believe, for whatever reason, that – here's how he'd put it – "H_2O is wet." This belief is an attitude toward (something like) the following proposition:

<that the substance that is composed of molecules consisting of two parts hydrogen and one part oxygen is wet>

It is metaphysically possible – as well as epistemically possible, for that matter – that the former of these propositions is true while the latter is false. So the former of these propositions does not entail the latter. Consequently, even assuming that Oscar is rational, we shouldn't expect him to be able to infer that – here's how he'd put it – "H_2O is wet" simply from his belief that – here's how he'd put it – "water is wet."

Note further that my analysis of Oscar's "water" belief – along with Oscar's would be "H_2O belief" – is an externalist analysis. For, given my analysis of Oscar's "water" belief, I ought to characterize *Twin* Oscar's corresponding "water" belief as being an attitude toward the following private proposition:

<that the clear, drinkable liquid in the rivers and lakes in my_{TO} environment is wet>,

where 'TO' picks out Twin Oscar. This proposition is distinct from – and differs in its representational content from – the proposition toward which Oscar's "water" belief is an attitude, i.e., the proposition <that the clear, drinkable liquid in the rivers and lakes in my_O environment is

wet>. Consequently, at least given (I), Oscar's "water" belief and Twin Oscar's "water" belief have distinct representational contents. This entails externalism since, by stipulation, Oscar and Twin Oscar are exactly similar with respect to their intrinsic, qualitative features.

Conclusion

The first tenet of the traditional view of belief – that is, claim (I) – says that, necessarily, if two thoughts share (representational) content (of some sort), then they inherit that content from a single proposition (I). (I), all by itself, entails that relativized propositions are impossible, that content fixes truth (CFT), that externalism (about representational mental content) is true, and that egocentric beliefs are attitudes toward private propositions. Since I endorse the first tenet of the traditional view of belief, I endorse all of these claims that – I have argued – this first tenet entails. I already have begun the work of defending these positions against some of their most pressing objections. My defense will continue in the final section of Chapter Three. But, first, to set myself up for that, I shall devote the majority of Chapter Three toward drawing out an implication of CFT and, consequently, of (I) – namely, that we cannot think about nonexistents.

Three

Thinking as Relational

The most natural way to characterize my thinking about something is as the instantiation of a relation – call it the *thinking about* relation – between me and whatever I am thinking about. But this way of thinking about thought faces a serious worry: it seems to imply that we cannot think about nonexistents.

In what follows I argue that, nevertheless, the relational way of thinking about thought is the right one. For the relational way of thinking about thought follows, I argue, from Content Fixes Truth (CFT), the claim that, necessarily, if two thoughts share representational content (of some sort), then they have the same truth value. I conclude by considering how, if at all, we can make sense of our ability – or our supposed ability – to think about nonexistents.

1 -- *Thinking as Relational* and Thinking about Nonexistents

Thinking as Relational (TR) is the following claim:

> **TR** Necessarily, whenever a thinker thinks about something, her doing so *just is* her standing in the *thinking about* relation to that thing.[1]

And by 'just is' I mean *is numerically identical to*. I do not mean *is grounded in* or *is explained by*. Nor do I mean *is conceptually analyzable as*. TR says that, necessarily, whenever a thinker thinks about something, the event of her doing so is numerically identical to the event of her standing in a certain relation – in particular, the *thinking about* relation – to that entity.[2]

[1] To qualify, by 'the thinking about relation', I really mean *the or a thinking about relation.* For the moment we can set this complication aside. I discuss it explicitly in §8.

[2] TR may *seem* to presuppose an events ontology. But TR is not about events, *per se*. It's just that – given the subject matter – events talk provides us with a way of characterizing TR that (a) is fairly straightforward and (b) is available both to proponents of TR and to those who deny it.

We can thus further characterize TR by appeal to the following examples, each involving a paradigmatic case of one's thinking about something. Presently, I'm thinking about my neighbor's goat, Goofus. Specifically, I'm having the belief that – here's how I'd put – "Goofus really loves carrots." TR says that this particular instance of my thinking about Goofus *just is* an instance of my standing in the *thinking about* relation to him. Earlier today, my friend Matthew was lamenting the fact that – here's how he put it – "the current prime minister of the UK is a Tory." In so doing, he was thereby thinking about Theresa May. TR says that Matthew's doing so just is (or just was) his standing in the *thinking about* relation to May. Possibly, my neighbor McCailin is wondering – here's how she'd put it – "whether any members of the Richmond Flying Squirrels will be called up to Fresno by the end of the season." In a possible world in which McCailin wonders this, she is thereby thinking about all of the current members (in that possible world) of the Richmond Flying Squirrels.[3] TR says that her doing so just is her standing in the *thinking about* relation to each of member. (Notice, then, that TR is not just limited to cases of someone's thinking *directly* about something.[4])

Now, to clarify, in speaking of *something* about which a subject is thinking, I do not *mean*, as a matter of definition, some *entity*. Nor, in speaking of *the thing* that a given subject is thinking about, do I mean, *per say*, *the entity* that the subject is thinking about. Nor do I mean *whatever falls under the existential quantifier*. I just mean *whatever the subject is thinking about*. So I use terms like 'something' and 'thing' in an ontologically noncommittal way. Now, perhaps, one

[3] Or at least this is the case in those possible worlds in which she thinks this and in which there is such a team as the Richmond Flying Squirrels. I address this complication shortly.

[4] Again, I think about something *indirectly* just in case I think about it *via* thinking about certain of its features (or purported features). And I think about something *directly* just in case I think about it but not *via* thinking about certain of its features (or purported features). To be absolutely clear that McCailin is thinking about the members of the Squirrels indirectly, let's imagine that she hasn't ever seen (or otherwise perceived) any of them or heard them mentioned by name.

cannot think about something without thinking about something that *exists*, where to exist – let us stipulate – is just to have being in some sense, i.e., to be an entity. In fact, later in this chapter, I'll argue for this position. I use 'thing' and 'something' in an ontologically noncommittal way in part to ensure that, in arguing for this position, I am arguing for something substantial.[5]

With this in mind, we can further illustrate TR by appeal to the following example. Say that, on Christmas Eve, my eight-year-old cousin Cooper thinks to himself – here's how he'd put it – "Santa Claus is coming to town." Cooper would say, if you asked him, that he is thinking about Santa Claus and that Santa Claus exists. Now, say that Cooper is right about the first point. Then TR says that, regardless of whether Cooper is right about the second point, Cooper's thinking about Santa Claus *just is* his standing in the *thinking about* relation to Santa Claus.

A final point of clarification here has to do with the term 'relation'. Consider the sentence 'Cooper is thinking about Santa Claus'. The surface grammar of this sentence suggests that Cooper is standing in a certain relation – in particular, the *thinking about* relation – to Santa Claus. Perhaps there is a sense of 'relation' according to which, if the preceding sentence is true, then it is automatically true that Cooper is standing in the *thinking about* relation to Santa Claus. Call this the *pleonastic sense* of 'relation'. (See Crane 2013: §3.4.) But surface grammar alone does not reveal to us anything about the underlying ontological structure of whatever state of affairs (if there is such a state of affairs) the preceding sentence represents (or purports to represent). Accordingly, it does not tell us whether Cooper stands in the *thinking about* relation (or some other

[5] One way of reading Quine (1961) is to say that, for Quine, we cannot use terms like 'thing' and 'something' to make such substantial claims, at least not insofar as we intend our interlocutors to interpret us in the way we intend to be interpreted. I am not sure this is the right way of reading Quine (1961). But, say that it is. Then Quine is mistaken. For surely you can understand me when I say that, by 'thinks about something', I do not mean 'thinks about some entity'. Now perhaps there is no theoretically neutral way to capture what I am saying in first-order logic. Fortunately, I don't need to be able to do that for you to understand me.

relation, for that matter) in the *substantial* sense of 'relation'. In the substantial sense of 'relation', to say that Cooper stands in a certain relation to Santa is to say that there is a way things are, objectively speaking, that cannot fully be accounted for simply by Cooper's exemplifying a property F or standing in a relation R, unless Cooper's doing so also involves (or occurs in concert with) *Santa's* exemplifying some property G or relation T. (Perhaps, in this case, Cooper's exemplifying F and Santa's exemplifying G together amount to Cooper's standing in an internal relation to Santa Claus. Or perhaps Cooper simply stands in an external relation to Santa. See §6.) Here and going forward, I shall use 'relation' in the substantial sense. Moreover, in my above articular of TR, I am using 'relation' in the substantial sense.

This above clarification should help us to see how strong a claim TR is. For consider *Serious Actualism*. Serious Actualism (SR) is the following claim:

> **SR** Necessarily, whatever stands in a relation (or otherwise exemplifies a property) exists.[6]

Again, I am using 'relation' in the substantial sense of the term. With this in mind, Serious Actualism seems particularly hard to deny. To deny Serious Actualism, one would have to say that – at least possibly – reality is a certain way that can be accounted for only by the exemplification of a property by something that is not part of reality in any sense. To put it differently, to deny Serious Actualism one would have to insist that – at least possibly – reality somehow includes, not just the property exemplifications of what is part of reality, but also the property exemplifications of what is not part of reality. But this seems incoherent. For it seems that the only way to make

[6] Since relations are just polyadic properties, Serious Actualism is equivalent to the claim that, necessarily, whatever exemplifies a property exists. Since I will be focusing on cases involving relations, I will often speak of Serious Actualism just as the claim involving relations. I speak this way just to avoid verbosity.

sense of something's being part of reality (i.e., its being an entity, its existing) as opposed to its not being part of reality is to say that reality includes that things' exemplifying properties.

Given Serious Actualism, TR implies that we cannot think about nonexistents. For say that TR is true. Then it is impossible to think about something without standing in a relation to it. Consequently, given Serious Actualism, it is impossible to think about something unless that something exists. So, for instance, if TR and Serious Actualism are both true, then either Santa Claus exists or we can't think about Santa Claus.[7]

It should not be surprising, then, that many philosophers deny TR (or something weaker than TR) precisely because they accept Serious Actualism. Among these philosophers are Franz Brentano (1874: 271), Richard Cartwright (1960: 639), Alvin Plantinga (1983: 10), John Searle (1983: 17-19), Kent Bach (1987: 12), Uriah Kriegel (2007), Katalin Farkas (2008), Jodi Azzouni (2010: 44), Tim Crane (2012a, 2012b, 2013), John Hawthorne and David Manley (2012), Trenton Merricks (2012: 74, 2015), and Terence Horgan (2013: 234-235).[8]

But some philosophers accept TR. For instance, Alexius Meinong (1904: 76-77), Richard Sylvan (1980: 674), Kit Fine (1982: 97), and Graham Priest (2005: 57) all accept TR but deny Serious Actualism. Other philosophers – among them Héctor-Neri Castañeda (1972) and Edward

[7] The idea that there is an inconsistency between a relational conception of thought and our ability to think about nonexistents goes back at least to Franz Brentano (1874).

[8] Some philosophers on this list may seem to oppose TR only given certain disambiguations of TR. Searle (1983: 17-19), for instance, distinguishes between extensional content and intensional-with-an-s content. Searle would say that the utterance by which I characterize TR is true provided that by 'thinks about something' I mean *has as the extensional content of one's thought*, but he would say that the utterance is false provided that by 'thinks about something' I mean *has as the intensional-with-an-s content of one's thought.* Similarly, Azzouni (2010: 44) thinks that 'thinks about' is ambiguous between a *relational* sense applying only to our thoughts about *existents* and to a non-relational sense applying only to our thoughts about Picard, Vulcan, the fountain of youth, and so on. To clarify, TR is meant to be in conflict with Searle's and Azzouni's views. For, when I say that, necessarily, a thinker's thinking about something *just is* her standing in the *thinking about* relation to it, I mean to rule out that there are two distinct senses of *thinks about*, one of which is relational and the other of which isn't. It will be clear in §9 that I mean for TR to rule out these views. For the moment, however, I set this complication aside.

Zalta (1983, 1988: 10) – accept both TR and Serious Actualism (or at least they would accept Serious Actualism, as *I* characterize it). As such, they are committed to the claim that we cannot think about nonexistents.

I, too, accept both TR and Serious Actualism. Above I presented my case for Serious Actualism. And in §10 I shall briefly sketch an account of how to make plausible the idea that we cannot think about nonexistents. The bulk of this chapter consists of an argument for TR.

My argument for TR centers on two claims. One of these is Content Fixes Truth (CFT), the claim that, necessarily, if two thoughts share representational content (of some sort), then they have the same truth value. The other is *Thinking-about-an-Entity as Relational* (TER):

> **TER** Necessarily, whenever a thinker thinks about an entity, her thinking about that entity *just is* her standing in the *thinking about* relation to it.

TER says what TR says, with that exception that TER is restricted to our thinking about entities.

My argument proceeds as follows. In §2 I present a toy argument, which I call "the Magician Argument." On the reading of the Magician Argument I specify, the Magician Argument is uncontroversially valid. Moreover, on that reading, that argument's validity is explained by the fact that the conjunction of its premises is an instance of its conclusion. In §3 through §8, I argue that, to make sense of the Magician Argument's being valid for this reason, we must accept TER. In §9 I present an argument – one of whose premises is TER – for Thinking as Relational (TR). I show that this argument succeeds provided that CFT is true. I conclude that Thinking as Relational (TR) follows from CFT. Consequently, since I accept CFT, I accept Thinking as Relational.

2 – The Magician Argument

Right now, it is September 19, 2016, 9:30 p.m. (or so let's assume, for the sake of my story). And, right now, my friend Doug is thinking about the actor Neil Patrick Harris. I know he is thinking about Harris because, as I try to work, Doug is reminding me that, twenty-seven years ago to the minute, Harris first graced our T.V. screens as the eponymous character of ABC's celebrated series, *Doogie Howser, M.D.* Unbeknownst to Doug, Doug – in thinking about Harris – is thinking about a magician. For, in addition to being a celebrated actor, Harris is an award-winning magician (though Doug does not know this).

With this in mind, consider the following argument:

(1) Doug thinks about Harris.

(2) Harris is a magician.

Therefore,

(3) Doug thinks about a magician.

Call this argument the Magician Argument. To clarify, (1) is to be interpreted as the claim that Doug is thinking about someone, which someone is identical to Harris. (1) is not to be interpreted as the claim that Doug is thinking about someone *as* Harris.[9] Likewise, (3) is to be interpreted as the claim that Doug is thinking about someone, which someone is a magician. (3) is *not* to be interpreted as the claim that Doug is thinking about someone *as* a magician. (Again, Doug doesn't even know that Harris is a magician.[10]) Call these interpretations of (1) and (3) "the extensional

[9] Here and throughout, to avoid awkwardness in my writing, I use 'someone' when, strictly speaking, I should use 'something'. Accordingly, for the sake of my argument, 'someone' should be treated as simply meaning *something*, and not *something that is a person*.

[10] Compare: Miranda Hillard is unaware that her ex-husband Daniel is moonlighting as her children's nanny, but in thinking about Daniel she is thinking about Mrs. Doubtfire.

interpretations of (1) and (3).” Accordingly, call the interpretation of the Magician Argument that includes these interpretations “the extensional interpretation of the Magician Argument.”

Given this clarification, it should be easy to see that the Magician Argument is valid (and, for that matter, sound). Consider the following story:

> **NO-THINK.** Doug is not thinking about a magician. (That is, it’s not the case that Doug is thinking about someone, which someone is a magician.) But Doug is thinking about Neil Patrick Harris, and Neil Patrick Harris is a magician.

NO-THINK isn’t just false. It’s necessarily false. Indeed, NO-THINK is *incoherent*. So it’s impossible – incoherent, even – for (1) and (2) to be true while, at the same time, (3) is not true.

So the Magician Argument is valid. That is, it is modally valid. Moreover, there is something in particular that *explains* its being modally valid: the conjunction of its premises is an *instance* of its conclusion. In asserting the conclusion, we assert nothing over and above what we assert in asserting the conjunction of the argument’s premises. It’s just that we assert what we do at a more abstract level, such that some of the details aren’t filled in. Put differently, in stating the conclusion, we point to the same state of affairs (i.e., the same event) as we do in stating the conjunction of the premises; we just do so at a higher level of generality.[11]

[11] Trenton Merricks (2015) argues that the sense in which a given argument can be characterized as *valid* depends very much on the ontology of the argument – on whether, for instance, the argument is composed of events, of sentences, of mental states, or of abstract propositions. I want to remain neutral with regards to the ontology of the Magician Argument. To do so, I’ve characterized validity-by-generalization in terms of what we assert in asserting the argument’s premises, rather than in terms of the (potentially) intrinsic properties of the argument itself. And in the following sections, when we speak of the way in which the Magician Argument can be analyzed, we can understand all of this ultimately in terms of how *what the Magician Argument is about* can be represented. So the claims I make about the Magician Argument are explicitly claims about the Magician Argument’s subject matter and how it is represented. Thus, the extent to which the claims I make are about intrinsic features of the Magician Argument may very well depend on the Magician Argument’s ontology, but the truth values of my claims (and the validity of my inferences) do not.

Let's say that an argument is *valid by generalization* just in case the conjunction of its premises is an instance of its conclusion. We've just seen that the Magician Argument, at least on the extensional interpretation, is valid by generalization. And, to see this, we haven't had to presuppose anything about how the Magician Argument can (and can't) be formally analyzed. That means that we can appeal to this fact in defending claims about the Magician Argument's formal analysis. Let's do so presently.

3 – Formally Analyzing the Magician Argument

To formally analyze a state of affairs S is to use a certain *interpreted* formal expression $\mathscr{B}$ to (successfully) represent S. Likewise, to formally analyze a proposition P is just to use an interpreted formal expression $\mathscr{B}$ to represent whatever state of affairs P represents, and to formally analyze an argument is just to jointly analyze its premises and conclusion by appeal to an ordered set of formal expressions (which order corresponds to that of the argument itself).

My use of 'formal analysis' is derivative on my use of 'formally analyze'. I use 'formal analysis' to represent the *activity* wherein one formally analyzes a given claim or argument, and I also use it to represent the *formal expression* that results from that analysis (i.e., the analysans).

(Strictly speaking, then, these terms are success terms. To avoid verbosity, I'll often speak loosely in that I'll describe *candidate* formal analyses simply as formal analyses, even though it's still an open question whether they actually are formal analyses. Likewise, I'll often speak of our attempt to analyze a certain claim or premise in a certain way just as our analyzing it in that way.)

It's an open question to what extent we can read ontology off of formal analysis. For the next several sections, I'm just going to assume that we can. I'll assume that, for any claim, (a) there is one and only one way that it can be formally analyzed and (b) the ontological structure of that claim can be read directly off of that analysis. Say, then, that the Magician Argument's premise

(1) is analyzable as <T(d, h)>, where T(x, y) $=_{df}$ x thinks about y and where 'h' and 'd', as names, *denote* and thereby *pick out* Harris and Doug, respectively.[12] Then, since (1) is true, the state of affairs of Doug's thinking about Harris obtains, and furthermore that state of affairs *just is* Doug's standing in the *thinking about* relation to Harris. Say, on the other hand, that the one proper analysis of (1) is as <T^h(d)> such that 'd' denotes Doug and T^h(x) $=_{df}$ x thinks Harris-wise. While 'd' picks out Harris and while 'T^h' *mentions* Harris in that it is essentially *about* Harris, it doesn't *pick out* Harris. So if (1) is analyzable in this way, then Doug's thinking about Harris does not involve Harris's instantiating a property at all. Rather, Doug's thinking about Harris *just is* Doug's instantiating the monadic property *thinking about Harris*.

Over the next few sections, I'll show that, given this assumption, Doug's thinking about Harris just is his standing in the *thinking about* relation to Harris. I'll begin, in §4 through §6, by looking at three different candidate analyses of the Magician Argument. One of these is the most straightforward candidate analysis. The other two are alternative analyses suggested by the work of Uriah Kriegel (2007). Kriegel (2007) proposes a couple of ways of understanding what it is to think about something that are meant to be alternatives to the conception of thought as the instantiation of a relation between a thinker and what is being thought about. These alternatives,

[12] I use 'mentions' and 'picks out' as terms of art. A term *mentions* me just in case it is essentially about me, at least given its interpretation. (It thus *rigidly designates* me.) For instance, I am mentioned by the predicate 'N' such that N(x) = x is Nick Rimell, and I am also mentioned by the predicate 'D' such that D(x) = x is the oldest son of Donna in α (i.e., the actual world). While these terms mention me, neither pick me out. But another way to mention me *is* to pick me out. To pick me out is to mention me by essentially selecting me from the domain of the existential quantifier (given its interpretation). There's two ways in which this may happen. First, the term may denote me, as a name does. Second, the term may be a variable – 'x', say – such that I am its sole value in all possible worlds in which it has a value. To see what I mean, consider the formulas <$\exists$x S(x, d)>, <$\exists$x (S(x, d) ^ $\exists$y(S(y, d) ^ Y(y,x))>, and <$\exists$x (S_α(x, d) ^ $\exists$y(S_α(y, d) ^ Y_α(y,x))> such that, in all of these formulas, 'd' denotes Donna, S(x, y) $=_{df}$ x is a son of y, Y(x, y) $=_{df}$ x is younger than y, S_α(x, y) $=_{df}$ x is a son of y in α, and Y_α(x, y) $=_{df}$ x is younger than y in α. I am Donna's son, but I am not her only son. I have precisely one brother, who is younger than me. Of course, I *could* have had an older brother also. Thus, in the former of these two formulas, there is no variable that *picks me out*. But, in the latter of these formulas, there is: 'x'.

then, should really be thought of as alternatives to Thinking as Relational (TR). Nevertheless, for the sake of argument, I shall treat them as alternatives to TER.[13]

In what follows, we'll see that, for any one of these three proposed analyses, either that proposed analysis fails to accommodate the validity-by-generalization of the Magician Argument or it commits us to the view that Doug's thinking about Harris just is his standing in the *thinking about* relation to Harris. More importantly, we'll develop a picture of what it is for a candidate analysis to accommodate an argument's validity-by-generalization. We'll then apply this picture, in §7, towards showing that Doug's thinking about Harris just is his standing in the *thinking about* relation to Harris, at least given our assumption that we can read ontology off of formal analysis, In §8, we'll discharge this assumption.

4 -- The Straightforward Analysis

The second premise of the Magician Argument is that Neil Patrick Harris is a magician. This is essentially about Neil Patrick Harris. It is also about his being a magician. At face value, then, it is about the event of Harris's instantiating the property of *magicianhood*.

But maybe you don't think there is any such property. For instance, maybe you don't believe in abundant properties. Fair enough. Choose a property, any property, so long as Harris exemplifies it.[14] Whatever it is, feel free to substitute talk of that property for my talk of *magicianhood*, and take (2) to be about that. With this in mind, we can all agree that premise (2) is analyzable as (and therefore *must* be analyzed as) <M(h)>, such that 'h' denotes Harris and M(x)

[13] In §9 we'll see that, given CFT, the reasons why these accounts fail as alternatives to TER equally count as reasons why they fail as alternatives to TR.

[14] Assuming that Harris instantiates at least one monadic property, select among those. But, if for some reason, you think that all of Harris's properties are polyadic, choose one of those. What I say over §4 through §7 about how to analyze the Magician Argument can be easily tweaked to accommodate this.

$=_{df}$ x is a magician, at least *given* our assumption that we can read ontology off of formal analysis. (But, again, feel free to substitute some other predicate for 'is a magician'.)[15]

With this in mind, we should only take seriously those analyses of the Magician Argument that include <M(h)> as the analysis of (2). Here is one such analysis:

[1A] T(d, h)

[2A] M(h)

Therefore,

[3A] ∃x (T(d, x) ^ M(x))

Here, and henceforth, 'd' denotes Doug, 'h' denotes Harris, T(x, y) $=_{df}$ x thinks about y, and M(x) $=_{df}$ x is a magician. This analysis of the Magician Argument is the most straightforward candidate analysis. So let's call it the Straightforward Analysis.

If the Straightforward Analysis really is a proper analysis of the Magician Argument, then Doug's thinking about Harris just is his standing in the *thinking about* relation to Harris. For, on this analysis, premise (1) is analyzed as <T(d, h)>.[16] This includes the two terms 'd' and 'h', respectively denoting Doug and Harris, and both of these terms are in the argument of a predicate expression, which expression represents Doug's thinking about Harris.[17]

[15] A qualification: The term in the argument needn't be 'h'. It just needs to be a term that denotes Harris and thereby picks him out of the domain of the existential quantifier. Throughout this paper, I'm just going to use 'h' as a stand in for any such Harris-denoting term, since the ontological upshot is the same.

[16] And there are near-identical candidate analyses of the Magician Argument that also commit us to this metaphysical picture. The most obvious one is the candidate analysis that differs in that (1) is analyzed in terms of the triadic predicate expression $<T^3(d, h, a)>$ such that $T^3(x, y, t) =_{df}$ x thinks about y at t, while 'a' denotes the relevant time (in particular, June 15, 2015, 9:30 a.m.) And (3) is altered accordingly.

[17] A free logician might challenge my claim that 'd' and 'h' *denote* Harris. But note that, on this analysis, the conjunction of (1) and (2) follows by existential generalization. This requires that we read <T(d, h)> as being true just in case the open formula <T(x, y)> is satisfied by (Doug, Harris, such that Doug and Harris are both members of the domain of the existential quantifier. Thus, 'd' and 'h' really do *denote* Doug and Harris in that they are names that select them from among the members of the domain of the existential

In appealing to the Straightforward Analysis, we succeed in formally demonstrating that the Magician Argument is valid by generalization. To see this, consider any pair of formulas $\mathcal{B}$ and $\mathcal{C}$ such that $\mathcal{C}$ is a legitimate formalization of (3) and $\mathcal{B}$ is a legitimate formalization of (1 ^ 2) – i.e., of the conjunction of the Magician Argument's premises.[18] In offering $\mathcal{B}$ and $\mathcal{C}$ as such formalizations, we succeed in demonstrating the Magician Argument's validity-by-generalization *if and only if* $\mathcal{C}$ is also an instance of $\mathcal{B}$. After all, $\mathcal{C}$ represents the same state of affairs that is represented in our asserting the Magician Argument's conclusion. It thus follows that we can show that (1 ^ 2) is an instance of (3) by showing that the formula $\mathcal{B}$ is, likewise, an instance of the formula $\mathcal{C}$. But if we fail to show that this relation holds between $\mathcal{B}$ and $\mathcal{C}$, then – since $\mathcal{B}$ and $\mathcal{C}$ are supposed to represent (1 ^ 2) and (3), respectively – in presenting the argument these claims as such we fail to show that the former is an instance of the latter.

The most obvious way to show that a given formula $\mathcal{B}$ is an instance of another formula $\mathcal{C}$ is to show that $\mathcal{C}$ is demonstrable (i.e., provable) from $\mathcal{B}$ by *existential generalization*. Here, by 'existential generalization', I mean something quite broad: one or more instances of a term (name, predicate, etc.) is replaced with some variable *X* whose value falls under the unrestricted domain of quantification, and '∃*X*' plus the necessary parentheses are added *anywhere* in the formula, just so long as the instances of *X* fall within the quantifier's scope and the formula remains well-formed. (Thus, in this sense of 'existential generalization', <∃x D(x) ^ C(c)> is provable by

quantifier. This will be consistent with every other (purported) analysis that I present, provided that it accommodates the Magician Argument's validity-by-generalization.

[18] By 'formula', I mean – in particular – a well formed formula in some formal system or other that is truth tracking, at least in the sense that it can be used to correctly sort modally valid arguments from modally invalid arguments.

existential generalization from <D(b) ^ C(c)>, and <∃**ϕ** **ϕ**(b)> is provable by existential generalization from D(b).)[19]

The conjunction of [1A] and [2A] is clearly an instance of [3A], for [3A] is straightforwardly demonstrable from this conjunction by existential generalization, as so defined. The conjunction of [1A] and [2A] is <T(d, h) ^ M(h)>, and [3A], again, is <∃x (T(d, x) ^ M(x))>. And <∃x (T(d, x) ^ M(x))> is provable from <T(d, h) ^ M(h)> by existential generalization. This demonstrates that (1 ^ 2) is an instance of (3). So the Straightforward Analysis accounts for the Magician Argument's validity-by-generalization.

5 -- Kriegel's Monadic Adverbialism

Uriah Kriegel (2007) thinks that "representing something does not involve bearing a relation to it" (2007: 309). Kriegel proposes two alternative accounts of representation, the simplest of which is exemplified by the following:

> Although the surface grammar of "you are thinking of Bigfoot" casts the thinking as a relation between you and Bigfoot, the sentence can be paraphrased into "you are thinking Bigfootly," or perhaps more naturally, "you are thinking Bigfoot-wise." The latter casts the thinking as a non-relational property of yours. (314)

Thus, for Kriegel, a representation of Bigfoot is constituted by "the instantiation of a non-relational property of representing Bigfoot-wise" (315).

For the sake of argument, let us extend Kriegel's analysis of what it is to think about Bigfoot to an analysis of what it is to think about some entity. In doing so, we arrive at the

[19] Here I am relying on there being some intuitive sense of a formula's being provable from other formulas, without its simply being provable from those formulas within a particular proof system. Without giving a definition, we might characterize this roughly as follows: for any formula $\mathcal{B}$ and set Δ of formulas (possibly the null set), $\mathcal{B}$ is provable from the members of Δ just in case $\mathcal{B}$ is provable from the members of Δ in some proof system that is sound for a logic that is *truth preserving*, where a logic is truth preserving just in case it can be used, without fail, to identify modally valid arguments.

following picture: necessarily, whenever a thinker thinks about some entity E, her doing so just is her instantiating a monadic, adverbialist property of the form *thinking E-wise*. Call this view *Thinking as Monadic*. Given our assumption that we can read ontology off of formal analysis, *Thinking as Monadic* entails that the Magician Argument's premise (1) can be analyzed as follows:

[1B] $T^h(d)$

Here, $T^h(x) =_{df}$ x thinks Harris-wise, where it is further understood that to think about Harris just is to think Harris-wise.

Any analysis of the Magician Argument that includes [1B] as the analysis of (1) is, as such, both *monadic* and *adverbialist*. It is *monadic* in that it involves our formalizing Doug's thinking about Harris in terms of a monadic, rather than a dyadic, predication. And it is *adverbialist* in that the term that *mentions* Harris but does not *pick out* Harris.

Here is one such candidate analysis:

[1B] $T^h(d)$

[2B] M(h)

Therefore,

[3B] $T^M(d)$

Here, $T^M(x)$ = x thinks about a magician.[20]

[20] Remember that, in the interpretation relevant to the Magician Argument, to say of some x that it thinks about a magician is to say that x thinks about someone, which someone is a magician.

Again, Kriegel does not explicitly offer the above candidate analysis. But the above seems to me to be the most *natural* candidate analysis of the Magician Argument among those that include [1B] as an analysis of (1).[21] So I'll call it the Kriegel Monadic Analysis.

The Kriegel Analysis does not account for the Magician Argument's being valid by generalization. As we saw in §4, an analysis of the Magician Argument accounts for its being valid by generalization if and only if, on that analysis, there is a pair of formulas $\mathcal{B}$ and $\mathcal{C}$ such that (1 ^ 2) is analyzable as $\mathcal{B}$, (3) is analyzable as $\mathcal{C}$, and $\mathcal{B}$ is itself an instance of $\mathcal{C}$. But on the Kriegel Analysis, (3) is formalized as $<T^M(d)>$ while (1 ^ 2) is formalized as $<T^h(d) \wedge M(h)>$. And $<T^h(d) \wedge M(h)>$ is not an instance of $<T^M(d)>$.

To begin, $<T^M(d)>$ is not demonstrable from $<T^h(d) \wedge M(h)>$ by existential generalization, or even given multiple iterations of existential generalization. If it were, then $<T^M(d)>$ would include at least one instance of '∃', which of course it doesn't.

And there's only one other way – so far as I can think – whereby any given formula may plausibly be said to be an instance of another: It's what I shall call determinate-to-determinable generalization. Let's say that, for a pair of formulas $\mathcal{B}$ and $\mathcal{C}$, $\mathcal{C}$ generalizes from $\mathcal{B}$ by determinate-to-determinable generalization just in case $\mathcal{C}$ results from replacing one or more instances of single term in $\mathcal{B}$ with another term, such that the former term represents a determinate of the latter. (Thus, determinate-to-determinable generalization requires the terms in question to be interpreted.)

[21] But there is a related candidate, inspired by something else Kriegel says. Kriegel writes that, "if we were to read ontology off of grammar, [the sentence 'x represents Bigfoot-wise'] would suggest a state of affairs involving a particular x instantiating the non-relational property of representing, which in turn instantiates a non-relational property of occurring Bigfoot-wise" (2007: fn. 19 (314/331)). Accordingly, we might formalize Doug's thinking about Harris as $<\eta(T_1(d))>$ and Doug's thinking about a magician as $<\mu(T_1(d))>$, such that 'd' denotes Doug, $T_1(x) =_{df}$ x thinks, $\eta(\Phi(x)) =_{df}$ x Φs Harris-wise, and $\mu(\Phi(x)) =_{df}$ x Φs magician-wise. It will be clear from what follows that the according analysis of the Magician Argument fails to accommodate the Magician Argument's validity-by-generalization for precisely the same reason that the Kriegel Monadic Analysis fails. Furthermore, this sort of analysis is ruled out by the general argument presented in §7.

Consider the pair of formulas <D(b)> and <A(b)> where A(x) $=_{df}$ x is an animal and where everything else is interpreted as above. Necessarily, all dogs are mammals. In this very loose sense, doghood is a determinate of mammalhood. So <A(b)> results from replacing one or more instances of a single term in <D(b)> with another term (i.e., replacing 'D' with 'A'), such that the former term represents a determinate of the latter.

But $<T^M(d)>$ is not demonstrable from $<T^h(d) \wedge M(h)>$ by determinate-to-determinable generalization either, nor even by multiple iterations of it. For $<T^M(d)>$ to generalize in this way, it has to include some term that replaces a term in $<T^h(d) \wedge M(h)>$ such that the former picks out a determinable of that which the latter picks out. The only term in $<T^M(d)>$ that is a plausible candidate for such a replacement is 'T^M', and the only term in $<T^h(d) \wedge M(h)>$ that it plausibly replaces is 'T^h'. But such a replacement would violate the rule for determinate-to-determinable translation. For Harris is only *contingently* a magician, so *thinking about Harris* is not a determinate of *thinking about a magician.*[22]

Thus, we cannot appeal to the Kriegel Analysis to formally demonstrate that the Magician Argument is valid by generalization. There is no number of iterations of existential generalization by which $<T^M(d)>$ is demonstrable from $<T^h(d) \wedge M(h)>$, nor is there any such number of iterations of determinate-to-determinable generalization. Nor is there any *combination* of such iterations.

[22] Determinate-to-determinable generalization is inspired by a solution Kriegel (2007: 316) offers in response to a related worry. The worry Kriegel has in mind is based on Frank Jackson's (1977) objection to adverbialism with respect to experience. This worry has to do with the adverbialist's ability to accommodate the validity of the inference from 'Mental state M represents a purple rectangle' to 'Mental state M represents a rectangle'. *Prima facie*, the adverbialist cannot accommodate this inference, since the adverbialist must take *representing-purple-rectangle-wise* and *representing-rectangle-wise* to be different monadic properties. Kriegel's solution is to say that *representing-rectangle-wise* is a determinable of which *representing-purple-rectangle-wise* is a determinate (2007: 323-324). This is a plausible solution to the Jackson-inspired worry to which Kriegel responds. But, as we've just seen, it doesn't help Kriegel evade the worry I raise.

For, if there were, then $<T^M(d)>$ would still involve at least one instance of '∃'. And it doesn't. So $<T^h(d) \wedge M(h)>$ is not an instance of $<T^M(d)>$.

Of course, the Kriegel Analysis is not the *only* plausible analysis of the Magician Argument that is monadic or adverbialist or both. In the general argument I present in §7, these candidate analyses will be ruled out as well.

6 -- Thinking by Proxy

Graham Priest says that "the most common suggestion" for how to avoid commitment to intentional relations to nonexistents is to "reparse [the relevant intentional relation] as one between the agent and some surrogate object, especially some mental representation" (2005: 58). Priest is speaking here of intentional attitudes involving *non*existents, so he isn't speaking of *entities*, in my sense of the term. But we might plausibly extend this suggestion (not endorsed by Priest) to thoughts and other intentional attitudes involving entities. Thus, take *Thinking by Proxy* to be the view that, necessarily, there is some relation R* such that, whenever a thinker thinks about some entity, her doing so just is her standing in R* not to the entity itself, but rather to a certain proxy for that entity. Thus, Doug's thinking about Harris does not involve his standing in a relation R* to Harris himself. Rather, it involves his standing in a relation R* to some proxy x, such that in standing in R* to x he thereby thinks about Harris.

Versions of Representation-by-Proxy vary depending on what sort of thing the proxy is taken to be. One suggestion is that the proxies are themselves intentional entities. For instance, when Doug is thinking about Harris, Harris's proxy may be the proposition <that Harris was the

star of *Doogie Howser, M.D.*>, or it may be the name 'Harris' or that name's Fregean sense. Or perhaps it is the concept Harris.[23] Or maybe it is something else.

But the proxies needn't be intentional. Perhaps, instead, they are states of affairs. Say, for instance, that Doug is thinking about Harris in virtue of his recalling that Harris has hosted the Tonys four times. Perhaps the proxy is the state of affairs of Harris's having hosted the Tonys four times. Or perhaps, instead, it is a property. More specifically, perhaps it is Harris's haecceity – *Harrisity*.[24]

In the previous section we were introduced to an account, provided by Uriah Kriegel, of how to make sense of our thinking about Bigfoot. Kriegel (2007) also offers another, more refined account. Here's how he characterizes the account, overall:

> The [account] has two parts, one for conscious representations and one for non-conscious ones. First, a conscious representation of Bigfoot does not represent Bigfoot in virtue of bearing a relation to Bigfoot, but in virtue of instantiating some non-relational property of representing Bigfoot-wise. Second, a non-conscious representation of Bigfoot does not represent Bigfoot in virtue of bearing a relation to Bigfoot, but in virtue of bearing a relation to conscious representations of Bigfoot. More generally, non-conscious representations derive their intentionality from conscious representations, not from relations to intentional objects, while conscious representations have an altogether non-relational intentionality. (2007: 318)

Thus, Kriegel's account has two parts. It involves a characterization of *non-derivative* intentionality, which Kriegel takes to be the intentionality of conscious mental states. It also involves a separate account of *derivative* intentionality, only the latter of which is relational.

---

[23] These particular suggestions don't seem very promising, for they just push back the problem the denier of TER is trying to address in the first place – namely, how can we account for *representation* without positing a relation between the representor and what it represents?

[24] This suggestion is motivated by (though not entailed by) Plantinga's (1979: 109 – 111) suggestion that we take singular propositions to have, as literal constituents, not whatever they are directly about but rather the haecceities of whatever they are directly about.

Kriegel's overall account inspires the following version of *Thinking by Proxy*: necessarily, whenever a thinker thinks about some entity, her doing so just is her standing in a certain relation $\boldsymbol{\phi}$ to a certain conscious mental state, which mental state non-derivatively represents that entity.

This account suggests the following formal analysis of Doug's thinking about Harris:

[1C] $\exists x\,(G(d, x) \wedge H^{h}(x))$,

such that G(x, y) = x G's y (where 'G' is an arbitrary constant for whatever the above-mentioned surrogacy relation is) and such that $H^{h}(x) =_{df}$ x represents Harris-wise.

Furthermore, analyzing (1) as [1C] recommends a specific analysis of the Magician Argument, and it puts pressure on us to accept that analysis. For, if we analyze (1) as [1C], consistency demands that we analyze (3) as $\langle \exists x\,(G(d, x) \wedge H^{M}(x))\rangle$, such that $H^{M}(x) =_{df}$ x represents a magician. And, again, we must analyze (2) as $\langle M(h)\rangle$. We thus have the following analysis of the Magician Argument:

[1C] $\exists x\,(G(d, x) \wedge H^{h}(x))$

[2C] M(h)

[3C] $\exists x\,(G(d, x) \wedge H^{M}(x))$

Call this analysis the Kriegel Proxy Analysis.

The Kriegel Proxy Analysis does not account for the Magician Argument's being valid by generalization. If it did, then the conjunction of [1C] and [2C] would be an instance of [3C]. In other words, $\langle \exists x\,(G(d, x) \wedge H^{h}(x)) \wedge M(h)\rangle$ would be an instance of $\langle \exists x\,(G(d, x) \wedge H^{M}(x))\rangle$. And we've determined over §4 and §5 that a formula $\boldsymbol{\mathcal{B}}$ is an instance of a formula $\boldsymbol{\mathcal{C}}$ if and only if $\boldsymbol{\mathcal{C}}$ is provable from $\boldsymbol{\mathcal{B}}$ by one or more iterations either of existential generalization or of determinate-

to-determinable generalization (or of some combination thereof). But <∃x (G(d, x) ^ H^M(x))> is not provable from <∃x (G(d, x) ^ H^h(x)) ^ M(h)> in any of these ways.

It is certainly not provable by one or more iterations of existential generalization, or even by a combination of iterations *some of which* are of existential generalization. Of course, [3C] does involve an instance of '∃'. But the one instance of '∃' in [3C] closes the same instances of 'x' that are closed by the '∃' in <∃x (G(d, x) ^ H^h(x)) ^ M(h)>. So the single instance of '∃' in [3C] does not result from existential generalization on the conjunction of [1C] and [2C]. So [3C] is not provable from the conjunction of [1C] and [2C] by any iterations of existential generalization.

Nor is it provable by any instances of determinate-to-determinable generalization. The only plausible candidate for such an iteration is the replacement of 'H^h' with 'H^M'. But, again, determinate-to-determinable generalization requires that the term being replaced represent a determinate of that which is represented by the replacing term, and the determinate-to-determinable relation (as I have defined it) is an *essential* one. But, again, Harris is only contingently a magician.

So the conjunction of [1A] and [1B] is not an instance of [3C]. So the Kriegel Proxy Analysis fails to accommodate the Magician Argument's validity-by-generalization.

But perhaps there are other plausible analyses of (1) that that Kriegel version of *Thinking by Proxy* recommends. [1C] contains the predicate 'H^h' as its Harris-mentioning term. Might there be an analysis of (1) that differs in that it includes a *name* for Harris instead? Or might there be one that simply doesn't include a Harris-mentioning term at all?[25]

[25] These options are not exhaustive. Another option is formalize (1) in terms of an expression that involves non-predicate, non-name Harris-mentioning terms. For example, we might analyze (1) as ∃x (G(d, x) ^ η(H(x))), where H(x) =$_{df}$ x represents and η(*Φ*(x)) =$_{df}$ x *Φ*s Harris-wise. Or we might analyze (1) simply as <∃x R(d, x)>, such that R(x, y) =$_{df}$ x Rs y, where 'R' is an arbitrary constant representing a Harris-mentioning surrogacy relation, e.g., *having the Harris-representing mental state*. This latter suggestion corresponds to what Priest (2005: 58-59) himself considers on behalf of those who attempt to characterize

Let's begin with the latter option. The problem here is that premise (1) is about Harris. So premise (1) is analyzable in terms of a certain formula $\mathcal{B}$ only if Harris is among that which $\mathcal{B}$ represents. But, again, a formula $\mathcal{B}$ represents Harris only if $\mathcal{B}$ includes a Harris-mentioning term.[26]

The former option looks more promising. For instance, we might attempt to analyze (1) as

$\exists x\ (G(d, x) \wedge R(x, h))$,

such that $G(x, y) =_{df}$ x instantiates the mental state y and $R(x, y) =_{df}$ x is about y. And if we analyze (1) in this way, then we should analyze (3) accordingly: namely, as $\langle \exists y\ (\exists x\ (G(d, x) \wedge R(x, y)) \wedge M(y)) \rangle$. And, of course, since we must analyze (2) as $\langle M(h) \rangle$, then in analyzing (1) as above we must thereby analyze $(1 \wedge 2)$ as $\langle \exists x\ (G(d, x) \wedge R(x, h)) \wedge M(h) \rangle$. And $\langle \exists y\ (\exists x\ (G(d, x) \wedge R(x, y)) \wedge M(y)) \rangle$ generalizes from this formula.

But this candidate analysis is not at all faithful to Kriegel's project. In fact, if (1) really is analyzable as $\langle \exists x\ (G(d, x) \wedge R(x, h)) \rangle$, then Doug's thinking about Harris just is Doug's standing in the *thinking about* relation to Harris. It's just that, in standing in the *thinking about* relation to

---

intentional states about nonexistents in terms of relations to proxies. Priest's arguments against this sort of account begin as my own argument has begun: the problem comes when we attempt to apply this account of analysis towards claims regarding intentional objects' having certain features (e.g., *being a magician*). What Priest doesn't do is present a general account of why such an analysis will fail as a proper analysis of anything. This is what I am attempting to do. I've begun the process above, and I'll complete it in §7. In so doing, I'll rule out these other alternatives as well.

[26] **Objection.** In saying that the proper analysis of (1) must include a Harris-mentioning term, I am assuming that premise (1) is *directly* about Harris rather than *indirectly* about Harris. Of course, if premise (1) is just the *sentence* 'Doug thinks about Harris', then it *is* directly about Harris. But, otherwise, it's not obvious it is (provided this distinction applies to the premise in the first place).

Response. I am not assuming this at all. I am simply saying that (1) is *essentially* about Harris. And it is. For instance, it is not just the proposition *that Doug thinks about the star of Doogie-Howser, M.D.* This proposition is not the same as the proposition *that Doug thinks about Harris*. If it were, it would have the same truth conditions. But it doesn't. For there is a (very distant) possible world in which Doug thinks about Fred Savage and where Savage, not Harris, is the star of *Doogie Howser, M.D.*

Harris, Doug thereby stands in an *internal* relation to Harris. Two entities stand in an *internal* relation just in case (a) they stand in a relation and (b) their doing so can be analyzed in terms of each individual's instantiating a separate property (which may or may not be monadic). Meanwhile, two individuals stand in an *external* relation just in case (a) they stand in a relation and (¬b) it is *not* the case that their doing so can be analyzed in terms of each individual's instantiating a separate property. Say, then, that (1) is analyzable as <∃x (G(d, x) ^ R(x, h))>. Then Doug's thinking about Harris just is Doug's standing in the internal *thinking about* relation to Harris. And this in turn just is Doug's standing in the external *instantiating* relation to some mental state and that mental state's standing in the external *being about* relation to Harris.[27]

7 – Generalizing about Generalizing

In §4 through §6, we saw that the Straightforward Analysis of the Magician Argument accommodates the Magician Argument's validity-by-generalization, whereas the candidate analyses based on Kriegel's accounts of thinking do not. More importantly, we learned a great deal along the way about what a proper analysis of the Magician Argument will (and *must*) look like.

For instance, we saw that a proper analysis of the Magician Argument must include Harris-mentioning terms in its analyses of all of those claims that are essentially about Harris. And, of course, if Harris is *not* essentially among the content of the claim, then the formula by which it is analyzed had better *not* include a Harris-mentioning term. Thus, the proper analysis of (1) must include a Harris-mentioning term, just as the proper analysis of (2) must. Meanwhile, the proper analysis of (3) must *not* include a Harris-mentioning term.

[27] Ultimately, what what's controversial is simply whether one's thinking about Harris involves *Harris's* instantiating a property. For it's this that threatens our purported ability to think about nonexistents.

We also developed a picture of how to formally demonstrate validity-by-generalization (or its lack thereof). In particular, we saw that an argument is valid by generalization just in case there's a pair of formulas $\mathcal{B}$ and $\mathcal{C}$ such that (a) the conjunction of the argument's premises is analyzable as $\mathcal{B}$, (b) the argument's conclusion is analyzable as $\mathcal{C}$, and (c) $\mathcal{C}$ is provable from $\mathcal{B}$ in one of the following ways:

(i) $\mathcal{C}$ is provable from $\mathcal{B}$ by one or more iterations of existential generalization,

(ii) $\mathcal{C}$ is provable from $\mathcal{B}$ by one or more iterations of determinate-to-determinable generalization, or

(iii) $\mathcal{C}$ is provable from $\mathcal{B}$ by some combination thereof.

To make things easier on us, let's use the word 'generalizes' to capture this relation between $\mathcal{B}$ and $\mathcal{C}$. That is, let's stipulate that a formula $\mathcal{C}$ generalizes from a formula $\mathcal{B}$ just in case $\mathcal{B}$ and $\mathcal{C}$ meet either of (i), (ii), or (iii).

Let's apply this constraint to the Magician Argument, which, again, is the following:

(1) Doug thinks about Harris.

(2) Harris is a magician.

Therefore,

(3) Doug thinks about a magician.

Remember also that a legitimate analysis of the Magician Argument must include, as an analysis of (2), the formula <M(h)> such that 'h' denotes Harris and $M(x) =_{df}$ x is a magician. So any analysis of the Magician Argument that accounts for its validity by generalization will include this as an analysis of (2). It follows that, on *any* analysis of the Magician Argument that accounts for its being valid by generalization, there is a pair of formulas $\mathcal{B}$ and $\mathcal{C}$ such that, on the analysis in

question, $\mathcal{B}$ is analyzable as (1), $\mathcal{C}$ is analyzable as (3), and $\mathcal{C}$ generalizes from the formula $\mathcal{B}$ ^ M(h), the latter being the analysis of (1 ^ 2). Take the pair of formulas $\underline{\mathbf{B}}_{(1)}$ and $\underline{\mathbf{C}}_{(3)}$ to be an *arbitrary* instance of one such pair. Thus, (1) is analyzable as $\underline{\mathbf{B}}_{(1)}$, (3) is analyzable as $\underline{\mathbf{C}}_{(3)}$, and $\underline{\mathbf{C}}_{(3)}$ generalizes from $\underline{\mathbf{B}}_{(1)}$^ M(h). We'll see that it follows that Doug's thinking about Harris just is his standing in the *thinking about* relation to Harris. (Again, we're just taking for granted, for the time being, that we can read ontology transparently off of formal analysis.)

My argument involves two steps. The first is to show that $\underline{\mathbf{C}}_{(3)}$'s generalizing from $\underline{\mathbf{B}}_{(1)}$^M(h) includes – and *only* includes – a single iteration of existential generalization wherein multiple instances of a single denoting term for Harris are replaced.[28] The second step is to infer that Doug's thinking about Harris just is his standing in the *thinking about* relation to Harris.

Here's the first step. Since $\underline{\mathbf{C}}_{(3)}$ contains *no* Harris-mentioning term, our generalizing from $\underline{\mathbf{B}}_{(1)}$^M(h) to $\underline{\mathbf{C}}_{(3)}$ must include an iteration of generalization wherein (at the very least) the 'h' of 'M(h)' is replaced. This leaves open the question of whether it is the *only* iteration of generalization involved in our getting from $\underline{\mathbf{B}}_{(1)}$^M(h) to $\underline{\mathbf{C}}_{(3)}$.

Assume, for *reductio*, that it isn't the only iteration involved in our generalizing from $\underline{\mathbf{B}}_{(1)}$^M(h) to $\underline{\mathbf{C}}_{(3)}$. Then there is some other iteration of generalization involved in our move from $\underline{\mathbf{B}}_{(1)}$^M(h) to $\underline{\mathbf{C}}_{(3)}$, and so it must *this* iteration of generalization wherein the Harris-mentioning term of $\underline{\mathbf{B}}_{(1)}$ is replaced. And here we have two options. Either this other iteration includes *after* the replacement of the 'h' in 'M(h)', or it occurs *before* it. It turns out that taking either of these options leads to a contradiction.

Assume, first, that in our generalizing from $\underline{\mathbf{B}}_{(1)}$^M(h) to $\underline{\mathbf{C}}_{(3)}$, we begin by replacing the 'h' of 'M(h)' and *then* replace the Harris-mentioning term of $\underline{\mathbf{B}}_{(1)}$, whatever that is. On this

[28] I speak of 'h', for simplicity's sake. But any term that denotes Harris will do.

assumption, $\underline{\mathbf{C}}_{(3)}$ generalizes from some formula $\mathcal{D}$ that results simply from our replacing the far right 'h' of $\underline{\mathbf{B}}_{(1)}$ ^ M(h). But if $\mathcal{D}$ results simply from our replacing the far right 'h' of $\underline{\mathbf{B}}_{(1)}$ ^ M(h), then $\mathcal{D}$ expresses a claim that includes Doug's thinking about Harris but does not assert that Harris *himself* is a magician. Rather, it just involves, in some sense, *magicianhood* or *someone's* being a magician. And it doesn't say of *whoever* is a magician that *that individual* is thought about by Doug. But, then, there's a possible world where $\mathcal{D}$ is true even though Doug does *not* think about a magician. (This is a world where Doug thinks about Harris and where there are magicians, but where Harris is not one of them.) So $\mathcal{D}$ does not generalize to $\underline{\mathbf{C}}_{(3)}$ after all. So we have a contradiction.

Assume, instead, that our replacing the Harris-mentioning term of $\underline{\mathbf{B}}_{(1)}$ precedes our replacing the 'h' of 'M(h)'. Then $\underline{\mathbf{C}}_{(3)}$ generalizes from a formula $\mathcal{E}$ that results from our replacing, the Harris-mentioning term in $\underline{\mathbf{B}}_{(1)}$ *and nothing else*. Here, too, we face a contradiction. For if $\mathcal{E}$ results simply from our replacing the Harris-mentioning term of $\underline{\mathbf{B}}_{(1)}$ and nothing else, then $\mathcal{E}$ expresses a claim that includes Harris's being a magician but does *not* present this magician as the object of Doug's thought. Rather, it just involves some more general claim about Doug's thinking about so and so. And this claim – unlike $\underline{\mathbf{C}}_{(3)}$ – is consistent with Doug's *not* thinking about a magician. So $\underline{\mathbf{C}}_{(3)}$ does not generalize to $\mathcal{E}$. Contradiction.

Let's review. Our generalizing from $\underline{\mathbf{B}}_{(1)}$^M(h) to $\underline{\mathbf{C}}_{(3)}$ requires, at the bare minimum, one iteration wherein we replace the 'h' of 'M(h)'. We assumed for *reductio* that this particular iteration involves nothing else. On this assumption, then, there must be some other iteration wherein the Harris-mentioning term in $\underline{\mathbf{B}}_{(1)}$ is replaced. This iteration will occur either before or after the iteration wherein we replace the far-right 'h' of $\underline{\mathbf{B}}_{(1)}$^M(h). But we discovered that, either way, we encounter a contradiction. So we must reject what we assumed for *reductio*. We must

insist, instead, that our generalizing from $\underline{\mathbf{B}}_{(1)}$^M(h) to $\underline{\mathbf{C}}_{(3)}$ involves only a single iteration of generalization, wherein one or more instances of 'h' is replaced.

That completes the first step of my demonstration. The second step begins with the observation that the Harris-mentioning term in $\underline{\mathbf{B}}_{(1)}$ is also 'h', such that $\underline{\mathbf{C}}_{(3)}$ generalizes from $\underline{\mathbf{B}}_{(1)}$^M(h) by a single iteration of existential generalization. Again, $\underline{\mathbf{C}}_{(3)}$ does not include any Harris-mentioning terms. But $\underline{\mathbf{B}}_{(1)}$ does. So our single iteration of generalization must involve our replacing the Harris-mentioning term in $\underline{\mathbf{B}}_{(1)}$ as well. So the Harris-mentioning term in $\underline{\mathbf{B}}_{(1)}$ must be 'h'. And since $\underline{\mathbf{B}}_{(1)}$ is well-formed, 'h' is a member of the argument of some well-formed predicate expression (with or without a free variable). So $\underline{\mathbf{B}}_{(1)}$ includes an atomic predicate expression wherein 'h' is a member of the argument, and the move from $\underline{\mathbf{B}}_{(1)}$^M(h) to $\underline{\mathbf{C}}_{(3)}$ involves our simultaneously replacing both the 'h' in $\underline{\mathbf{B}}_{(1)}$ and the 'h' in <M(h)>. It thus involves a single iteration of existential generalization. So $\underline{\mathbf{C}}_{(3)}$ is identical to a formula $\exists x(\mathcal{B}(x) \wedge M(x))$ such that $\mathcal{B}(x)$ results from replacing the 'h' in $\underline{\mathbf{B}}_{(1)}$ with x (where x is free in $\mathcal{B}_{(x)}(x)$).

This tells us a great deal about how (1) must be analyzed. We know that $\underline{\mathbf{B}}_{(1)}$ includes a predicate expression whose argument has 'h' as a member. And we know that 'h' is a denoting term, for we know that it picks out a member of the domain of the existential quantifier. (Otherwise, $\underline{\mathbf{C}}_{(3)}$ wouldn't follow from $\underline{\mathbf{B}}_{(1)}$^M(h) by existential generalization.)[29] And it seems uncontroversial to add that $\underline{\mathbf{B}}_{(1)}$ also includes an atomic predicate expression whose argument includes 'd' (or some other Doug-denoting term). After all, (1) is analyzable as $\underline{\mathbf{B}}_{(1)}$, and (1) is the claim that Doug thinks about Harris. So there has to be a Doug-mentioning term in $\underline{\mathbf{B}}_{(1)}$ as well, and in a way that allows for $\underline{\mathbf{B}}_{(1)}$ to represent Doug's thinking about Harris. The obvious candidate is the name 'd', and – since $\underline{\mathbf{B}}_{(1)}$ is well-formed – it must be that $\underline{\mathbf{B}}_{(1)}$ includes an atomic predicate

[29] Thus, even free logicians must say that 'h' picks out a member of the existential quantifier.

expression wherein 'd' is a member of the argument.[30] Thus, since $\underline{\mathbf{B}}_{(1)}$ is an arbitrary constant for whatever (1) is analyzable as, we know that – for any formula $\mathcal{B}$ – (1) is analyzable as $\mathcal{B}$ only if $\mathcal{B}$ includes genuine denoting terms both for Harris and for Doug, which denoting terms are members of the argument (or arguments) of one or more predicate expressions.

This leaves open the details of what exactly $\underline{\mathbf{B}}_{(1)}$ is. The most obvious candidate is that the denoting terms for Doug and Harris are members of a single argument. So, for instance, it may be that (1) is analyzable as <T(d, h)>, or it may be that (1) is analyzable as <T(d, h, a)>, such that 'a' denotes the relevant time (September 19, 2016, 9:30 p.m.). Or it may be something else. But, in all cases, in reading our ontology off of this formula, we characterize Doug's thinking about Harris as Doug's standing in the *thinking about* relation to Harris (and perhaps to something else – e.g., a time).[31] And this relation is an *external* one.

But perhaps, instead, the denoting terms for Doug and Harris are members of different arguments, each of which is part of the term $\mathcal{B}$ such that (1) is analyzable as $\mathcal{B}$. For instance, it

---

[30] This, of course, is not to *prove* that $\underline{\mathbf{B}}_{(1)}$ includes 'd'. But consider the following slightly different argument:

(1) Doug thinks about Harris.
(2) Harris is a magician.
(4) Doug is a vegan.

Therefore,

(5) A vegan thinks about a magician.

It should be clear that this argument (at least on an extensional reading) is also valid by generalization and that we could run a near identical demonstration of how we can and can't analyze this argument to accommodate its being valid by generalization. And it should be clear that, here, we could prove that the formula by which we analyze (1) must include 'd' as well as 'h'. I'll take it for granted that this is obvious enough. And so I'll take myself as being just as justified in what I say of 'd' as I am in what I say of 'h'.

[31] This assumes that the predicate expression is not a negation of or a disjunct of the formula $\mathcal{B}$. But we needn't *assume* this. For, if the predicate expression *were* a negation or disjunct of the formula $\mathcal{B}$, then $\mathcal{B}$ wouldn't express the state of affairs of Doug's thinking about Harris. Instead, it would express the state of affairs of Doug's *not* thinking about Harris or the state of affairs of *either* Doug's thinking about Harris or something else's occurring.

may be the formula ∃x (G(d, x) ^ R(x, h)), such that G(x, y) $=_{df}$ x instantiates the mental state y and R(x, y) $=_{df}$ x is about y. Even still, Doug's thinking about Harris just is his standing in the *thinking about* relation to Harris. It's just that this relation is internal.

So, for any formula $\mathcal{B}$ such that (1) is analyzable as $\mathcal{B}$, $\mathcal{B}$ is true only if Doug's thinking about Harris is identical to Doug's standing in the *thinking about* relation to Harris (assuming that we can read ontology off of formal analysis).

8 -- The Magician Argument *Arguments*

In §2, we saw that the Magician Argument is valid by generalization. In §3 through §7 we assumed that ontology can be read off of the (proper) formal analysis of a claim or argument. And we saw that, based on this assumption, the only candidate analyses of (1) that can accommodate the Magician Argument's validity-by-generalization are those analyses that characterize Doug's thinking about Harris as Doug's standing in the *thinking about* relation to Harris. It follows that, *if* we can read ontology off of formal analysis in the straightforward way characterized in §3, *then* Doug's thinking about Harris just is his standing in the *thinking about* relation to Harris.

But maybe we can't read ontology straight off of formal analysis, at least not unconditionally. It doesn't matter. We're still committed to the claim that Doug's thinking about Harris just is his standing in the *thinking about* relation to Harris.

For, even if formal analysis *needn't* be ontologically committal in the way characterized in §3, it certainly *can* be. We can always just *stipulate* this commitment. Thus, let us say that, by definition, to *ontologically analyze* a claim or argument is just to provide a formal analysis with the express intention of representing the claim's or argument's underlying ontology. More precisely, for any claim P and formula $\mathcal{B}$, P is ontologically analyzable as $\mathcal{B}$ just in case (a) P is formally analyzable as $\mathcal{B}$ and (b) P's underlying ontology can be straightforwardly read off of $\mathcal{B}$

just as we've been doing for the past several sections. So there's a species of formal analysis called ontological analysis. And this is true just by stipulation.

Now, here is a substantial claim. Every claim – at least every *true* claim – is ontologically analyzable as something. Every event has an ontological structure. It may involve, for instance, two entity's standing in a certain relation, or it may involve one entity's exemplifying a monadic property, or it may involve something else. But, for any event, there is some fact of the matter with regards to what its underlying ontology is. So there will be *some* formal expression corresponding to that ontology, even if we don't know what that is (and even if we never *could* know). So, for every event, there's some formula that accurately represents its ontology. Now consider the formula that does this and that, furthermore, makes no attempt at representing the event in any other way. By definition, this is the formula by which we can properly *ontologically analyze* that event, or any claim representing that event. Thus, since every true claim represents an event, every true claim has a proper ontological analysis.

It follows that Doug's thinking about Harris just is his standing in the *thinking about* relation to Harris. And we don't need to make any assumptions to see this. (1) is true. So it has a proper ontological analysis. And ontological analysis is a subspecies of formal analysis. So there is at least one formal analysis of (1) such that we can read ontology straightforwardly off of this analysis. And, for any such analysis, that analysis characterizes Doug's thinking about Harris as Doug's standing in the *thinking about* relation to Harris.

We thus have the following sound argument for an *instance* of TER:

(A) The Magician Argument is valid by generalization.

(B) If the Magician Argument is valid by generalization, then for any formula $\mathscr{B}$ such that Doug's thinking about Harris is analyzable as $\mathscr{B}$, then – if we can

straightforwardly read ontology off of $\mathcal{B}$ – Doug's thinking about Harris just is his standing in the *thinking about* relation to Harris.

Therefore,

(C) For any formula $\mathcal{B}$ such that Doug's thinking about Harris is analyzable as $\mathcal{B}$, then – if we can straightforwardly read ontology off of $\mathcal{B}$ – Doug's thinking about Harris just is his standing in the *thinking about* relation to Harris. (from A, B)

(D) There is at least one formula $\mathcal{B}$ such that Doug's thinking about Harris is formally analyzable as $\mathcal{B}$ and such that we can straightforwardly read ontology off of $\mathcal{B}$. (Call this analysis an *ontological* analysis.)

Therefore,

(E) Doug's thinking about Harris just is his standing in the *thinking about* relation to Harris. (from C, D)

Call this argument the Doug-Harris Magician Argument. I'd like to make one qualification, which applies to (B) through (E). All of this time, I've been assuming that, if there is such a relation as the *thinking about* relation, there is a single *thinking about* relation. This assumes that properties are abundant. But perhaps they are not. If so, then this at least opens up the possibility that, in our speaking of Doug's thinking about Harris on September 19, 2016, 9:30 p.m., it is misleading to say that Doug stands in <u>the</u> *thinking about* relation to Harris. Perhaps, instead, it is simply <u>a</u> *thinking about* relation. Let's name it the *thinking about*$^{(D\text{-}H)}$ relation. So, in the above presentation of the Doug-Harris Magician Argument, sparse-property adherents should read 'the *thinking about* relation' as 'the *thinking about*$^{(D\text{-}H)}$ relation'.[32]

The Doug-Harris Magician Argument is sound. It is straightforwardly valid.[33] In §2, we saw that (A) is true. In §3 through §7, we saw that (B) is true. And we've just seen that (D) is true.

[32] Note here that I am not suggesting that the *concept* of thinking about something is not univocal. (In §9, I consider something along these lines.) I am just suggesting that (or at least acknowledging the possibility that) thinking about something is multiply realizable.

[33] (C) follows from (A) and (B) by *modus ponens*. It follows from (C) that, if Doug's thinking about Harris is analyzable by at least one formula off of which we can straightforwardly read our ontology, then Doug's

And what I've said about Doug's thinking about Harris generalizes to any instance of a thinker's thinking about an entity, as well as to any possible instance of a thinker's thinking about an entity. Doug's thinking about Harris is a perfectly arbitrary case of a thinker's thinking about an entity. And all entities, as such, exemplify properties. (Again, whether or not the relevant property is *magicianhood* is entirely incidental.) Consider any other such instance. For that instance, there is a sound argument that perfectly corresponds to the Doug-Harris Magician Argument, with the details changed in the obvious ways. And consider any non-actual possible world in which thinkers think about entities. Corresponding arguments may be run in those worlds as well. Thus, for any *possible* instance of a thinker's thinking about an entity, there is a sound argument for the claim that that thinker's thinking about that entity just is her (his, their, etc.) standing in the (or a) *thinking about* relation to that entity. Thus:

TER Necessarily, whenever a thinker thinks about an entity, her doing so *just is* her standing in the *thinking about* relation to that entity.

9 – Why Content Fixes Truth (CFT) entails Thinking as Relational (TR)

Consider the following argument:

TER Necessarily, whenever a thinker thinks about an entity, her doing so *just is* her standing in the *thinking about* relation to that entity.

Link Necessarily, if – whenever a thinker thinks about an entity – her doing so just is her standing in the *thinking about* relation to that entity, then – whenever a thinker thinks about something – her doing so just is her standing in the *thinking about* relation to that thing.

Thererfore,

TR Necessarily, whenever a thinker thinks about something, her doing so *just is* her standing in the *thinking about* relation to that thing.

---

thinking about Harris just is Doug's standing in the *thinking about* relation to Harris. (D) tells us that there is at least one such formula.

Call this argument "the Argument for Thinking as Relational." The Argument for Thinking as Relational is valid. And, as we've seen above, TER is true. So the Argument for Thinking as Relational is sound provided that Link is true.

My defense of Link is as follows. First, assume, for *reductio*, (a) that – whenever a thinker thinks about an entity – her doing so just is her standing in the *thinking about* relation to that entity and (b) that there is at least one instance of a thinker's thinking about something that is not identical to that thinker's standing in the *thinking about* relation to that thing. Call the thinker involved in this latter case "S," and call the thing that S is thinking about "A." So S is thinking about A. Moreover, given our assumption, A is not an entity, and S's thinking about A is not identical to S's standing in the *thinking about* relation to A. But, surely, whatever it is for a thinker to think about a non-entity is, at the very least, constitutively involved in what it is for a thinker to think about an entity.[34] But, then, for a thinker to think about a certain entity is not *just* for it to stand in the *thinking about* relation to that entity. This contradicts what we assumed for *reductio*. So what we assumed for *reductio* is false. Moreover, our argument for this relies on nothing contingent. So, we may conclude, what we assumed for *reductio* is necessarily false. So Link is true. And, again, TER is true, and the Argument for TR is valid. So the Argument for TR is sound and its conclusion, TR, is true.

The success of the above defense of Link turns on the falsity of the following claim, which I'll call "Two Senses":

TS There are two senses of *thinking about*. One sense of *thinking about* is the sense employed in the extensional interpretation of the Magician Argument

[34] For instance, if a thinker's thinking about a certain non-entity just is the thinker's exemplifying a certain monadic property of the form *thinking-such-and-such-wise*, then a thinker's thinking about an entity also involves such a monadic property exemplification, even if it is not identical to it.

> and applies only toward one's thinking about some entity. Another, *conceptually prior* sense of *thinking about* applies toward one's thinking about something, whether or not that something is an entity.

To see that the above defense of Link turns on the falsity of Two Senses, note that my argument for T*E*R is successful if and only if I am employing the same sense of *thinking about* in arguing for TER (and in articulating TER) that I employ in my initial interpretation of the Magician Argument. So let us just stipulate this as far as my defense of TER – and my articulation of TER itself – goes. I do not wish to stipulate this when it comes to TR, however. For I want TR to capture whatever plausible sense of *thinking about* we might employ in characterizing – here's how we can put it – "Somebody's thinking about something." Given this, the success of my defense of Link turns on there being no such plausible sense of *thinking about*, unless thinking about something, in that particular sense, at the very least constitutively involves thinking about something in the sense that I employ in the extensional interpretation of the Magician Argument I employ in §2.[35] So my defense of Link turns on there not being a plausible sense of *thinking about* on which the sense of *thinking about* I employ in my analysis of the Magician Argument is derivative. That is, the success of the above defense of Link depends on the falsity of Two Senses.

Content Fixes Truth (CFT) entails that Two Senses is false. Again, CFT is the following:

> **CFT** Necessarily, if two thoughts share representational content (of some sort), then they have the same truth value.

To see that CFT entails that Two Senses is false, assume, for conditional proof, that Two Senses is true. Call the sense of *thinking about* I employ in the extensional interpretation of the Magician Argument "the extensional sense of *thinking about*." What I have shown in §2 through §8, then, is

[35] In this case, we might say, this posterior sense of *thinking about* does not capture what it is, *fundamentally*, to think about something, which is really what TR is meant to capture.

that – in the extensional sense of *thinking about* – a particular thinker's thinking about a particular entity just is that thinker's standing in the *thinking about* relation to that entity. Now, if there is a conceptually prior sense of *thinking about*, then there is an event E such that – in this conceptually prior sense of *thinking about* – (a) a particular thinker S's thinking about something *just is* this particular event E and (b), if (and only if) that something exists, then the occurrence of E partly constitutes S's standing in the *thinking about* relation to that something. So, in this case, thinking about something *in any sense* always involves the occurrence of some event E, and only in certain cases does the occurrence of this event E partly constitute a thinker's standing in the *thinking about* relation to something. The only way to make sense of this, it seems to me, is to say that, in the extensional sense of *thinking about*, the facts about whether someone is thinking about something depend both on features that are intrinsic to the thinker and on features that are not intrinsic to the thinker, whereas – on the supposedly conceptually prior sense of *thinking about* – the facts about whether someone is thinking about something depend only on features of the former sort. But, then, on the supposedly conceptually prior sense of *thinking about*, the facts about what one is thinking about are *narrow*. Now, this does not by itself entail internalism about representational content. But it does entail that CFT is false. For, if the facts about what one is thinking about are narrow, then what one is referring to, in thought, is not fixed by what one is thinking about. And, if what one is referring to, in thought, is not fixed by what one is thinking about, then the truth value of one's thought is not fixed by what one is thinking about.[36] Again, CFT says that, necessarily, if two thoughts share representational content *of some sort*, then they have the same

[36] Say that Oscar thinks to himself – here's how he'd put it – "The Canadian PM is made up of about 50% H_2O." Meanwhile, say that Twin Oscar thinks to himself – here's how he'd put it – "The Canadian PM is made up of about 50% H_2O." The referent of Oscar's belief is the prime minister of Canada, whereas the referent of Twin Oscar's belief is the prime minister of Twin Canada. Canada's prime minister is made up of about 50% H_2O. Twin Canada's prime minister is not. So Oscar's belief is true whereas Twin Oscar's belief false. Of course, this is just one example. But cases like this are ubiquitous.

truth value. So, if there is a sense of *thinking about* according to which the facts regarding what one is thinking about are narrow, then CFT is false. So, to discharge our assumption for conditional proof: if Two Senses is true, then CFT is false. So Two Senses is false provided that CFT is true. Since my argument is based on nothing contingent, we may conclude that CFT entails that Two Senses is false.

So, to conclude: the Argument for Thinking as Relational is sound provided that Link is true; my defense of Link turns on the falsity of Two Senses; and CFT entails that Two Senses is false. So it follows from CFT that Link is true and, consequently, that the Argument for Thinking as Relational is sound. So it follows from CFT that Thinking as Relational (TR) is true.

10 – What We're Thinking about When We Think We're Thinking about Nonexistents

It follows from CFT that, necessarily, whenever a thinker thinks about something, her doing so just is her standing in the *thinking about* relation (or *a thinking about* relation) to that thing (§2 - §9). CFT is true (Ch. 2). So, necessarily, whenever a thinker thinks about something, her doing so just is her standing in the *thinking about* relation (or *a thinking about* relation) to that thing. And if follows from this claim – that is, from Thinking as Relational (TR) – that, if Serious Actualism is true, we cannot think about nonexistents (§1). Serious Actualism is true (§1). So we cannot think about nonexistents.

In accepting this claim I open myself up to a challenge: explain what we're doing when we think we're thinking about nonexistents. For instance, it may seem obvious to you that, right now, you are thinking about Santa Claus and, thereby, are thinking about something that doesn't exist. I deny that you are doing this. So I owe you an explanation of what it is that you are actually doing.

My preferred explanation is to say that you are thinking about properties.[37] In particular, you are thinking about certain properties that we associate with Santa Claus – e.g., the property of *being a right jolly old elf* and the property of *having a belly that shakes like a bowlful of jelly*. Now say that you think to yourself – here's how you might put it – "Santa Claus doesn't exist." In thinking this, you are thinking about certain properties, and you are identifying those properties as being non-co-exemplified. Or say that you think to yourself – here's how you might put it – "I am thinking about the fountain of youth, which doesn't exist." Here you are thinking about certain properties (*being a fountain*, etc.), you are characterizing these properties as being non-co-exemplified, and you are also characterizing yourself as standing in the *thinking about* relation to these properties.[38]

The account I have in mind does not just apply to those episodes in which we appear (misleadingly) to be thinking about *objects* that do not exist. The account also applies to episodes in which we appear (misleadingly) to be thinking about other sorts of things, e.g., events. Say, for instance, that I am eagerly awaiting the arrival of a fruit-of-the-month basket that, as it turns out, never arrives. It may seem that, while I eagerly await the arrival of the fruit-of-the-month basket, I am thereby thinking about a nonexistent event – namely, the arrival of the fruit-of-the-month basket. But, again, what I am really thinking about are certain properties – the property of *being a fruit*, the property of *being a basket*, etc. I am also thinking about these properties as being related to one another in a complicated way. This may even involve my thinking about these properties as

[37] The rough picture here goes back by Russell (1905), though Russell largely focuses on singular thought (and singular assertion).

[38] This simplifies matters somewhat. You are not simply thinking about yourself as standing in the *thinking about* relation to these properties. You are thinking about yourself as having a single thought wherein you stand in the *thinking about* relation to these properties in a certain, unified way. A full account of what it is to think about something would need to provide details on the exact sense in which the way you are standing in the *thinking about* relation to these things is unified.

jointly participating in an event, where this amounts to my thinking about these properties as standing in a certain complex relation to *eventhood*.

Finally, the account I have in mind is not a disjunctive one. Or, more carefully: the account I have in mind is not simply one of the disjuncts of a disjunctive account of what it is to think, more generally. The account I have in mind says that, in general, we are always thinking about properties. It adds to this that, sometimes, when those properties are related to each other in certain ways, we are also thereby thinking about things that those properties exemplify. Say, for instance, that I am thinking about Ruth Bader Ginsberg. At the moment at which I am thinking about Ginsberg, I am also thinking about certain properties (*being a US Supreme Court Justice*, *wearing glasses*, *being quick-witted*, etc.).[39] So I stand in the *thinking about* relation to these properties, and I do so in a certain, unified way. Moreover, it is in virtue of my standing in the *thinking about* relation to these properties (and in a certain, unified way), along with these properties' being coexemplified, that I stand in the *thinking about* relation to Ginsberg herself. But my thinking about these properties is not identical to my thinking about Ginsberg.

The account I have just proposed is an *error theory* in that it is meant to explain away whatever it is that leads us, wrongly, to think that we can (and, sometimes, do) think about nonexistents.

Of course, there is much to be worked out with this error theory. And there are objections that this sort of error theory faces. While I do not have the space to offer a direct defense of this error theory here, I will briefly offer – or at least gesture at – an indirect defense of it.

My indirect defense of this error theory is as follows: the best alternative to this error theory is to say that Santa Claus, the fountain of youth, etc., really do exist even if they are not concrete.

[39] Of course, this simplifies things. I am really thinking about the United States Supreme Court, along with the relation of *being a member of*, and so on.

To make sense of this account, we should say that, in general, things exist necessarily but that some things are accidentally nonconcrete. We should add that, often, when we think we're thinking about nonexistent things, we are just thinking about accidentally nonconcrete things, and we are picking them out by the traits that they supposedly would have if they were concrete.[40]

It seems to me that anyone who accepts the above picture of thought should – on pain of being unprincipled, accept necessitism, the claim that, necessarily, whatever exists exists necessarily. Necessitism thus plays a key role in what I take to be the best alternative to the error theory I have proposed.

In the following chapter, I shall consider an argument for necessitism that, it may seem, my own arguments in Chapter Three make particularly compelling. If this is right, then my belief in the above-presented error theory runs the risk of being self-undermining. I shall assuage this worry by arguing that the argument for necessitism being discussed is itself self-undermining. Moreover, I shall show how my argument for this diagnosis can be extended to show that necessitism itself is false. If the argument here is successful, then the main alternative to the above-proposed error theory fails.

Conclusion

Thinking as Relational (TR) says that that, necessarily, whenever a thinker thinks about something, her doing so just is her standing in the *thinking about* relation (or *a thinking about* relation) to that thing. I have argued that Thinking as Relational follows from Content Fixes Truth (CFT), the claim that, necessarily, if two thoughts share representational content (of some sort), then they have the same truth value. I endorse Content Fixes Truth. Consequently, I endorse Thinking as Relational.

[40] This is only part of the story. It does not fully account for fictional objects or for impossibilia.

And, since I also endorse Serious Actualism, I am committed to the view that we cannot think about nonexistents. In light of this commitment, I have offered the bare bones of an error theory meant to explain what it is that we're thinking about when we think we're thinking about nonexistents. The central claim of this error theory is that, in these cases (as in all cases), we are thinking about properties.

Four

Necessitism and the Nature of Propositions

Necessitists claim that, necessarily, whatever exists exists necessarily. This claim may seem tough to swallow. But, according to Timothy Williamson, necessitism follows from a proper understanding of the nature of propositions. For, necessarily, if an individual A does *not* exist, the singular proposition *that A does not exist* is true. But propositions must exist to be true. And, at least for a *singular* proposition to exist, whatever that proposition is about must exist as well. Or so Williamson argues.

Williamson's argument may at first seem to introduce a worry about the particular combination of views and arguments I defended in the previous chapter. To begin, as I noted in Chapter Three, belief in necessitism threatens to undermine the motivation for the error theory that I sketched in the conclusion of that chapter. This may seem to be especially problematic for me since – as we'll see in what follows – Williamson's argument for necessitism at least *seems* to be much harder to resist for those who are persuaded by my arguments in Chapter Three. This gives rise to the worry that my arguments and assertions in the previous chapter threaten to undermine each other.

But this worry, it turns out, can be dismissed. For Williamson's argument is self-undermining. Consequently, it should not be persuasive for anyone, regardless of whether she accepts the arguments and assertions I defend in the previous chapter. Williamson's argument is self-undermining since, while there are various views on the nature of propositions one might sensibly adopt, any such view has the following feature: adopting that view either undercuts the motivation for a premise of Williamson's argument or commits one to denying necessitism. So, to

be motivated to accept necessitism on the strength of Williamson's argument, we must accept a view of propositions that commits us to denying the argument's conclusion.

Or so I shall argue in what follows. I shall then offer a sketch of how my argument can be extended to show that necessitism is false.

1 – The Argument from Propositions

Recall from Chapter Two, §1, that – while there is no single, universally-accepted *definition* of the term 'proposition' – philosophers (in the analytic tradition) generally agree that propositions (i) are the entities toward which our thought tokens are *attitudes* and (ii) are bearers of truth and falsity.[1] There's no consensus on what exactly this first feature amounts to. Of course, in previous chapters, I have laid out my own account of what exactly this amounts to. For present purposes, however, it will be enough simply to note that propositions, as the entities toward which our thought tokens are attitudes, are not *themselves* thought tokens (or token speech acts).

So propositions are bearers of truth and falsity and are the things toward which our thought tokens are attitudes. Propositions, as such, are *about* things.

Singular propositions are propositions that are *directly* about things. (See Chapter Two, §6.2.) For instance, plausibly, the proposition *that Michael Jordan is an avid golfer* is a singular proposition since, plausibly, it is directly about Michael Jordan.[2] On the other hand, the proposition *that the greatest professional basketball player of all time is an avid golfer* is not a singular proposition, for – while it is about Jordan – it is not directly about Jordan (or anything else, for that matter). Rather, it is about Jordan indirectly. It is about Jordan in virtue of its representing a

[1] See, e.g., John Perry (1979: 5 - 6), Robert Stalnaker (1981: 129 - 130), Timothy Williamson (2002: 235), Trenton Merricks (2015: 21), and Matthew McGrath (2014).
[2] To avoid inconsistency with quotations from others, I italicize the names of propositions and – rather than putting them in brackets – in this chapter.

certain feature of Jordan – in this case, the feature of being the greatest professional basketball player of all time.[3]

There are good reasons to think that, necessarily, a singular proposition exists only if whatever it is directly about exists as well. One such reason centers on what is arguably the orthodox view about the nature of singular propositions. According to *Russellianism* about singular propositions, singular propositions have whatever they are directly about as *constituents*.[4] Add to this the more-than-plausible view that something has to exist to be a constituent of something. It follows that, if Russellianism about singular propositions is true, then – necessarily – singular propositions exist only if the things they are directly about exist as well.

Here is another reason to think that, necessarily, a singular proposition exists only if whatever it is directly about exists as well. Recall from Chapter Three, §1, that Serious Actualism says that, necessarily, whatever stands in a relation (or otherwise exemplifies a property) exists.[5]

---

[3] See, e.g., Alvin Plantinga (1979: 109), Robert Adams (1981: 6-7), David Kaplan (1989: 483), Kit Fine (2007: 54), and Trenton Merricks (2015: 158) for this way of defining 'singular proposition'. Some philosophers define 'singular proposition' differently. For instance, Greg Fitch and Michael Nelson begin their 2014 *SEP* article "Singular Propositions" as follows: "Singular propositions (also called 'Russellian propositions') are propositions that are about a particular individual in virtue of having that individual as a direct constituent." This use of 'singular proposition' is common. But this is not how *I* use the term. (If I were to use the term this way, certain views I take to be substantial would come off as trivial.)

[4] As its name suggests, this view goes back to Russell (1903), though – from 1905 on – Russell himself would have denied that propositions like *that Michael Jordan is an avid golfer* are really singular. Contemporary philosophers who endorse the Russellian view of singular propositions include David Kaplan (1970, 1989: 494), Robert Adams (1981: 6 – 7), Nathan Salmon (1986: 1), Scott Soames (1987), David Braun (1993), and Jeffrey King (2007: 3, 76). (Alvin Plantinga (1983) endorses a related view, according to which singular propositions have, as literal constituents, the *essences* of whatever they are directly about.) Trenton Merricks – who does not endorse Russellianism – calls this view the "received view" about singular propositions. Merricks writes, "The received view about singular propositions is that each singular proposition has the relevant entity – the entity that that proposition is directly about – as a constituent" (2013: 158). (Those who follow Fitch and Nelson's use of 'singular propositions', as discussed in fn. 3, should characterize Russellianism about singular propositions as the view that "there are singular propositions.")

[5] Defenders of Serious Actualism include A.N. Prior (1967), Robert Adams (1981), Alvin Plantinga (1983), Christopher Menzel (1991), Timothy Williamson (2002: 240 – 241), and Trenton Merricks (2015). John Pollock (1985: 134 – 140), Nathan Salmon (1987: 95, 1998: 286), and Scott Soames (2008: 319) all reject Serious Actualism.

Plausibly, to explain why a certain singular proposition is directly about one thing as opposed to something else, we need to understand direct aboutness as being constituted by a certain relation between the proposition and the thing that the proposition is directly about. If so, and if Serious Actualism is true, then, necessarily, a singular proposition exists only if whatever it is directly about exists as well.[6]

Timothy Williamson agrees, for the above-presented reasons, that – necessarily – a singular proposition exists only if whatever it is directly about exists as well (2002: 240 – 244). And, in light of this, Williamson endorses the following argument:

(1) Necessarily, if Williamson doesn't exist, then the proposition *that Williamson doesn't exist* is true.

(2) Necessarily, if the proposition *that Williamson doesn't exist* is true, then the proposition *that Williamson doesn't exist* exists.

(3) Necessarily, if the proposition *that Williamson doesn't exist* exists, then Williamson exists.

Therefore,

(4) Necessarily, if Williamson doesn't exist, then Williamson exists. (1, 2, 3)

Therefore,

(5) Necessarily, Williamson exists. (4)[7]

[6] Robert Adams is the first to offer this argument from Serious Actualism (1981: 12). But the seeds of this argument are in Prior (1957: 29 – 40 and 1967: 149 - 151). And the related view that aboutness (direct or otherwise) is constitutively relational goes back at least to Franz Brentano (1874: 271 – 272). (Again, Brentano himself does not endorse this view of aboutness.)

[7] This is almost word-for-word how Williamson presents the argument (2002: 233 - 234). (Williamson uses 'I' rather than his name.) Williamson's argument is inspired by related arguments defended, for instance, by Prior (1957: 29 – 40 and 1967: 149 – 151), by Fine (1977: 149 – 150 and 1985: 160 – 180), and by Plantinga (1983: 9 – 10 and 1985: 341 – 349). Merricks also defends a related argument (2015: 157 – 190). What these arguments have in common is, with minor variations, a stock of mutually inconsistent claims

The conclusion of this argument is that Williamson exists necessarily, i.e., necessarily *has being*. (So 'exists' is being used in a very broad sense.) But nothing is special about Williamson, here. If this argument for Williamson's necessary existence is sound, then so is the analogous argument for my necessary existence, the analogous argument for your necessary existence, and so on. And if these arguments are all sound, then so are the corresponding arguments, for any possible world W, having to do with the existents of W. (Of course, if Williamson is right, these will be precisely the same arguments.) So, if the above argument is sound, then the following is true:

Necessitism: Necessarily, everything that exists exists necessarily.

In light of all this, we can treat the above argument as an argument for necessitism, in general. So I shall call the above argument "the Argument from Propositions for Necessitism" – for short, "the Argument from Propositions."[8]

Necessitism is counterintuitive, to say the least. For, intuitively, you and I exist but might not have. So, intuitively, there are contingent beings. So, intuitively, contingent beings are *possible*. That is, intuitively, the following is true:

Contingentism: Possibly, something exists that might not have existed.

from which they draw. Each argument results from identifying all but one of these claims as premises and then identifying the contradictory of the remaining claim as the conclusion.

[8] A view analogous to necessitism is *permanentism*, which states that "always everything is always something" (Williamson 2013: 4). As Williamson notes, there is – relative to the Argument for Necessitism – a "parallel proof" of permanentism (2002: 234). For this and other reasons, Williamson endorses permanentism (2002, 2013). And Meghan Sullivan endorses permanentism as well (2012). My arguments regarding Williamson's argument for necessitism can be applied to the parallel argument for permanentism. Likewise, my argument in §5 against necessitism itself can be applied to permanentism itself.

But Williamson and other necessitists – for instance, Bernard Linsky and Edward Zalta (1994: 445 - 450, 1996) – argue that necessitism is not as unpalatable as it may at first appear to be.[9] According to Linsky, Zalta, and Williamson, what makes us inclined against necessitism, initially, is the presumption that you and I and other concrete beings are *essentially* concrete, i.e., concrete in every possible world in which we exist. But – so they say – we are *not* essentially concrete. For, while we exist necessarily, we are only contingently concrete. And whatever may have at first lead us to hold that we are contingent beings may now be seen as a reason only to hold that we are contingently concrete. For instance, we may have thought that we are essentially human. But – according to Linsky, Zalta, and Williamson – we are not essentially human. Rather, we are (at least plausibly) essentially human-if-concrete.[10]

So necessitism is not as tough to swallow as it may at first seem to be. Moreover, at least initially, the Argument from Propositions seems well motivated. To begin, the argument is straightforwardly valid.[11] And premise (1) just seems intuitively true. And (2) follows from Serious Actualism. So (1) and (2) are both well motivated. And the motivation for (3) we've more or less already seen. Again, there are good reasons to think that, necessarily, a singular proposition exists only if whatever it is directly about exists as well. (3) follows from this view, along with the plausible claim that singular propositions are *essentially* directly about whatever they are *actually* directly about. So the Argument from Propositions at least *appears* to be well motivated.

---

[9] Linsky and Zalta would not call themselves necessitists, but here the difference seems to be just in terminology.

[10] Zalta would not put the point this way, however. As I am using the term 'essentially', an entity x is essentially F just in case, necessarily, if x exists, x is F. But others use 'essentially' differently. For instance, Zalta (2006) proposes a use of 'essentially' according to which, for any x and any F, x is essentially F iff_{df}, necessarily, if x is concrete, then x is F.

[11] As Williamson notes, (4) follows validly from (1), (2), and (3) given the transitivity of strict implication (2002: 234). And (5) follows validly from (4) since, in light of (4), Williamson's "nonexistence strictly implies a contradiction and is therefore impossible" (2002: 234).

Note, further, that the Argument from Propositions appears to be especially well-motivated for those who accept my arguments and assertions from the previous chapter. For, to begin, in the previous chapter I endorse Serious Actualism. Moreover, I argue in the previous chapter that, given CFT (which I endorse), we ought to characterize a subject's thinking about so and so (or such and such) as that subject's standing in the (or a) *thinking about* relation to so and so (or such and such). I then conclude from this, given Serious Actualism, that we cannot think about nonexistents. If my arguments here are sound, then there are analogous arguments – also sound – that lead us to the conclusion that propositions (singular or otherwise) cannot be about nonexistents. But, then, those who accept my arguments from previous chapters appear to have very few options for resisting Williamson's Argument from Propositions. Consequently, the Argument from Propositions should be especially compelling for those who accept my arguments from the previous chapter.

But here appearances are deceiving. For the Argument from Propositions is self-undermining. Consequently, regardless of what one thinks about the nature of propositions (or about Serious Actualism), one should not be (even defeasibly) moved to endorse the Argument from Propositions, unless one is just antecedently committed to necessitism.

I shall demonstrate this in what follows. To do so, it will be important that I treat as a live option each view of propositions that, it seems to me, *is* a live option, even setting aside my arguments from previous chapters. With this in mind, I shall now set aside the conclusions from previous chapters. (So, for instance, I shall remain agnostic over Serious Actualism and over the question of whether a singular proposition can exist only if whatever it is directly about exists.) I shall not return to them until the conclusion of this chapter.

My defense of the claim that the Argument from Propositions is self-undermining consists of two steps. The first step is to draw out a particular consequence of necessitism that should be of

independent interest to those who are, at the very least, open to necessitism. The second step will be to appeal to this consequence to show that the Argument from Propositions is self-undermining.

2 – Necessitism, Necessitarianism, and Events

Right now, it is 12:00 p.m., and I am at my kitchen table typing. But I could have been doing something else at this time, either at my kitchen table or elsewhere. So *necessitarianism* – the claim that whatever is the case *necessarily* is the case – is false. (And Williamson, Linsky, and Zalta – given that they think we are only contingently concrete – agree.) The following argument, then, is sound:

(6) If necessitism is true, then either there are no such entities as events or necessitarianism is true.

(7) Necessitarianism isn't true.

Therefore,

(8) If necessitism is true, then there are no such entities as events. (6, 7)

Call this argument "the Events Argument." The Events Argument is valid, and (7) is true.

And (6) is true as well. To see that (6) is true, consider the following principle:

Events: Necessarily, if events exist, then – for any entity x and property F, along with any spatial location l and time t – x is F at l and t if and only if the event of x's being F at l and t exists.

If Events is true, then so is (6). To see this, assume – for conditional proof – that Events is true. And say also that events exist. It then follows by Events that there exists the event of my typing at my kitchen table at 12 p.m. Now, for any event E, an essential feature of E is that it is an event.[12]

[12] Moreover, for any entity x, property F, location l, time t, and event E such that E is the event of x's being F at l and t, E is essentially the event of x's being F at l and t.

Say, then, that necessitism is true. Then, since there exists the event of my typing on my computer at 12 p.m. at my kitchen table, there necessarily exists this event. It then follows by Events that, necessarily, I am typing at my kitchen table at 12 p.m. More generally, consider any entity x and property F, along with any spatial location l and time t. If x is F at l and t, then – given Events – there exists the event of x's being F at l and t. So, given necessitism, there necessarily exists the event of x's being F at l and t. So, given Events, it's necessarily the case that x is F at l and t. So, to conclude: given the existence of events and the truth of necessitism (along with our assumption for conditional proof), necessitarianism is true – i.e., whatever is the case.[13] So, to discharge our assumption for conditional proof: (6) is true provided that Events is true.

And Events is true. Indeed, Events follows straightforwardly from a proper understanding of what it is for such and such an event to exist. For recall that, by 'exists', I just mean *has being*. And for there to be such and such an event just is for that event to occur. So for an event to exist just is for it to occur. (Indeed, Williamson himself suggests this point.[14]) It follows that, if Events is false, then – for some x, along with some property F, location l, and time t – at least one of the following is possible: either (a) x is F at l and t but the event of x's being F at l and t does not occur

[13] Let *weak necessitarianism* be the claim that – for any entity x, property F, location l, and time t – x is F at l and t only if, necessarily, x is F at l and t. One may object that, so far, I've shown only that, if Events is true, then weak necessitarianism is true as well. And weak necessitarianism is strictly weaker than full-blown necessitarianism, for weak necessitism doesn't account for the possibility that there are events that do not constitutively involve locations in space and time. Fair enough. But weak necessitarianism is no more plausible than necessitarianism. To keep things simple, I'll continue to focus on necessitarianism, but the reader is welcome to substitute my use of 'necessitarianism' for 'weak necessitarianism'.

[14] Williamson suggests this point in the context of warning his readers not to misinterpret necessitism for a claim about existence, in a much more restricted sense of 'exists'. Williamson writes the following: "…[O]ne might naturally say 'Events do not exist; they occur'. In saying so, one does not imply that there are no events, as one would on an unrestricted use of 'exist'" (2013: 18 – 19).

(even though, in general, events do occur), or (b) the event of x's being F at l and t occurs, but x is not F at l and t.[15] But neither (a) nor (b) is possible. Indeed, both (a) and (b) are incoherent.

Or so I say. But perhaps the necessitist will insist otherwise. I can think of one (and only one) way for the necessitist to do this. The necessitist can remind us of the important distinction between something's existing and its being concrete, and then the necessitist can insist that, while Events is false, the following related principle is true:

> **Events*:** Necessarily, if events exist, then – for any entity x and property F, along with any spatial location l and time t – x is F at l and t if and only if the event of x's being F at l and t exists *and is concrete*.

Assume that events exist, and recall that at 12 p.m. I am typing at my kitchen table. It follows, by Events*, that there exists the event of my typing at my kitchen table at 12 p.m. and that this event is concrete. Given necessitism, this event exists necessarily. Even still, the necessitist may say, this event is only contingently concrete. So there are possible worlds – among them, the actual world – in which this event is concrete, but there are also possible worlds in which this event is nonconcrete. In any possible world W in which this event is concrete, I am typing at my kitchen table at 12 p.m. But, in any possible world W* in which this event is *non*concrete, it is *not* the case that I am typing at my kitchen table at 12 p.m. Or so the necessitist may say.

This response on behalf of the necessitist fails. To see that it fails, assume, for *reductio*, that Events is false while Events* is true. The only principled way to account for this difference is to say that events exist necessarily but that many events are only contingently concrete (or contingently nonconcrete, as the case may be). Now consider the following:

[15] For simplicity's sake, I am treating the universal quantifier as ranging over a constant domain across possible worlds. Many contingentists will find this move problematic. The *necessitist*, however, will not (and cannot).

> **Kitchen.** At 12 p.m., I am at my kitchen table eating (not typing). Meanwhile, my friend Alex is at my kitchen table typing.

It follows from what we assumed for *reductio* that, in any possible world W in which Kitchen is true, the following is also true:

> **Kitchen Events.** The following events exist: (A) the event of my eating at my kitchen table at 12 p.m.; (B) the event of Alex's typing at my kitchen table at 12 p.m.; and (C) the event of my typing at my kitchen table at 12 p.m. All three of these events exist. But, while events A and B are concrete, event C is nonconcrete.

Now, there are possible worlds in which Kitchen is true. (At least, anyone who denies necessitarianism should say that there are possible worlds in which Kitchen is true.) So there are possible worlds – the same possible worlds – in which Kitchen Events is true. Consider one such possible world. Call it "W*." In W*, both Kitchen and Kitchen Events are true. So, in W*, I am sitting at my kitchen table eating. So, in W*, I exist and am concrete. And, in W*, the property of *typing* is instantiated. And the property of *typing* – whether or not *it* is concrete – is the sort of thing the instantiation of which is *caused* by something and *causes* things. And event C involves my typing at a particular *time* and at a particular physical *location*. And event C exists in W*. So, in W*, event C exists and is constituted by the instantiation of a property F by an individual x at a time t and location l, where the individual x in question is concrete in W* and where the property F in question is the sort of thing the instantiation of which is *caused* by something and *causes* things. Given all this, insofar as it makes sense to think of an event as concrete in the first place, event C surely qualifies as being concrete in W*. But, again, Kitchen Events is true in W*, and according to Kitchen Events event C is *non*concrete. So, in W*, event C both is and is not concrete. So there is a possible world in which a single event both is and is not concrete. But there is no such possible world. So we have a contradiction. So we must reject what we assumed for *reductio*. It is

not the case that Events* is true while Events is false. So the response we've been considering on behalf of the necessitist fails.[16]

Events is true. So (6) is true. And, again, (7) is true. So both premises of the Events Argument are true. And the Events Argument, again, is valid. So it is sound. So, if necessitism is true, then there are no such entities as events. Moreover, other than the examples I've used, nothing in my defense of the Events Argument relies on anything contingent. So, necessarily, the Events Argument is sound. So, necessarily, if necessitism is true, then there are no such entities as events. In other words, necessitism entails that there are no such entities as events.[17]

3 – Williamson's Argument and the Nature of Propositions

In light of what I've just argued, one might already be suspicious of the Argument from Propositions. For, if one is to find the Argument from Propositions persuasive, one must begin by accepting that there are such *entities* as propositions. But the ontological status of propositions is very much in dispute. On the other hand, most philosophers seem happy to accept the existence of

[16] Here is another way to put the point. It follows from Kitchen Events that the event of my typing at my kitchen table at 12 p.m. exists and, as such, occurs. And it also follows from Kitchen Events that I exist and am concrete. Moreover, typing is a kind of causal interaction with things that takes place in space and in time. All of this should lead us to believe that, insofar as it makes sense to say of *any* event that it is concrete, the event of my typing at my kitchen table at 12 p.m. is concrete, or at least it is concrete in any possible world in which Kitchen Events is true. But Kitchen Events also says that the event of my typing at my kitchen table at 12 p.m. is not concrete. So Kitchen Events implies a contradiction. So, necessarily, Kitchen Events is false.

[17] Williamson (2013) defends an argument that has much in common with the Events Argument. This is Williamson's argument for the claim that necessitism is inconsistent with the *truthmaker* principle, according to which "every truth is made true by something" (2013: 391). At the heart of Williamson's argument is the observation that the truthmaker principle and necessitism jointly entail necessitarianism (2013: 391 - 397). So we might expect Williamson to accept the Events Argument happily, given that it centers on an analogous observation. But Williamson seems at times to take for granted that events exist. (See, e.g., Williamson 2013: 18 - 19, 337, & 420.) And there is also an important difference between the Events Arguments and Williamson's argument for the inconsistency of necessitism and the truthmaker principle. Whereas the Events Argument crucially turns on Events, Williamson's argument for the inconsistency of necessitism and the truthmaker principle simply turns on the claim that, necessarily, if the truthmaker principle is true, then – for any proposition P – there is an entity x such that, necessarily, P is true if and only if x exists. And this is more or less just a definition of what the truthmaker principle is.

events. So Williamson's argument presupposes a controversial ontological claim even though, if his argument is sound, we must reject a much less controversial ontological claim. Of course, this is not in itself to show that we should reject the Argument from Propositions. But we should at least be suspicious of the argument.

We should also be moved to ask, "How should one's preferred view of the nature of propositions inform one's evaluation of Williamson's argument?" In what follows, I shall consider four rival (realist) views on the nature of propositions. For each such view, I shall argue that proponents of that view should find the Argument from Propositions to be, at best, unmotivated.

3.1 – Propositions as Sets of Possible Worlds

One prominent view of propositions takes it that propositions are sets of possible worlds. More precisely, for any proposition P, P is the set of all and only the possible worlds in which P is true.[18]

One inescapable implication of this view is that propositions represent what they do partly in virtue of how they are interpreted. My argument here is based on a related argument defended by Joseph Moore (1999), Michael Jubien (2001: 48 – 49), and Trenton Merricks (2015: 94 – 98).[19] In Chapter Two, §4, I presented an analogous argument having to do with the conception of

[18] Proponents of this view of propositions include Robert Stalnaker (1976: 79 – 80) and David Lewis (1979: 515, 1986: 53). (Some proponents of this view may wish to characterize a proposition P as the set of possible worlds *at* which P is true. For my present purposes, this difference is irrelevant. But see fn. 23.) A closely related view of propositions – which Stalnaker, in his 1981 (134), endorses – says that propositions are functions from possible worlds to truth values. More specifically, on this view, a proposition P is a function *I* mapping possible worlds to truth values such that, for any possible world W, *I*(W) = *true* (*false*) if and only if P is true (false) in (or at) W. (David Chalmers 1996: 63 – 65 also seems to endorse this view, but Chalmers – at least as of 2009 – is explicitly anti-realist about ontology.) In this section, to keep things simple, I'll argue explicitly that those who say that propositions are sets of possible worlds (and who are not already inclined to accept necessitism) should deny premise (1) of the Argument from Propositions, or at least they have no reason to accept this premise. But everything I have to say here regarding the view that propositions are sets of possible worlds may be said also of the view that propositions are functions from possible worlds to truth values. (For instance, see fn. 20).

[19] The argument by Moore, etc. is much more contentious, as it seeks to show that propositions are *not* sets of possible worlds (or functions from possible worlds to truth values, for that matter). (Also, cf. Paul Benacerraf 1965 for an argument targeting the view that *numbers* are sets of a certain sort.)

relativized propositions as sets of centered worlds. Consequently, the argument I present here will be in a somewhat condensed form.

Here is the argument. Assume that propositions are sets of possible worlds, as characterized above. And consider, for instance, the proposition *that snow is white*, which of course represents snow as being white. (Like Williamson 2002: 235, I follow the convention of naming propositions by what they represent.) On the view we're considering, this proposition is identical to the set of possible worlds in which snow is white. But, we might reasonably ask, why is it that *this* set of possible worlds is the one that represents snow as being white? Why isn't it, instead, that the proposition representing snow as being white is the set of all and only the possible worlds in which it is *not* the case that snow is white?[20] It seems the difference here is entirely arbitrary, unless of course we simply *interpret* the set of possible worlds in which snow is white as the set that represents snow as being white. So, if the proposition *that snow is white* really is the set of possible worlds in which snow is white, then that proposition represents whatever it represents (and even represents *at all*) at least partly in virtue of its being interpreted in a certain way. And my argument for this conclusion has nothing in particular to do with the proposition *that snow is white*, nor does it rely on any contingent claims. So, in general, on the view that propositions are sets of possible worlds, propositions are such that, necessarily, a proposition represents what it does (and represents *at all*) at least partly in virtue of its being interpreted in a certain way.[21]

[20] The analogous issue arises for the friend of propositions as functions from possible worlds to truth values. On this view, the proposition *that snow is white* is the function I_{SW}, such that, for any possible world W, $I_{SW}(W)$ = *true* if and only if, in W, snow is white. Why is it that snow is represented as being white by this function as opposed to, say, the function I_{NSW}, such that, for any possible world W, $I_{NSW}(W)$ = *true* if and only if, in W, it is *not* the case that snow is white?

[21] Lewis likely accepts (or likely would accept) this implication of the view that propositions are sets of possible worlds. For Lewis seems happy to allow that sets of possible worlds (and other mathematical entities) represent – or at least can represent – in virtue of how they are interpreted (1986: 144).

Here's why this matters. It shows us that, if we endorse the view that propositions are sets of possible worlds, we should be entirely unpersuaded by Williamson's argument for necessitism. For recall premise (1) of the Argument from Propositions:

(1) Necessarily, if Williamson doesn't exist, then the proposition *that Williamson doesn't exist* is true.

Now assume that propositions are sets of possible worlds. Then, in that case, propositions represent what they do partly in virtue of how we interpret them. Moreover, necessarily, propositions represent *at all* only if there are thinkers around to interpret them. Say, then, that there's a possible world W in which (a) Williamson does not exist and (b) there are no thinkers. In W, the proposition *that Williamson doesn't exist* is not interpreted, so in W the proposition *that Williamson doesn't exist* doesn't represent anything. (In fact, in W, the "proposition" in question doesn't even properly count as a proposition!) So *that Williamson doesn't exist* isn't true in W. Therefore, since Williamson doesn't exist in W, W is a world in which it is *not* the case that, if Williamson doesn't exist, the proposition *that Williamson doesn't exist* is true.[22] Now those who are inclined initially to reject necessitism (and who further adopt the view that propositions are sets of possible worlds) will say that *there is* such a possible world W. In this case, then, those who are inclined initially to reject necessitism (and who further adopt the view that propositions are sets of possible worlds)

[22] But, since propositions are named by what they represent, isn't it essential to *that Williamson doesn't exist* that it represents Williamson as not existing? No. At least, this doesn't follow. For when we decide to treat the clause 'that Williamson doesn't exist' as a name for the proposition that represents (i.e., that *actually* represents) Williamson's nonexistence, we are thereby appealing to a certain feature of the proposition in question (namely, to the feature of being the proposition that represents Williamson's nonexistence) in order to fix the referent of 'that Williamson doesn't exist'. If propositions really are sets of possible worlds, then in fixing the referent of 'that Williamson doesn't exist' in this way, we are picking out a particular set – i.e., the set of all and only the possible worlds in which Williamson doesn't exist. The expression 'that Williamson doesn't exist' then functions as a rigid designator, one that refers to the particular set in question across all possible worlds (or at least across all possible worlds in which the set exists), even if that set doesn't represent Williamson's nonexistence across all such worlds.

should reject (1). Now of course Williamson and other necessitists will maintain that, regardless of what we think about propositions, (1) is just trivially true since, for them, there is no possible world in which Williamson doesn't exist. But that's all beside the point. What matters is that those who say that propositions are sets of possible worlds *and* who are not already inclined to accept necessitism should deny (1), or at least they have no reason to accept it.[23]

Before moving on I want to discuss a couple presuppositions of the above argument, which presuppositions will run throughout the length of this chapter.

One presupposition is that the nature of propositions is not a contingent matter. So, for instance, if propositions are sets of possible worlds, then *necessarily* propositions are sets of possible worlds. My arguments could be complicated in ways to avoid this presupposition.[24] But

[23] A proposition P is true *in* a possible world W just in, were W actual, P would be true. A proposition P is true *at* a possible world W just in case whatever P (actually) represents as being the case is the case in W. For instance, assume, for the sake of demonstration, that there is a possible world in which there are no propositions (and, therefore, in which there are no true propositions). The proposition *that there are no propositions* is not true *in* such a possible world, but it is true *at* such a possible world. With this in mind, one who is committed to the claim that propositions represent what they do partly in virtue of how they are interpreted might deny (1) but nonetheless still accept the following:

> (1*) The proposition *that Williamson doesn't exist* is true *at* every possible world in which Williamson does not exist.

(See, e.g., Robert Adams 1981: 20 – 24 and Jeffrey King 2007: 8 - 96. Both Adams and King would accept (1*) while denying (1), though neither Adams nor King do so because they take propositions to represent what they do partly in virtue of how they are interpreted.) Accepting (1*), however, is of no use to the necessitist who wishes to argue for her view by defending an argument analogous to Williamson's argument from propositions. For, on that analogous argument, the second premise would be the following:

> (2*) The proposition *that Williamson doesn't exist* exists at every possible world at which it is true.

But this is false, or at least we have no reason to accept this unless we are just antecedently committed to the necessary existence of propositions.

[24] For instance, I have just argued, in effect, that if propositions are sets of possible worlds then, in a possible world W in which Williamson doesn't exist and there are no thinkers, there are no true propositions. But I could have argued, instead, that in a possible world W in which (i) propositions are sets of possible worlds, (ii) Williamson doesn't exist, and (iii) there are no thinkers, there are no true propositions.

such a complication, it seems to me, is unnecessary. For it seems implausible that the nature of propositions is only a contingent matter.

The other presupposition is that propositions, in general, represent things as being certain ways. But some may wish to deny that propositions represent things as being certain ways. I seem to owe these philosophers an argument for my view. Here is one argument. First, propositions are bearers of truth and falsity. Second, something has to represent things as being a certain way in order for it to be true or false. Conclusion: propositions represent things as being certain ways. And here is an argument for a related conclusion. To begin, singular propositions are those propositions that are directly about things. So singular propositions are about things. And a proposition is not the sort of thing that could just be about something without presenting it as being a certain way. So singular propositions represent things as being certain ways. And recall that, at the end of the day, all that really matters for us is what we should say of singular propositions – in particular, the singular proposition *that Williamson doesn't exist*.

Now I am persuaded by the above arguments. But say that you aren't. No matter. Again, propositions are bearers of truth and falsity. (As discussed in §1, this is a point of general agreement. More importantly, the fact that propositions are bearers of truth and falsity is an obvious presupposition of premise (2) of Williamson's argument.) We could easily modify my argument in this present section so that it relies not on the presupposition that propositions represent things as being certain ways but, rather, on the presupposition that propositions are bearers of truth and falsity. For instance, above I argued that, if propositions are sets of possible worlds, then either (i) it is objectionably arbitrary which propositions represent what or (ii) propositions represent what they do (and *at all*) in virtue of their being interpreted in certain ways. I could have argued instead that, if propositions are sets of possible worlds, then either (i*) it is

objectionably arbitrary which propositions have which truth conditions or (ii*) propositions have the truth conditions they have (and have truth conditions *at all*) in virtue of their being interpreted in certain ways. My argument would go through just the same.

Going forward, I'll continue to speak of propositions' representing things as being certain ways, but the reader is welcome to "translate" such talk as talk of propositions' having certain truth conditions.[25] My arguments will not turn on the difference here. (But, wherever the difference may *appear* to be relevant, I'll give an account – in a footnote – of how the argument could instead be put in terms of propositions' having certain truth values.)

3.2 – Propositions as Structured

Let us say that a proposition P is *structured* just in case (a) P has *constituents*, (b) P's constituents are *arranged* in a certain way, and (c) it is (at least partly) in virtue of P's having the constituents it has – and those constituents' being arranged in the way in which they are arranged – that P represents what it represents. One popular view of propositions is that they are structured.

We can divide defenders of structured propositions (roughly) into two camps. Russellians about propositions hold that propositions have, as constituents, the things that they are about (more carefully, the things they are *directly* about, along with the *properties* that they represent). Fregeans about propositions hold that the constituents of propositions are the *senses* (or *concepts*)

[25] A difficulty may seem to arise when it comes to *naming* propositions. For those who deny that propositions represent things as being certain ways cannot name propositions in terms of what they represent. The obvious alternative is to name propositions in terms of their truth conditions. There is a potential problem with this way of naming propositions: it seems to limit us to only one name for all necessarily true propositions (and one name for all necessarily false propositions). But those who hold that propositions do not represent things as being certain ways should probably endorse the proposition analogue to representationalism (as discussed in Chapter Two, §7.1, fn. 39) and, in light of this, say that there is only one necessarily true proposition (and one necessarily false proposition) anyway. (An exception may apply to those philosophers – e.g., Russell 1903 – who hold that propositions are events of a certain sort. But proponents of this view can appeal to Lagadonian representation. See §3.2, pg. 15.)

that represent the various things that the propositions are about. Consider, for instance, the proposition *that Bill loves Damon*. The Russellian about propositions says that this proposition consists of Bill, the *loving* relation, and Damon (arranged in a particular way). The Fregean about propositions says that this proposition consists of the senses Bill, loving, and Damon (arranged in a particular way), where Bill represents Bill, loving represents the relation (or activity) of loving, and Damon represents Damon.[26]

Here is one implication of the view that propositions are structured: either (a) propositions represent what they do (and *at all*) at least partly in virtue of their being interpreted in certain ways or (b) propositions are events of a certain sort.

To begin to see why there is this implication, consider the following argument – defended, for instance, by Trenton Merricks – against a particular version of the view that propositions are structured. On the particular version of this view in question, propositions are sets consisting of the things they are about, ordered in a particular way. So, for instance, the proposition *that A loves B* is an ordered set whose members include A, the *loving* relation, and B. But which set, exactly, is it? According to Merricks:

> The ordered set <A, *loving*, B> is neither a better nor a worse candidate for being the proposition *that A loves B* than is, for example, <B, *loving*, A> or <A, B,

[26] I'm simplifying things. First, there may be mild variations on how a given Russellian or Fregean about propositions understands what exactly the constituents of this proposition are. Second, one might plausibly hold that different sorts of propositions (e.g., singular propositions versus general propositions) are structured differently. Of course, for our purposes, what's relevant is one's view regarding *singular* propositions. See fn. 4 for defenders of the view that singular propositions are Russellian propositions. Defenders of singular propositions as Fregean include, in addition to Frege (1948, 1956), Gareth Evans (1982: 7 - 41) and John McDowell (1984). (Perhaps some will resist the characterization of certain Fregean propositions as singular since, for any Fregean proposition P about some entity x, P's being x is mediated by a certain sense, which P has and which represents x. But to say this is not to say that such a proposition P cannot be directly about P. For, as I discuss in §1, to say of a proposition P that it is directly about an entity x is to say that it is about x and is not about x in virtue of its representing a certain property. So it is a mistake to say that a Fregean proposition could not be a singular proposition. This mistake arises, I suspect, out of a tendency to characterize Fregeanism – inaccurately – as a kind of descriptivism. See Recanati 1993: Ch. 2 for a helpful discussion of this issue.)

> *loving*>. More generally, there will be many ordered sets that are all equally good candidates to be any given proposition (cf. Bealer, 1998, 6-7; Jubien, 2001; King, [2007], 7-8; Schiffer, 2003, 16). But the ordered set that is that proposition must be the best candidate for being that proposition. So it is false that any one ordered set is that proposition. (2015: 142)

And Merricks's argument generalizes to any ordered set taken to be a proposition. For any such set, it will be arbitrary that *it* is the proposition representing such and such whereas some distinct set with the same membership is not.

Plausibly, we can generalize on this argument much further. Note that the view of propositions discussed above is a version of the Russellian view of propositions. Merricks's argument may be generalized so as to target the analogous Fregean view as well. Here, too, there is no reason to say, for instance, that the proposition representing A's loving B is the ordered set <A, loves, B> as opposed to, say, the ordered set <B, loves, A> or <A, B, loves>. And we can generalize even further. We've so far only considered versions of the view that propositions are structured according to which constituency amounts to set membership. But, we may argue, the problem we've been discussing really has nothing to do with exactly how we understand what constituency amounts to. However exactly we understood what constituency amounts to (set membership, parthood, instantiation, or something else), there will always be certain arbitrary choices we'll have to make with regards to how the constituents of a given proposition are to be arranged. But, then, it will be arbitrary, at least to some extent, what represents what. And this is unacceptable.

Or so we may argue. There are two options for resisting this argument.

One option is to borrow a line from the friend of propositions as sets of possible worlds. One can say that propositions represent what they represent (and *at all*) partly in virtue of our interpreting them in certain ways. For instance, one might stipulate that, for any x and y and two-

place relation R, <x, R, y> is to be interpreted as the proposition that x R's y, whereas <y, R, x> is to be interpreted as the proposition that y R's x. In this case, it is not at all arbitrary what represents what.[27]

The other option is to say that propositions are events. In this case, to be the constituent of a certain proposition is to be a participant of a certain event. For instance, one might hold that the proposition *that A loves B* – i.e., the proposition that represents (at least in a Lagadonian sense of 'represents') A's loving B – is simply the event of A's loving B.[28] In this case, it is not at all arbitrary what represents what.

These are the only two plausible options for resisting the above argument against the view that propositions are structured.[29] I conclude, then, that the friend of structured propositions is committed to saying that propositions either (a) represent what they do (and *at all*) partly in virtue of their being interpreted in certain ways or (b) are events.

Here's why this matters. First, as we've already seen, if we take propositions to represent what they do (and *at all*) partly in virtue of how they're interpreted, then we shouldn't accept premise (1) of the Argument from Propositions, at least not unless we're antecedently committed

[27] David Lewis suggests an account of *singular* propositions along these lines (1986: 57 – 59). I hesitate to include Lewis among the other philosophers – listed in §1, fn. 4 – who are Russellians with respect to singular propositions. For Lewis is most often associated with the view – which I attribute to Lewis in §3.1, fn. 18 – that propositions are sets of possible worlds. There may seem to be an inconsistency, here, in Lewis's thinking. I suspect that Lewis would reply to this charge as follows: to say that propositions are entities of some kind K is really just to say that entities of that kind K perform (or can be interpreted as performing) certain *propositional roles*; and there is no inconsistency in taking it that entities of different kinds can perform (or can be interpreted as performing) different propositional roles.

[28] This seems to be Russell's (1903) view of propositions, though Russell does not explicitly use the term 'events' (1903: 139 - 140). (See Merricks 2015: 123 – 129 for discussion.) Jeffrey King also holds that propositions are events (or what he calls "facts"). But King would hold that the proposition *that A loves B* is not simply the event of A's loving B but, rather, is a much more complex event, one that involves A's and B's and *loving*'s all standing in certain semantic relations (2007: 24 – 64 (Ch. 2)).

[29] Merricks considers and argues against both of these responses (2015: 123 - 129 & 135 – 139). For our purposes, we needn't settle the question of whether Merricks's arguments here are successful.

to necessitism.[30] Second, as we've also already seen, necessitism entails that there are no such entities as events. So, if we take propositions to be events, then we should hold that necessitism is false and that the Argument from Propositions is unsound. Say, then, that we take it that propositions are structured. Then we shouldn't be (even defeasibly) motivated to accept necessitism by the Argument from Propositions.

3.3 – Propositions as Action Types

According to Peter Hanks (2013, 2015) and Scott Soames (2015), propositions are action types of a certain sort. For instance, the proposition *that snow is white* is an action type, tokens of which include (i) my believing – at 5:11 p.m., March 21, 2017 – that snow is white; (ii) your wondering – at 3:00 a.m., July 4, 2000 – whether snow is white; and so forth.[31]

If propositions are action types, then events exist, in which case necessitism is false. One way to see this is to look at how – if propositions are action types – propositions come to have the semantic features (e.g., representational features, truth conditions, and truth values) that they have. Here is what Hanks has to say about this:

> Propositions get their truth conditions from particular acts of judgment and assertion, which are themselves the original or primary bearers of truth and falsity. The source of truth conditions is to be found in the acts of representation we perform when we make judgments and assertions, not in the propositional contents we use to classify and individuate these actions. More precisely, the source is to be found in acts of *predication* through which, in the simplest cases, people attribute properties and relations to objects. The explanation for why propositions have truth conditions must appeal to these acts of predication. (2015: 3-4)

[30] Of course, I might have also argued – and in precisely the same way – that, in order to avoid the charge of arbitrariness with respect to a proposition's *truth conditions*, the friend of structured propositions must say either that (a*) propositions have the truth conditions they have partly in virtue of how they are interpreted or (b*) propositions are events.

[31] Recall from fn. 4 that Soames (1987) accepts Russellianism about singular propositions. Soames (2015) marks a departure from Soames (1987). (See, in particular, Soames 2015: 13, where Soames explicitly criticizes the view that propositions are structured.)

And Soames writes the following:

> In addition to assessing an agent's habitual or overall accuracy in representing his or her environment, we also need to assess the accuracy of the agent's sayings or cognitions, one by one. For this, we need truth and falsity, plus cognitive doings that *represent* things as being various ways, *where the sense in which they represent things is simply that performing them guarantees that agents represent those things.* When to perceive or think of o as P is to represent o as it is, we identify an entity – a particular sort of perceiving or thinking – plus a property that entity has when this sort of perceiving or thinking is accurate. The entity is a proposition, which is the cognitive act of representing o as P. The property is truth, which the act has iff to perform it is for one to represent o as o really is. (2015: 18)

So, for both Hanks and Soames, propositions' semantic features are derived from the related features of their instances or – at the very least – of their possible instances. But, then, at the very least, there are possible worlds in which these propositions have instances. And, since these propositions are action types, their instances are actions. And actions are events of a certain sort. So, if propositions are action types, and if Hanks and Soames are right about how these action types inherit their semantic features, then – at least *possibly* – there are events. And, if *possibly* there are events, then *actually* there are events. For things are happening all around us. I am (or at least was) typing. You are reading. And so on. Now, just to observe that these things are happening is not in *itself* to say that there are events. (See §5, fn. 46.) But if it is in principle possible for there to exist such things as events, then surely there exists (or existed) the event of my typing, the event of your reading, etc. So, if possibly events exist, then events actually exist. It follows that, if propositions are action types, then – if Hanks and Soames are right about how these action types inherit their semantic features – events exist. And, as we saw in §2, it follows from the existence of events that necessitism is false.

But say that propositions are action types but that Hanks and Soames are wrong about how these propositions inherit their semantic features. This seems incredibly unlikely to me, but no matter. For, surely, if action types exist, then it must be at least in principle possible for action types to have instances. (Note that this is weaker than the claim that, if action types exist, then – for every action type T – T possibly has an instance.) And, necessarily, the instances of action types are events. And, again, if it is in principle possible for events to exist, then certainly there actually are events. So, on the assumption that propositions are action types, events exist and necessitism is false.[32]

Say, then, that we endorse the view that propositions are action types. Then we should say that necessitism is false and that the Argument from Propositions is unsound.

3.4 –Propositions as Simple

Say that an entity is *simple* just in case it neither has *constituents* nor has *instances*. And say that an entity is *complex* just in case it is not simple. All of the views of propositions we've discussed so far say that propositions are complex. But one might instead hold that propositions – even singular propositions – are simple. As far as I know, only one contemporary philosopher – Trenton Merricks – explicitly endorses this view of propositions (2015: 205 - 207).[33]

Recall that premise (3) of the Argument from Propositions is the following:

[32] Actually, the move from the possible existence of events to the actual existence of events is overkill. For recall from §2 that, *necessarily*, if there are events, then necessitism is false. It thus follows from the possible existence of events that, possibly, necessitism is false. And necessitism is true only if it is necessarily true. So, if possibly there are events, then necessitism is false.

[33] An alternative view to Merricks's is the view that (a) propositions have logical form – e.g., the proposition *that Tammy is tall and Sherry is short* is of the logical form *P^Q*, whereas the proposition *that Tammy is tall* is only of the logical form *P* – and that (b) the simple propositions are all and only those propositions that are only of the logical form *P*. Strictly speaking, on this view the proposition *that Williamson does not exist* is not simple. Nevertheless, it should be clear from what follows that the relevant consequences I lay out for the view that propositions in general are simple applies equally to this alternative view.

(3) Necessarily, if the proposition *that Williamson doesn't exist* exists, then Williamson exists.

Let 'o' be a stand-in for any term that directly refers to some entity (i.e., to some existent), and let 'P(o)' be a stand-in for any 'that'-clause that includes the particular term for which 'o' is a stand-in. So o is some entity or other, and P(o) is some proposition or other that is directly about o. As Williamson notes, (3) is an instance of the following:

(3+) Necessarily, if P(o) exists then o exists.[34]

We should accept (3) only insofar as we are prepared to accept (3+). But belief in propositions as simple undermines the motivation for (3+), at least for those of us who are not antecedently committed to necessitism.

Or so I shall now argue. In arguing for this claim I shall make use of the notion of *ontological dependence*, characterized as follows:

> x *ontologically depends* on y iff$_{df}$, necessarily, (a) x exists only if y exists and (b), *were* y not exist, x would not exist, and x's nonexistence would be *explained* by y's nonexistence.

Three points of clarification are in order. First, unless I state otherwise, by 'explain' I mean *wholly explain*. (This also applies *mutatis mutandis* to the derivatives of 'explain'.) Second, given (b), condition (a) is redundant. I include condition (a) because I think it helps to make clear how condition (b) should be understood, but in what follows I shall often omit explicit reference to condition (a). Third and finally, the subjunctive conditional in the definition of ontological

[34] See Williamson (2002: 240). There are slight differences in how Williamson and I word (3+).

dependence should be interpreted as expressing a counterpossible. (In fact, subjunctive conditionals throughout this section should be interpreted in this way.)[35]

So understood, ontological dependence includes but is more than mere modal correlation. For instance, the number 2 exists necessarily. (Or so let us assume.) It follows that, necessarily, if I exist, the number 2 exists. But I am not ontologically dependent on the number 2. For even if it were the case (per impossibile) that neither I nor the number 2 existed, 2's nonexistence would not *explain* my nonexistence. On the other hand, assume that there are singleton sets, including the set {NKR}, i.e., the set whose sole member is me. It plausibly follows from this assumption that, necessarily: (a*) {NKR} exists only if I exist and (b*), were I not to exist, {NKR} would not exist, and {NKR}'s nonexistence would be explained by my nonexistence. That is to say: it plausibly follows that {NKR} ontologically depends on NKR.

Again, I shall argue that, for those of us who are not antecedently committed to necessitism, belief in propositions as simple undermines the motivation for (3+). My argument consists of two claims. First, in order for those of us who have no antecedent commitment to necessitism (and also have no antecedent commitment to the view that, necessarily, there are no such things as singular propositions) to be motivated to accept (3+), we must hold that P(o) ontologically depends on o. Second, if propositions are simple, P(o) does not ontologically depend on o.

To see that the former of these claims is true, consider the circumstances under which one could have genuine motivation for endorsing (3+). Consider, first, philosopher *A*, who thinks that (necessarily) there are no such things as singular propositions. *A* should say that the conditional *if*

[35] E.J. Lowe (1998: 145) initially proposes the appeal to explanation as a way of characterizing ontological dependence, though he does not ultimately adopt this proposal. Meanwhile, Benjamin Schnieder (2005) offers a rigorous characterization of ontological dependence in terms of explanation. My understanding of ontological dependence is importantly different from Schnieder's (and from Lowe's related proposal). But this is *not* to say that I am disagreeing with Schnieder (or with Lowe). The term 'ontological dependence' is a semi-technical term and, as such, can be used in many different ways.

P(o) exists then o exists has a necessarily false antecedent, and from this *A* should infer that (3+) is true.[36] Next, consider philosopher *B*, who is a necessitist. *B* should say that the conditional *if P(o) exists then o exists* has a necessarily true consequent, and from this *B* should infer that (3+) is true. But now consider philosopher *C*, who is not a necessitist and who does not believe that (necessarily) there are no such things as singular propositions. Philosopher *C* has motivation for accepting (3+) only insofar as she holds that, necessarily, were o not to exist, then P(o) would not exist either, and P(o)'s nonexistence would be *explained* by o's nonexistence. (Otherwise, it would be wholly arbitrary that, in any possible circumstance in which o does not exist, P(o) does not exist.) It follows that – for those of us who are neither antecedently committed to necessitism (and also have no antecedent commitment to the view that, necessarily, there are no such things as singular propositions) – to be motivated to accept (3+) we must hold that P(o) ontologically depends on o.[37]

---

[36] Recall from §1 that, following Williamson, I use 'exist' in a very broad sense. For me, to exist is simply to have being, i.e., to be an entity.

[37] In §1, I presented the motivation that Williamson (2002) offers for (3), which motivation, in effect, consists of two independent arguments for (3+), along with the observation that (3) is an instance of (3+). For those of us who have no antecedent commitment to necessitism (and also have no antecedent commitment to the view that, necessarily, there are no such things as singular propositions), to be motivated to accept either of these arguments for (3+) we must hold that P(o) ontologically depends on o. Here, to conserve space, I offer only a rough sketch of why this is.

To begin, both of these arguments for (3+) may be characterized as instances of the following schema:

> [3a+] Necessarily, P(o) – if it exists – is directly about o, and its being directly about o is constituted (at least in part) by its standing in a relation of class *S* to o.
>
> [3b+] Necessarily, if P(o) stands in a relation of class *S* to o, then o exists.
>
> Therefore,
>
> (3+) Necessarily, if P(o) exists then o exists.

S is a variable ranging over sets of relations. Let *the Argument from Russellianism* be the instance of the above schema for which the value of *S* is the set whose sole member is the *having as a constituent* relation. (Note that this argument is clearly at odds with the belief that propositions are simple.) Let *the Argument*

But, again, if propositions are simple, then P(o) does not ontologically depend on o. To begin to see why, imagine the following:

> **Blinking.** Presentism is true. In other words, that which exists *simpliciter* exists presently (if it exists in time at all). Now, at a time t_0, both P(o) and o exist. A few seconds past t_0, o goes out of existence. Then a few seconds later o comes back into existence. Then a few seconds later o goes out of existence. And the pattern continues until time t_1, which is one hour past t_0. So, over the course of one hour, o is "blinking" in and out of existence.

Now assume that Blinking is true. And assume, for *reductio*, that propositions are simple and that P(o) ontologically depends on o. Since P(o) is a proposition, P(o) is simple. So P(o) neither has constituents nor has instances. A *fortiori*, o is neither a constituent of P(o) nor an instance of P(o). So o is wholly distinct from P(o).[38] But, since presentism is true, it follows that, whenever o goes

---

from Serious Actualism be the instance of the above schema for which the value of *S* is the set containing all relations (and nothing else).

For those of us who have no antecedent commitment to necessitism (and also have no antecedent commitment to the view that, necessarily, there are no such things as singular propositions), to be motivated to accept any instance of this argument schema we must hold that P(o) ontologically depends on o. To see this, let S be some set or other over which *S* ranges, let the S Argument be the instance of the above schema for which the value of *S* is S, and let (3a+$_S$) and (3b+$_S$) be the S Argument's first and second premises, respectively. For those of us who are not antecedently committed to necessitism (and who are not antecedently committed to the view that, necessarily, there are no such things as singular propositions), to be motivated to believe that P(o) is *essentially* directly about o we must hold that it is simply part of the nature of P(o) that it is directly about o (i.e., that part of *what it is* to be the particular entity P(o) is to be directly about o). It follows that – for those of us who lack the above-mentioned antecedent commitments – to be motivated to accept (3a+$_S$) we must hold that, necessarily, *were* P(o) not to stand in a relation of class S to o, P(o) would not exist, and P(o)'s nonexistence would be *explained* by the fact that it does not stand in such a relation to o. Meanwhile, those of us who are not committed to necessitism have reason to accept (3b+$_S$) only insofar as we have reason to believe that, necessarily, were o not to exist, P(o) would not stand in a relation of class S to o. It follows, by the transitivity of (complete) explanation, that – if those of us who are not antecedently committed to necessitism (and who also are not antecedently committed to the view that, necessarily, there are no such things as singular propositions) are to be jointly motivated in accepting both premises of the S Argument – we must say that, necessarily, were o not to exist, P(o) would not exist, and in this case P(o)'s nonexistence would be explained by o's nonexistence. So, to generalize (and to apply the definition of 'ontological dependence'): the motivation for any instance of the above argument schema commits us to holding that P(o) ontologically depends on o, at least insofar as we have no antecedent commitment to necessitism (and also have no antecedent commitment to the view that, necessarily, there are no such things as singular propositions).

[38] Here I am assuming, for simplicity's sake, that P(o) is not identical to o (or a constituent – e.g., a proper part – of o). It is probably clear that, given the dialectic in which I am in engaged, this assumption is

out of existence, P(o) goes out of existence as well. But P(o) also comes back into existence when o comes back into existence. So P(o) blinks in and out of existence along with o. And, since P(o) ontologically depends on o, P(o) "blinks" in and out of existence precisely *because* o does.[39] To put the point (even more) metaphorically: P(o) follows o in and out of existence.

This consequence of Blinking is unacceptable. P(o) and o are wholly distinct entities. So it would be utterly arbitrary for P(o) to follow o in and out of existence. Now, perhaps, one might insist, P(o) and o stand in a certain relation to each other – the *being directly about* relation, say – when, and only when, they exist. (See fn. 37.) But it is implausible to think that this somehow could explain P(o)'s following o in and out of existence. For, given that o is neither a constituent of P(o) nor an instance of P(o), any such relation in which P(o) stands to o is *extrinsic* (as opposed to *intrinsic*) to P(o).[40] And, for any entity x and relation R, it is implausible to think that P(o)'s going in and out of existence can be *explained* by P(o)'s going back and forth from instantiating a relation R, provided that R is extrinsic to it. I conclude, then, that Blinking – when combined with our assumption, for *reductio*, that propositions are simple and that P(o) ontologically depends on o – has an unacceptable consequence.

---

harmless. However, for the moment let me just acknowledge this assumption and note that, in fn. 42, I discharge it.

[39] Technically, nothing in the definition of 'ontological dependence' guarantees that, in this case, P(o) comes back into existence when o does (and, furthermore, that P(o)'s coming back into existence is explained by o's coming back into existence). But, in the present context, it would be wholly unprincipled to deny this. For, again, P(o) and o both start out as existing, and the only variable is o's existence (or lack thereof). Moreover, anyone who accepts (3+) and who also accepts presentism ought to accept that, necessarily, P(o) exists only when o exists. And, insofar as P(o)'s going out of existence explains o's going out of existence, P(o)'s coming back into existence ought to be explained (at least in part) by o's coming back into existence.

[40] As I define 'extrinsic', a property F (which, again, may be either monadic or polyadic) is extrinsic to an entity x just in case x's exemplifying F (provided that x does exemplify F) constitutively involves some y's exemplifying some property G, where y is neither a constituent of x nor an instance of x (nor identical to x). As I define 'intrinsic', a property F is intrinsic to an entity x just in case x exemplifies F and F is not extrinsic to x. (Both 'extrinsic' and 'intrinsic' are semi-technical terms, and how I use the terms is slightly different from how others sometimes use them. For an overview of the standard uses of these terms as applied to properties, see Weatherson and Marshall 2014.)

Now, in fairness, Blinking is entirely made up. And it is controversial whether that which Blinking depicts is even possible (for some specification of what 'P(o)' and 'o' stand for). For it is controversial whether presentism is possibly true, and it is also at least somewhat controversial whether, even assuming the truth of presentism, things can go in and out of existence.[41] But whether or not Blinking is possible is not really the point. For Blinking is not *incoherent*. Imagine, then, that – perhaps *per impossibile* – Blinking is true. If we are to imagine further that propositions are simple and that P(o) ontologically depends on o, we find ourselves with the result that P(o) blinks in and out of existence precisely because o does. And, again, this result is unacceptable. Moreover – and this is the key – we don't need to endorse presentism or to deny that things can go in and out of existence to see that this result is unacceptable. What's unacceptable has nothing to do with o's existence or lack thereof; what's unacceptable is that P(o)'s existence is tied to o's in the way that our assumption for *reductio* requires it to be. In particular, what is unacceptable is that there is an explanatory link between o's going in and out of existence and P(o)'s going in and out of existence. So we should deny our assumption for *reductio*, even if we hold that Blinking is impossible. That is, we should deny the following conjunction: propositions are simple, and P(o) ontologically depends on o.

I conclude that, if propositions are simple, then P(o) does not ontologically depend on o. But recall that, for those of us who do not have an antecedent commitment to necessitism (and also do not have an antecedent commitment to the view that, necessarily, there are no such things as singular propositions), to be motivated to endorse (3+) we must say that P(o) ontologically depends

[41] To deny that objects can go in and out of existence even given the truth of presentism is to accept the necessitation of permanentism. (See fn. 8.) And, plausibly, permanentism is necessarily true if true at all. Moreover, necessitists should be especially attracted to permanentism. So it would be dialectically infelicitous for me to assume that, given presentism, objects can go in and out of existence.

on o. It follows that, for those of us who are not antecedently committed to necessitism, belief in propositions as simple undermines the motivation for (3+).

And recall, further, that to be motivated to accept (3) we must be motivated to accept (3+). So the belief in propositions as simple undermines the motivation for (3), at least for those of us who are not antecedently committed to necessitism.[42]

4 – The Argument from Propositions Is, at Best, Unmotivated

If Williamson (2002) is right, then reflection on the nature of propositions should lead us to endorse necessitism. If my arguments above are sound, then Williamson is not right.

For, to begin, among the four above-considered accounts of the nature of propositions, accepting any such account either undermines the motivation for one of the premises of the Argument from Propositions or commits us to the existence of events. And, if there are events, then necessitism is false and the Argument from Propositions is unsound. So, regardless of which of these four views of propositions one accepts, one should not be (even *defeasibly*) motivated by the Argument from Propositions to accept necessitism.

But might we adopt another realist view of the nature of propositions, one besides the four considered above? The problem is that there's very little conceptual space available for alternative views. And, for any sensible alternative view that *does* fall within this space, that view is bound to

[42] Recall from fn. 38 that, in arguing that if propositions are simple P(o) is not ontologically dependent on o, I made the assumption that o and P(o) are not identical (and that P(o) is not a constituent – e.g., a proper part – of o). I hereby discharge this assumption, and I now conclude that, for those of us who are not antecedently committed to necessitism, belief in propositions as simple undermines the motivation for (3+*), according to which, for any entity x and proposition P such that P is directly about x: if P is not identical to x (or to a constituent – e.g., a proper part – of x), then, necessarily, if P exists, then x exists. The proposition *that Williamson doesn't exist* is not identical to Williamson (nor is it a constituent of Williamson). So we should accept (3) only insofar as we are prepared to accept (3+*). So the belief in propositions as simple undermines the motivation for (3), at least for those of us who are not antecedently committed to necessitism.

lead us to the same conclusions to which the views we've already considered have lead us. For, propositions, if they exist, are either complex or simple. And, if they are complex, then either they have constituents or they have instances. Finally, if they have constituents, then those constituents either are arranged in a certain way, partly in virtue of which they represent whatever they represent, or they are not. We've already considered one view according to which propositions have constituents that are not so arranged. This is the view that propositions are sets of possible worlds. Proponents of any other (sensible) view according to which propositions have unarranged constituents (or constituents whose arrangements play no role in the propositions' representing what they do) face the same problem faced by those who say that propositions are sets of possible worlds. In particular, for proponents of such (potential) views, the only way to avoid its being objectionably arbitrary what propositions represent what is to insist that a proposition represents what it does at least partly in virtue of how it is interpreted.[43] And, as we've seen, those who do insist on this (or at least are committed to this) must deny – or at least have no reason to accept – premise (1) of the Argument from Propositions.

I thus conclude that any sensible realist view of the nature of propositions either undermines the motivation for one of the premises of the Argument from Propositions or commits us to denying necessitism. Of course, accepting that there are no such things as propositions also undermines the motivation for one of the premises of the Argument from Propositions. In particular, it undermines the motivation for (1). So any sensible view of the nature of propositions – including the view that there are no such things as propositions – either undermines the

[43] We've already seen that this is the case for a close cousin of the view that propositions are sets of possible worlds. This close cousin, again, is the view that propositions are functions from possible worlds to truth values. (See §3.1, fn. 18 and fn. 20.)

motivation for one of the premises of the Argument from Propositions or commits us to denying necessitism. So the Argument from Propositions is, at best, unmotivated.

5 – The Argument from Propositions is Unsound

I have argued that the Argument from Propositions is, at best, unmotivated. But, for all I've said so far, the Argument from Propositions is sound. For recall from §3.1 and §3.4, respectively, that if necessitism is true then premises (1) and (3) are *trivially* true. And (2) also is trivially true provided that necessitism is true. And the Argument from Propositions is valid. So, if necessitism is true, then the Argument from Propositions is trivially sound.

And, for all I've said so far, necessitism is true. Of course, there are *prima facie* compelling arguments against necessitism. But necessitists have a plausible way to resist these arguments – or at least the standard arguments – against their view. At the heart of this move, again, is the rejection of a key presupposition of the standard arguments – namely, the presupposition that whatever is concrete is essentially concrete.[44] I find this response from the necessitist reasonable. Moreover, while I have argued that the Argument from Propositions should not motivate one to be a necessitist, there are other candidate considerations that, at least for all I've said so far, *should* motivate one to accept (or at least to seriously consider accepting) necessitism. So I conclude that necessitism, at the very least, isn't *obviously* false.[45]

So, for all I've said so far, the Argument from Propositions is sound and necessitism true.

---

[44] Sullivan (2014: §3) proposes (but neither endorses nor rejects) an argument against necessitism that motivates the claim that whatever is concrete is essentially concrete. (The argument against necessitism that I am about to propose is independent of Sullivan's proposed argument.)

[45] For instance, some of the most *prima facie* attractive quantified modal logics seem to commit us to the truth of necessitism. See Linsky and Zalta (1994) and Williamson (1998, 2013) for defenses of necessitism based on this observation. For helpful (often critical) discussions of this sort of defense of necessitism, see Kripke (1963), Plantinga (1974), Fine (1977), Adams (1981), Marcus (1986), Salmon (1987), Menzel (1991, 2016), Deutsch (1991, 1994), and Sullivan (2014).

But I say that necessitism is false and that the Argument from Propositions is unsound. I say this because I endorse the following argument:

(8) If necessitism is true, then there are no such entities as events.

(16) There are such entities as events.

Therefore,

(17) Necessitism isn't true. (8, 16)

Call this argument "the Argument from Events against Necessitism" – for short, "the Argument from Events." The Argument from Events is clearly valid. And, as I argued in §2, (8) is true. So the Argument from Events is sound provided that (16) is true. I think that there are events. So I endorse (16). Consequently, I endorse the Argument from Events.

I think that there are events because I think that causation is a real phenomenon and that in order to account for this phenomenon we need to include events in our ontology.[46]

To see why I say this, consider the following:

> **Kick.** Mary roundhouse kicks the porcelain vase. Immediately after, and as a result, the vase shatters.

Kick involves an ordinary instance of causation (hence, the inclusion of 'as a result'). The correct analysis of Kick will be one that adequately and accurately characterizes what this instance of causation amounts to. The first candidate for such an analysis that comes to mind is the following:

[46] On the other hand, here is a bad argument for the existence of events. I am wearing a yellow shirt today. So the event of my wearing a yellow shirt today occurs. So the event of my wearing a yellow shirt exists. So there are events. The problem with this argument is that it is not at all obvious that the inference from the first claim to the second is valid. (Compare: this pen is blue; therefore, this pen instantiates the platonic universal *being blue*. It is not at all obvious that this inference is valid.)

> **Event Causation.** The event of Mary's roundhouse kicking the vase stands in the *being the cause of* relation to the event of the vase's shattering.

Event Causation entails that certain events are the relata of a certain relation. So Event Causation entails that there are events.[47] So, if the correct analysis of Kick is Event Causation, then there are events.[48] And the correct analysis of Kick either *is* Event Causation or, at the very least, is *relevantly similar* to Event Causation, where to say that an analysis of Kick is "relevantly similar" to Event Causation is to say that it, just like Kick, characterizes causation in terms of the instantiation of some relation or other between events. So, for instance, perhaps there is no single *being the cause of* relation. Perhaps, instead, there are just various causal relations. Then, plausibly, the correct analysis of Kick is one that differs from Event Causation, but only in that it identifies a *particular* causal relation – as opposed to simply the *being the cause of* relation – as being the relation in which the event of Mary's roundhouse kicking the vase and the event of the vase's shattering stand.[49]

Again, the correct analysis of Kick is Event Causation or, at the very least, something relevantly similar to Event Causation. Or so I say. And I say this because the most sensible alternative analyses are ones that, upon reflection, I reject.

---

[47] The entailment here is obvious given Serious Actualism. But even those who deny Serious Actualism should nonetheless say that there are properties (including relations) of certain sorts that can be instantiated only by things that exist. And causal relations are paradigmatic instances of such properties.

[48] Of course, I made up Kick. But feel free to replace Kick with an analogous story that represents something that is actually happening.

[49] Or perhaps, strictly speaking, there are no such entities as vases, only atoms arranged vase-wise. (See van Inwagen 1990: 98 - 107 (Ch. 10) and Merricks 2001: 2 – 8 (Ch. 1, §I).) In this case, the correct analysis of Kick will be much more involved than Event Causation. For, rather than involving the instantiation of a single relation between two events, both of which involve a vase, it will involve the instantiations of several relations between (or among) several different events, all of which involve, perhaps among other things, atoms that are (or at least recently have been) arranged vase-wise.

For instance, I reject counterfactual analyses of Kick (and of claims about causation, more generally). I reject, for instance, the following:

> **Counterfactual Causation.** There is some pair of times t and t* such that (i) t* is after t, (ii) Mary roundhouse kicks the vase at t, (iii) the vase shatters at t*, and (iv), in the nearest possible world in which Mary does not roundhouse kick the vase at t, the vase does not shatter at t*.[50]

Counterfactual Causation does not entail (or at least does not *obviously* entail) that there are causal relations and that the relata of these relations are (even in some cases) events. In fact, Counterfactual Causation does not entail (or at least does not *obviously* entail) that there are causal relations at all.[51] But Counterfactual Causation is not the correct analysis of Kick. Consider the question, "Why is it that, in the nearest possible world in which Mary does not roundhouse kick the vase at t, the vase does not shatter at t*"? An appropriate and *informative* response to this question is to say that Mary's roundhouse kicking the vase at t *causes* the vase to shatter at t*. But, then, Counterfactual Causation does not reveal to us what causation is; rather, it identifies, at best, a symptom of causation. So Counterfactual Causation does not adequately characterize what the instance of causation depicted in Kick amounts to. So it is not the correct *analysis* of Kick. And the same may be said of any other attempted counterfactual analysis of Kick. Such an attempted analysis will identify, at best, a *symptom* of causation. It will not, itself, provide an *analysis.*

[50] See Lewis (1973) for an early development of a counterfactual account of causation. (Lewis refines this account in, e.g., his 2004.) Ned Hall (2004) defends a two-factor view of causation, according to which there are two distinct phenomena worthy of the title 'causation', *one* of which can be successfully characterized by appeal to counterfactuals. Perhaps Hall is right. What I am defending, at present, is just the claim there is a *kind* of causation that cannot adequately be characterized by appeal to counterfactuals.

[51] Friends of counterfactual accounts of causation do often speak as if there are such things as events and as if they are the relata of causal relations. (See, for instance, L.A. Paul and Ned Hall (2013), as well as Peter Menzies's (2014) *SEP* article "Counterfactual Theories of Causation.") To my mind, however, one *needn't* posit such causal relations in adopting a counterfactual account of causation.

I also reject analyses of Kick according to which the relata of causal relations are entities other than events. For instance, I reject the following analysis of Kick:

> **Agent Causation.** Mary – perhaps in conjunction with certain properties she exemplifies – causes the vase to shatter. And Mary's causing the vase to shatter constitutively involves Mary's standing in a certain complex causal relation to the lamp, along with, perhaps, certain properties (plausibly, the *roundhouse kicking* relation and the *shattering* relation) and perhaps also to certain times and locations.

I reject Agent Causation because we cannot generalize from Agent Causation in the way we ought to be able to generalize, given that causation is a genuine phenomenon. One mark of a genuine phenomenon, it seems to me, is that we can characterize it in a way that is not wildly disjunctive. But, if we're to accept Agent Causation, we're committed to a wildly disjunctive analysis of what constitutes causation. To see what I mean, consider the following:

> **Brick.** Mary throws the brick into the porcelain vase. Immediately after, and as a result, the vase shatters.

Kick and Brick are sufficiently similar that we should expect any analysis of causation that (a) is not wildly disjunctive and that (b) posits genuine causal relations to be one according to which (c) both Kick and Brick can be understand as representing the instantiation of the same causal relation. So, if Agent Causation adequately explains what Kick represents, there ought to be a plausible explanation of what Brick represents that centers on the instantiation of the same relation that – according to Agent Causation – is instantiated in the scenario that Kick represents. But there is no such explanation. The closest we can get to one, it seems to me, is this:

> **Agent & Brick Causation.** Mary – along with the brick, and perhaps in conjunction with certain properties she exemplifies – causes the vase to shatter. And Mary's causing the vase to shatter constitutively involves Mary's standing in a certain complex causal relation to the brick and the lamp, along with, perhaps,

> certain properties (plausibly, the three-placed *throwing into* relation and the property of *shattering*) and perhaps also to certain times and locations.

But now let 'R*' name the complex relation in which – according to Agent Causation – Mary, etc. stand. And let 'R**' name the complex relation in which – according to Agent & Brick Causation – Mary, the brick, etc. stand. R* is not identical to R**. And to see this we don't need to know anything in particular about what exactly R* and R** are. It is enough to note that R** includes a *role* for the brick (or any other object one might throw into a vase in order to shatter it) whereas R* does not.[52] I conclude, then, that we cannot generalize from Agent Causation, as an analysis of Kick, to a plausible overall account of causation that is *not* wildly disjunctive.[53] And, again, the correct analysis of Kick will be one that *does* allow us to generalize to such an account. For causation, again, is a genuine phenomenon, and genuine phenomena can be adequately characterized by accounts that are not wildly disjunctive. So Agent Causation is not the correct analysis of Kick.

So I reject both Counterfactual Causation and Agent Causation as analyses of Kick. And, to my mind, Counterfactual Causation is the best candidate analysis of Kick that does not imply (or at least does not obviously imply) that there are causal relations. Likewise, to my mind, Agent Causation is the best candidate analysis of Kick according to which the relata of causal relations

[52] The reason is *not* – or at least not *just* – that the respective adicities of these two candidate relations will differ (though, almost surely, they will).

[53] On the other hand, Event Causation does seem generalizable in this way. The proponent of Event Causation as an analysis of Kick can analyze Brick as centering on the same relation on which Event Causation centers. For the proponent of Event Causation can analyze Brick as follows:

> **Brick-Event Causation.** The event of Mary's throwing the brick into the vase stands in the *being the cause of* relation to the event of the vase's shattering.

Both Event Causation and Brick-Event Causation characterize an instance of causation as constitutively involving one event's standing in the *being the cause of* relation to another event.

are not events. I conclude, then, that the analysis of Kick that we ought to accept is Event Causation, or at least is something relevantly similar to Event Causation. So I conclude that there are causal relations (or perhaps simply the single relation *being caused by*) and that events stand in these relations (or this relation). And, if events stand in causal relations, then there are events. So I conclude that there are events.

Again, the Argument from Events is valid, and – as we saw in §2 – (8) is true. And, as I've just argued, (16) is true as well. So the Argument from Events is sound.

Or so I have argued. The necessitist, of course, has resources to resist my defense of the Argument from Events, in particular by resisting my defense of (16). For instance, the necessitist might reasonably resist my critique of counterfactual analyses of causation. Or the necessitist might defend an analysis of Kick that succeeds where Agent Causation fails but that – like Agent Causation, and unlike Event Causation – posits, as the relata of causal relations, entities other than events.[54]

But notice that, whereas the necessitist does have some resources to resist my argument, my argument is not vulnerable to the necessitist's go-to response to standard arguments against necessitism. For, unlike these standard arguments, the Argument from Events does not presuppose that whatever is concrete is essentially concrete.

[54] See Paul and Hall (2013) and Menzies (2014) for a nuanced account of the challenges that counterfactual accounts of causation face, as well as for strategies for overcoming (or at least for *attempting* to overcome) these challenges. (My own challenge to such accounts, which is not at all novel, more or less amounts to a common, gut reaction to counterfactual accounts of causation. And many of the challenges that Paul and Hall consider and resist may be seen as resulting from attempts to develop genuine arguments out of this common, gut reaction.) See Jonathan Schaffer's 2016 *SEP* article "The Metaphysics of Causation" (in particular, §1) for an overview of popular, competing characterizations of the relata of causal relations. And see, once again, Paul and Hall (2013) for a helpful discussion of one of the main reasons to doubt that events are the relata of causal relations – namely, that there are many things that seem to be caused, at least in part, by *omissions*.

At any rate, I endorse the Argument from Events, and I do so for the reasons given. I conclude, then, that necessitism is false and that the Argument from Propositions is unsound.

Conclusion

My main goal in this chapter has been to demonstrate that the Argument from Propositions should not give us (even defeasible) motivation to accept necessitism. To demonstrate this, I have argued that – for any sensible view of the nature of propositions one might endorse – endorsing that view either undermines one's motivation for one of the premises of the Argument from Propositions or commits one to denying necessitism. A secondary goal in this chapter has been to provide a sketch of an argument against necessitism itself.[55]

If I am right that necessitism is false, then the Argument from Propositions is unsound, and – since it is valid – at least one of its premises is false. Given what we've learned from §3, *certain* views of propositions tell us which premise this is. For instance, if propositions are simple, then (3) is false. On the other hand, if propositions either are structured but are not events or, instead, are sets of possible worlds (or, for that matter, are entities of any other sort that have constituents), then propositions represent what they do (and *at all*) partly in virtue of how they are interpreted.

[55] Let *an* argument from propositions (for necessitism) be any argument that seeks to establish – from plausible considerations regarding singular propositions – that you and I exist necessarily and that, more generally, necessitism is true. If my arguments above are sound, then there are sound arguments – only superficially different from my arguments above – which show that *any* argument from propositions is unmotivated (and, beyond this, unsound). The truth of this conditional should be fairly obvious, at least *provided that* the above presentation of my arguments was clear and illuminating. Even still, in Appendix B, I shall consider an argument from propositions that – at least on the surface – differs significantly from the Argument from Propositions proper. I shall demonstrate that the arguments I've presented in response to the Argument from Propositions may be easily altered to arrive at arguments which reveal that this alternative argument from propositions is, just like the Argument from Propositions itself, unmotivated (and, for that matter, unsound).

And, if this is true, then (1) is false. And, of course, if there are no such entities as propositions, then (1) is false.[56]

Of course, my arguments and assertions from the previous chapter put further constraints on how we should characterize where the Argument from Propositions goes wrong. To begin, since I accept Serious Actualism, the following premise seems especially hard to deny:

(2) Necessarily, if the proposition *that Williamson doesn't exist* is true, then the proposition *that Williamson doesn't exist* exists.

Moreover, since I am committed to saying that propositions (singular or otherwise) can only be about what exists, I seem to be committed to the following premise:

(3) Necessarily, if the proposition *that Williamson doesn't exist* exists, then Williamson exists.[57]

That leaves me with denying (1). But, as we have seen, (1) is by no means undeniable.[58]

[56] The other views of propositions we've discussed do not, by themselves, tell us which premise of the Argument from Propositions is false. But, as we've seen, these views do imply – for reasons unrelated to my argument in §5 – that there are events, and in virtue of this they confirm that at least one of the premises of the Argument from Propositions is false.

[57] One might argue that I can resist (3) by denying that propositions essentially represent whatever they represent. But, first, denying this would undermine my motivation for (1). Second, see Appendix B for a related version of the argument that does not rely on the claim that propositions essentially represent whatever they represent.

[58] Recall that the first tenet of belief, which I endorse, entails that – necessarily – propositions are the fundamental bearers of content. As the argument presented in §3.1 should make clear, this rules out my conceiving of propositions as sets of possible worlds or as structured propositions that are not events. Bit leaves open to me the Russell (1903) conception of propositions as events as well as the conception of propositions as action types defended by Hanks (2013, 2015) and by Soames (Soames (2015). (See §3.2 and §3.3, respectively.) Note that in accepting either of these conceptions of propositions I have an independent reason for denying necessitism – namely, that necessitism is inconsistent with the existence of events.

Conclusion

The first tenet of the traditional doctrine of belief says that, necessarily, if two thoughts share representational content (of some sort), then they are attitudes toward – and, consequently, inherit the representational content of – the same proposition. I have argued from this ontological claim to an account of mental representation consisting largely of the following theses. First, Content Fixes Truth (CFT) is true and, consequently, the representational contents of mental states are wide. Second, our egocentric beliefs are attitudes toward private propositions. Third, for a thinker to think about something just is for her to stand in the (or a) *thinking about* relation to it. One consequence of this final claim, I have argued, is that we cannot think about nonexistents.

In addition to defending the above-presented account of mental representation, I have considered an argument for necessitism that, it would appear, is made particularly compelling against the backdrop of this defense. I have argued, however, that this argument for necessitism is self-undermining, and I have provided a sketch of how to extend my argument to show that necessitism is false. In doing so, I have provided indirect support for my preferred explanation of what we are thinking about when we think we are thinking about nonexistents.

It is worth emphasizing that the account of mental representation that I endorse is not a collection of disparate claims. Most obviously, of course, the three theses that largely constitute this account all follow from the first tenet of the traditional doctrine of belief. What's more, they all follow from this tenet in virtue of their following (either directly or indirectly) from the impossibility of relativized propositions. In addition, these three theses are mutually supportive in that we can often appeal to one or two of them in order to respond effectively to a worry faced by one of the others. We saw this, for instance, in Chapter Two, §7.4, when we saw that the appeal to

private propositions provides the externalist with a compelling response to the objection that externalism cannot adequately make sense of our epistemic limitations.

Looking forward, a broader, more ambitious project might draw out further ways in which these claims are mutually supportive.

Appendix A: Two Revised Arguments for CFT

Content Fixes Truth (CFT) says that, necessarily, if two thoughts share representational content (of any sort), then they have the same truth value. In Chapter One, I present two traditional arguments for CFT – the Argument from How Things Are and the Argument from How Things Could Be. I also show, in Chapter One, that both Jackson's and Chalmers's versions of internalism come with the resources to resist these two arguments.

It turns out that there are worries that these arguments face even if – as I argue in subsequent chapters – internalism is false. In this appendix, I shall draw on some of the lessons we learn from Jackson and Chalmers in Chapter One in order to present modified versions of these two arguments and to show what these arguments turn on. In doing so, I shall not rely on the defense of externalism I present in Chapter Two. For that defense of externalism relies on a substantive view about how thoughts relate to propositions. My goal in this appendix is to present what I take to be the most compelling arguments for externalism that do not rely on any view about how thoughts relate to propositions.

To this end, let's revisit the Argument from How Things Are. The Argument from How Things Are centers on two claims. First, necessarily, if two thoughts share representational content (of any sort), then they have the same truth value. Second, necessarily, thoughts that have the same truth conditions thereby have the same truth value.

One lesson we should all take away from Chalmers (2002a) is that the Argument from How Things Are, as it stands, is ambiguous. (See Chapter One, §6.3.) The source of the ambiguity is the phrase 'truth conditions'. By 'truth conditions', we might mean, e.g., *subjunctive truth conditions* or, instead, *epistemic truth conditions* (or, instead, *either subjunctive truth conditions or epistemic truth conditions*).

In fact, to go beyond Chalmers, there are even more options for what we might mean by 'truth conditions'. When philosophers use the term 'truth conditions', they often have something in mind besides subjunctive truth conditions or epistemic truth conditions. Often, they have in mind truth conditions of a sort that can be captured by certain biconditionals. For instance, if asked to give the truth conditions of my belief that snow is white, I might say, "My belief is true if and only if snow is white." [1] So we might also disambiguate the Argument from How Things Are, roughly, in terms of *truth conditions of the sort given by biconditionals* and, more precisely, in terms specific to how exactly we wish to characterize the relevant biconditionals.

I want to propose – as a descendent to the Argument from How Things Are – a disambiguation of this latter sort. But first I'd like to acknowledge a complication. I doubt that there is any way of making precise what sort of biconditionals one wishes to capture truth conditions in terms of that both (a) maps onto our intuitions about what are – and what are not – genuine truth conditions and (b) does not presuppose internalism or externalism. To see why I say this, consider Jackson's apple-on-head belief, i.e., Jackson's belief that – as Jackson himself would put it – "There is an apple on my head." Again, this belief is true (or so we are assuming). Intuitively, we cannot capture the truth conditions (in any recognizable sense of 'truth conditions') of Jackson's belief simply by appealing to *any* true material biconditional beginning with 'Jackson's belief is true if and only if'. For instance, intuitively, we cannot capture the truth conditions of Jackson's belief by appeal to the material biconditional 'Jackson's belief is true if

[1] I imagine that Jackson has in mind truth conditions of roughly this sort when he says that beliefs can have the same (representational) content while differing in truth conditions. (See Chapter One, §5.) For this reason, while it is true that Jackson and Chalmers differ with regards to where, for each of them, the Argument from How Things Are goes wrong, the difference doesn't stem from any philosophically substantial difference in their accounts of representational mental content. (This satisfies a promissory note made in Chapter One, §7.)

and only if snow is white'.[2] But the one clear alternative is to appeal, more specifically, to a material biconditional beginning with 'Jackson's belief is true if and only if' and ending with some phrase that perspicuously captures the content of Jackson's belief. However, internalists and externalists will disagree over which biconditionals meet this criterion. For instance, externalists will assert – and internalists will deny – that 'Jackson's belief is true if and only if Jackson has an apple on his head' meets this criterion.[3]

In light of this complication, my plan is simply to stipulate a use of 'truth conditions', in terms of biconditionals of a certain sort, and just not to worry that my use of 'truth conditions' does not map onto our intuitions about what do – and what do not – count as genuine truth conditions. In particular, I shall appeal to what I'll call "material truth conditions." A representation's *material truth conditions* are, roughly, those truth conditions given by a material biconditional of the form *R is true iff S*. More carefully, for a representation R, R's material truth conditions are given by all and only those material biconditionals that meet the following two criteria:

> (a) They are of the form *R is true iff S*, where 'S' is to be replaced by a sentence and '*R*' is to be replaced by a referring expression for R (and where the replacement for '*R*' contains no antecedent for any referring term in the replacement for '*S*').
>
> (b) They are true.

[2] And note that strict biconditionals are no better. In fact, they're worse. Consider, for instance, the strict biconditional 'Necessarily, my belief is true if and only if snow is white', where 'my belief' denotes the belief token consisting of my believing that snow is white. If we can appeal to *any* strict biconditional to characterize the truth conditions of my belief that snow is white, surely we can appeal to this one. But we cannot appeal to this biconditional for this purpose. For, possibly, snow is white but I don't believe it.

[3] Of course, internalists might happily acknowledge that – in some intuitive sense of 'truth conditions' – this biconditional can be used to characterize the truth conditions of Jackson's belief. But, then, they will deny that the principle that allows us to appeal to this biconditional is the following: *characterize the truth conditions of a thought by appeal to biconditionals ending with whatever expression most perspicuously captures that thought's content.* (Note also that appealing to a principle involving disquotation won't help since here we're trying to decide how to capture the truth conditions of a mental state, not a sentence.)

For instance, the material truth conditions of Jackson's apple-on-head belief (which, we're assuming, takes place at 3 p.m.) are given by the following material biconditional:

> Jackson's apple-on-head belief is true if and only if (at 3 p.m.) Jackson has an apple on his head.

Of course, the material truth conditions of Jackson's apple-on-head belief are also given by the following material biconditional:

> Jackson's apple-on-head belief is true if and only if (at 3 p.m.) snow is white.

In light of this example, one may insist that so-called "material truth conditions" are not really truth conditions, at least not in any intuitive sense of the phrase 'truth conditions'. That's fine. My use of 'truth conditions', as in 'material truth conditions', is just by stipulation.

With all this in mind, consider the following argument:

> (1) Necessarily, if two thoughts share representational content (of any sort), then they have the same material truth conditions.
>
> (2) Necessarily, if two thoughts have the same material truth conditions, then they have the same truth value.
>
> Therefore,
>
> (3) **CFT:** Necessarily, if two thoughts share representational content (of any sort), then they have the same truth value.

I suggest that we replace the original argument form How Things Are with the above argument, which I'll call "the Revised Argument from How Things Are." Since (1) is explicitly about

material truth conditions, the Revised Argument from How Things Are is not ambiguous in the way that the original Argument from How Things Are is (or in any other way, for that matter).

The Revised Argument from How Things Are is sound provided that (1) is true. For, to begin, the argument is clearly valid. And (2) is true. Given my above characterization of a representation R's material truth conditions, two representations R* and R** have the same material truth conditions just in case, for any material biconditional that gives one of these representations' material truth conditions, there is a material biconditional that gives the other representation's material truth conditions such that these two material biconditionals differ only with respect to which of the two representations – either R* or R** – is being referred to. It follows that, for any two representations R* and R** with the same material truth conditions, R* is true if and only if R** is true. *A fortiori*, for any two thoughts T* and T** with the same material truth conditions, T* is true if and only if T** is true. And since the argument I've just given for this claim relies on nothing contingent, we can conclude the following: necessarily, if two thoughts have the same material truth conditions, then they have the same truth value. So (2) is true. And, again, the Revised Argument from How Things Are is valid. So the Revised Argument from How Things Are is sound provided that (1) is true.

Turn now to the Argument from How Things Could Be. Again, this argument begins with representationalism, which says – roughly – that we can identify the representational content of a thought by identifying which possibilities the thought is consistent with and which possibilities the thought rules out. The argument precedes by presenting a particular way in which this representationalist picture of content may be fleshed out – namely, the way that Stalnaker fleshes it out – and then infers from this specific version of representationalism that CFT is true.

So far we've only considered one strategy for resisting the Argument from How Things Could Be: the strategy given to us by Jackson and Chalmers. Both Jackson and Chalmers accept representationalism yet reject Stalnaker's version of it, preferring instead a version of the view that does not entail – and, in fact, contradicts – CFT.

But there is another promising strategy for resisting the Argument from How Things Could Be: reject representationalism. There's a compelling objection to representationalism that has nothing to do with its consequences for internalism. Representationalism commits us to individuating mental states, in terms of their content, in a way that – intuitively – is too course grained. If representationalism is true, then two thoughts differ in representational content if *and only if* there is some possibility with which one thought is consistent and which the other thought rules out. It thus follows from representationalism that all necessarily true thoughts (more carefully: all thoughts that represent necessary truths) have the same representational content as each other. Likewise, it follows from representationalism that all necessarily false thoughts (more carefully: all thoughts that represent impossibilities) have the same representational content as each other. More generally, it follows from representationalism that – for any two thoughts T* and T** that are necessarily coextensive – T* and T** have the same representational content. *Prima facie*, there's an abundance of counterexamples to this claim. For instance, *prima facie*, my belief that 2 + 2 = 4 differs in representational content from your belief that all bachelors are unmarried; but these two beliefs are both necessarily true (and, consequently, are necessarily coextensive). *Prima facie*, my uncle's belief that there is a highest prime number differs in representational content from my barber's belief that he has the ability to shave all and only those who don't shave themselves; but these two thoughts are both necessarily false (and, consequently, are necessarily coextensive). *Prima facie*, my belief that Angela Merkel is the chancellor of Germany differs in

content from your belief that either Merkel is the chancellor of Germany or there are married bachelors; but these two beliefs are necessarily coextensive. The list goes on.[4]

I want now to propose an alternative to the original Argument from How Things Could Be that is not vulnerable to the objection we've just discussed. I have in mind the following:

> (1*) Necessarily, if two thoughts share representational content (of any sort), then they have the same subjunctive intension (/subjunctive truth conditions).
>
> (2*) Necessarily, if two thoughts have the same subjunctive intension (/subjunctive truth conditions), then they have the same truth value.
>
> Therefore,
>
> (3) **CFT:** Necessarily, if two thoughts share representational content (of any sort), then they have the same truth value.

Call this argument "the Revised Argument from How Things Could Be." The Revised Argument from How Things Could Be does not presuppose representationalism. Of course, premise (1*) makes appeal to subjunctive intensions, which at least one representationalist – namely, Stalnaker – identifies with representational thought contents (and which another representationalist –

[4] For responses to this objection in defense of representationalism, see Stalnaker (1984: 23 – 25), Jackson (2003b: 101 – 102), and Chalmers (2002a: 619 & 2011). Given our usual way of speaking about possibility and necessity, it may seem that the objection applies only to those versions of representationalism that characterize representational content in terms of the *metaphysical* possibilities with which a thought is consistent. But the objection also applies to those versions of representationalism that characterize representational content in terms of the *epistemic* possibilities with which a thought is consistent. (It is also applies to dual content versions of representationalism, but in this case we'd need to complicate the presentation of the argument a little.) Consider, for instance, the above example of two necessarily true thoughts that, *prima facie*, have different representational contents. These thoughts of course are both *metaphysically* necessary, but they are also (in Chalmers's sense) *epistemically* necessary in that they are both in principle knowable *a priori*. Consider, also, the above example of two necessarily false thoughts that, *prima facie*, have different representational contents. These thoughts of course are both metaphysically impossible, but they are also epistemically impossible in that what each represents as being the case is something we can know *a priori* not to be the case. Consider, finally, my belief that Angela Merkel is the chancellor of Germany and your belief that either Merkel is the chancellor of Germany or there are married bachelors. These two beliefs are epistemically-necessarily coextensive in the sense that, for any epistemic possibility, either both beliefs are consistent with that epistemic possibility or neither belief is.

namely, Chalmers – identifies with representational thought contents of a certain *sort*.) But (1*) does not require that a thought's representational content *be* its subjunctive intension. It does not require that, necessarily, for any two thoughts, those thoughts share representational content (of some sort) *if and only if* they have the same subjunctive intension. (1*) requires only that, necessarily, for any two thoughts, those thoughts share representational content (of some sort) *only if* they have the same subjunctive intension. So (1*) does not presuppose representationalism. Nor does any other part of the Revised Argument from How Things Could Be.[5]

If (1*) is true, then the Revised Argument from How Things Could Be is sound. To begin, the argument is valid. And (2*) is true. Say that two thoughts T* and T** have the same subjunctive intension (/subjunctive truth conditions). Then, for any possible world W, T* is true at W, considered as counterfactual, if and only if T** is true at W, considered as counterfactual. For any thought T, T is true (i.e., is true *simpliciter*) if and only if – for whichever possible world W* is actual – T is true at W*, considered as counterfactual. Since T* and T** have the same truth value at every possible world, considered as counterfactual, they have the same truth value at the possible world W*, considered as counterfactual. So T* is true if and only if T** is true. So, to generalize: if two thoughts have the same subjunctive intension (/subjunctive truth conditions), then they have the same truth value. Moreover, since my argument for this conclusion relies on nothing contingent, we can further conclude that – *necessarily* – if two thoughts have the same subjunctive intension (/subjunctive truth conditions) then they have the same truth value. So (2*) is true. And, again, the Revised Argument from How Things Could Be is valid. So the argument is sound provided that (1*) is true.

[5] Note that (1*) is also consistent with the claim, which I defend in §6 and §7 of Chapter Two, that egocentric thoughts are attitudes toward private propositions.

I've just proposed that we replace the two standard arguments for CFT – at least as these arguments are presented in §4 of Chapter One – with two closely related arguments. I've also demonstrated what these proposed replacements turn on. The Revised Argument from How Things Are is sound if (and only if):

> (1) Necessarily, if two thoughts share representational content (of any sort), then they have the same material truth conditions.

Meanwhile, the Revised Argument from How Things Could Be is sound if (and only if):

> (1*) Necessarily, if two thoughts share representational content (of any sort), then they have the same subjunctive intension (/subjunctive truth conditions).

It follows from what I've demonstrated that CFT is true provided that either (1) or (1*) is true.[6]

[6] Of course, plausibly, (1) is true if and only if (1*) is. And, of course, I also give an independent argument in favor of CFT in Chapter Two.

Appendix B: Another Argument from Propositions

Consider the following argument:

(1/2*) Necessarily, if Williamson doesn't exist, then there is a proposition P such that (i) P is directly about Williamson, (ii) P represents Williamson as not existing, and (iii) P is true.

(3*) Necessarily, if there is a proposition P such that (i) P is directly about Williamson, then Williamson exists.

Therefore,

(4) Necessarily, if Williamson doesn't exist, then Williamson exists. (1/2*, 3*)

Therefore,

(5) Necessarily, Williamson exists. (4)

Once again, nothing is special about Williamson, here. If the above argument for Williamson's necessary existence is sound, then so is the analogous argument for my necessary existence, the analogous argument for your necessary existence, and so on. And if these arguments are all sound, then so are the corresponding arguments, for any possible world W, having to do with the existents of W. So, if the above argument is sound, then necessitism is true. So I shall call the above argument "the Other Argument for Necessitism from Propositions" – for short, "the Other Argument from Propositions."

The Other Argument from Propositions is, at best, unmotivated. And my account of why this argument is unmotivated is analogous to my account – presented in §2 through §4 of Chapter Five – of why the Argument from Propositions is unmotivated. Among the various plausible accounts of the nature of propositions, accepting any such account either undermines the motivation for one of the premises of the Other Argument from Propositions or commits us to the

existence of events. Consider, first, the view that propositions are simple. Those who hold this view (and who are not antecedently committed to necessitism) should deny (3*), or at the very least they have no reason to accept (3*). Consider, next, the view that propositions are structured but are not events, as well as the view that propositions are sets of possible worlds (or, for that matter, are entities of any other sort that have constituents). Again, those who accept these views should say that propositions represent what they do (and *at all*) only insofar as they are interpreted in certain ways. So those who accept these views (and who are not antecedently committed to necessitism) should deny premise (1/2*), or at least they have no reason to accept it. Now consider the view that propositions are structured and, furthermore, are *events*, along with the view that propositions are action types. Proponents of either of these views should say that there are events. Finally, those who simply deny that there are such things as propositions (and who are not antecedently committed to necessitism) should say that (1/2*) is false, or at least they have no reason to accept (1/2*). It follows that, for any plausible view of the nature of propositions we might have, either we should be unmoved by the Other Argument from Propositions or we should accept the existence of events. But, again, if there are events, then necessitism is false. So, if we're to be moved by the Other Argument from Propositions to accept necessitism, we must accept a view that entails that necessitism is false. So the Other Argument from Propositions is, at best, unmotivated.

Of course, simply in arguing that the Other Argument from Propositions is unmotivated, I have not thereby argued that it is unsound. But the Other Argument from Propositions is unsound. For, as I argue in §5 of Chapter Five, necessitism is false.

Since the Other Argument from Propositions is unsound but is nonetheless valid, at least one of the premises of the Other Argument from Propositions is false. And *certain* views of

propositions tell us which premise this is. For instance, if propositions are simple, then (3*) is false. On the other hand, if propositions either are structured but are not events or, instead, are sets of possible worlds (or, for that matter, are entities of any other sort that have constituents), then propositions represent what they do (and *at all*) partly in virtue of how they are interpreted. In this case, (1/2*) is false. And, of course, (1/2*) is also false if there just are no such entities as propositions.

Bibliography

Adams, Robert Merrihew (1981). Actualism and Thisness. *Synthese* 49 (1): 3 - 41.

Azzouni, Jody (2010). *Talking About Nothing: Numbers, Hallucinations, and Fictions*. Oxford: Oxford University Press.

Bach, Kent (1987). *Thought and Reference*. Oxford: Oxford University Press.

Barnes, Elizabeth & Cameron, Ross (2009). The Open Future: Bivalence, Determinism and Ontology. *Philosophical Studies* 146 (2): 291-309.

——— (2011). Back to the Open Future. *Philosophical Perspectives* 25 (1): 1-26.

Bealer, George (1998). Propositions. *Mind* 107 (425): 1-32.

Benacerraf, Paul (1965). What Numbers Could Not Be. *Philosophical Review* 74 (1): 47-73.

Boghossian, Paul A. (1989). Content and Self-Knowledge in Philosophy of Mind. *Philosophical Topics* 17 (1): 5-26.

Bourget, David & Chalmers, David J. (2014). What Do Philosophers Believe? *Philosophical Studies* 170: 465-500.

Braun, David (1993). Empty Names. *Noûs* 27 (4): 449-469.

Brentano, Franz (1874/2005). *Psychology from an Empirical Standpoint.* Translated by Antos C. Rancurello, D.B. Terrell, and Linda L. McAlister. London and New York: Routledge.

Brogaard, Berit (2012). Context and Content: Pragmatics in Two-Dimensional Semantics. In Keith Allan & Kasia Jaszczolt (eds.), *Cambridge Handbook of Pragmatics.* Cambridge: Cambridge University Press. 113-134.

Brown, Curtis (2016). Narrow Mental Content. *The Stanford Encyclopedia of Philosophy* (Summer 2016 Edition), Edward N. Zalta (ed.), URL = {https://plato.stanford.edu/archives/sum2016/entries/content-narrow/}.

Burge, Tyler (1979). Individualism and the Mental. *Midwest Studies in Philosophy* 4 (1): 73-122.

Cameron, Ross P. (2015). *The Moving Spotlight: An Essay on Time and Ontology*. Oxford: Oxford University Press.

Cappelen, Herman & Dever, Josh (2013). *The Inessential Indexical: On the Philosophical Insignificance of Perspective and the First Person*. Oxford: Oxford University Press.

Capps, David; Lynch, Michael P. & Massey, Daniel (2009). A Coherent Moral Relativism. *Synthese* 166 (2): 413 - 430.

Cargile, James (1979). *Paradoxes, a Study in Form and Predication*. Cambridge: Cambridge University Press.

Cartwright, Richard (1960). Negative Existentials. *The Journal of Philosophy* 57 (20/21): 629-639.

Castañeda, Héctor-Neri (1972). *Thinking and the Structure of the World: Discours d'Ontologie*. *Crítica* 6 (18): 43-86.

Chalmers, David J. (1996). *The Conscious Mind: In Search of a Fundamental Theory*. Oxford: Oxford University Press.

———— (2002a). The Components of Content (Revised Version). In David Chalmers (ed.), *Philosophy of Mind: Classical and Contemporary Readings*. Oxford: Oxford University Press. 608-633.

———— (2002b). On Sense and Intension. *Philosophical Perspectives* 16 (16): 135-182.

———— (2002c). Does Conceivability Entail Possibility? In Tamar S. Gendler & John Hawthorne (eds.), *Conceivability and Possibility*. Oxford: Oxford University Press. 145-200.

———— (2003). The Nature of Narrow Content. *Philosophical Issues* 13 (1): 46-66.

———— (2006). The Foundations of Two-Dimensional Semantics. In Manuel Garcia-Carpintero & Josep Macia (eds.), *Two-Dimensional Semantics: Foundations and Applications*. Oxford: Oxford University Press. 55-140.

———— (2009). Ontological Anti-Realism. In David John Chalmers, David Manley & Ryan Wasserman (eds.), *Metametaphysics: New Essays on the Foundations of Ontology*. Oxford: Oxford University Press. 77-129.

———— (2011a). The Nature of Epistemic Space. In Andy Egan & Brian Weatherson (eds.), *Epistemic Modality*. Oxford: Oxford University Press. 60-107.

———— (2011b). Propositions and Attitude Ascriptions: A Fregean Account. *Noûs* 45 (4): 595-639.

Chisholm, Roderick M. (1973). Beyond Being and Nonbeing. *Philosophical Studies* 24 (4): 245 - 257.

Crane, Tim (1991). All the Difference in the World. *Philosophical Quarterly* 41 (January): 1-25.

——— (2001). *Elements of Mind: An Introduction to the Philosophy of Mind*. Oxford: Oxford University Press.

——— (2012a). What is the Problem of Non-Existence? *Philosophia* 40 (3): 417-434.

——— (2012b). Existence and Quantification Reconsidered. In Tuomas E. Tahko (ed.), *Contemporary Aristotelian Metaphysics.* Cambridge: Cambridge University Press.

——— (2013). *The Objects of Thought*. Oxford University Press.

Davies, Martin & Humberstone, Lloyd (1980). Two Notions of Necessity. *Philosophical Studies* 38 (1): 1-31.

Deutsch, Harry (1990). Contingency and Modal Logic. *Philosophical Studies* 60 (1-2): 89-102.

——— (1994). Logic for Contingent Beings. *Journal of Philosophical Research* 19: 273-329.

Egan, Andy (2006). Secondary Qualities and Self-Location. *Philosophy and Phenomenological Research* 72 (1): 97-119.

Einheuser, Iris (2008). Three Forms of Truth-Relativism. In Manuel Garcia-Carpintero & Max Kölbel (eds.), *Relative Truth*. Oxford: Oxford University Press. 187-203.

——— (2012). Relativized Propositions and the Fregean Orthodoxy. *Philosophy and Phenomenological Research* 84 (3): 590-603.

Evans, Gareth. (1979). Reference and Contingency. *The Monist* 62 (2): 178-213.

——— (1982). *Varieties of Reference*. Oxford: Oxford University Press.

Farkas, Katalin (2003). What is Externalism? *Philosophical Studies* 112 (3): 187-208.

——— (2008a). *The Subject's Point of View*. Oxford: Oxford University Press.

——— (2008b). Phenomenal Intentionality Without Compromise. *The Monist* 91 (2): 273-293.

Feit, Neil (2008). *Belief About the Self: A Defense of the Property Theory of Content*. Oxford: Oxford University Press.

Fine, Kit (1977). Prior on the Construction of Possible Worlds and Instants. Postscript to Prior and Fine, *Worlds, Times, and Selves*. Massachusetts: Duckworth. 116-161.

——— (1985). Plantinga on the Reduction of Possibilist Discourse. In Tomberlin, James & van Inwagen, Peter (eds.), *Alvin Plantinga (Profiles, Vol. 5)*. Dordrecht: D. Reidel Publishing Company. 145-186.

———— (1982). The Problem of Non-Existents. I. Internalism. *Topoi* 1. 97-140.

———— (2007). *Semantic Relationism*. Oxford: Blackwell Publishing.

Fitch, Greg, and Nelson, Michael (2014). Singular Propositions. *The Stanford Encyclopedia of Philosophy* (Fall 2014 Edition), Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/fall2014/entries/propositions-singular/>.

Fodor, Jerry A. (1987). *Psychosemantics: The Problem of Meaning in the Philosophy of Mind*. Cambridge, Mass.: MIT Press.

———— (1991). A Modal Argument for Narrow Content. *Journal of Philosophy* 88 (1): 5-26.

Frege, Gottlob (1892/1948). Sense and Reference. *Philosophical Review* 57 (3): 209-230.

———— (1956). The Thought: A Logical Inquiry. *Mind* 65 (259): 289-311.

Gertler, Brie (2012). Understanding the Internalism-Externalism Debate: What is the Boundary of the Thinker? *Philosophical Perspectives* 26 (1): 51-75.

Grice, H. P. (1957). Meaning. *Philosophical Review* 66 (3): 377-388.

Hall, Ned (2004). Two Concepts of Causation. In John Collins, Ned Hall & Laurie Paul (eds.), *Causation and Counterfactuals*. Cambridge, Mass.: MIT Press. 225-276.

Hanks, Peter W. (2013). First-Person Propositions. *Philosophy and Phenomenological Research* 86 (1): 155-182.

———— (2015). *Propositional Content*. Oxford: Oxford University Press.

Harman, Gilbert (1990). The Intrinsic Quality of Experience. *Philosophical Perspectives* 4: 31-52.

Hawthorne, John, & Manley, David (2012). *The Reference Book*. Oxford: Oxford University Press.

Horgan, Terence (2013). Original Intentionality is Phenomenal Intentionality. *The Monist* 96 (2): 232-251.

Jackson, Frank (1977). *Perception: A Representative Theory*. Cambridge: Cambridge University Press.

———— (1998). *From Metaphysics to Ethics: A Defence of Conceptual Analysis*. Oxford: Oxford University Press.

———— (2001). Responses. In panel discussion of: *From Metaphysics to Ethics: A Defence of Conceptual Analysis*. *Philosophical and Phenomenological Research* 62 (3): 653-664.

——— (2003a). Narrow Content and Representation—or Twin Earth Revisited. *Proceedings and Addresses of the American Philosophical Association* 77 (2): 55-70.

——— (2003b). Representation and Narrow Belief. *Philosophical Issues* 13 (1): 99-112.

——— (2004). Why We Need A-intensions. *Philosophical Studies* 118 (1-2): 257-277.

——— (2006). The Knowledge Argument, Diaphanousness, Representationalism. In Torin Alter & Sven Walter (eds.), *Phenomenal Concepts and Phenomenal Knowledge: New Essays on Consciousness and Physicalism*. Oxford: Oxford University Press. 52-64.

Jubien, Michael (2001). Propositions and the Objects of Thought. *Philosophical Studies* 104 (1): 47-62.

Kallestrup, Jesper (2011). *Semantic Externalism*. London and New York: Routledge.

Kaplan, David (1970). What is Russell's Theory of Descriptions? In Wolfgang Yourgrau & Allen D. Breck (eds.), *Physics, Logic, and History*. New York: Plenum Press. 277-295.

——— (1989). Demonstratives. In Joseph Almog, John Perry & Howard Wettstein (eds.), *Themes from Kaplan*. Oxford: Oxford University Press. 481-563.

King, Jeffrey C. (2007). *The Nature and Structure of Content*. Oxford: Oxford University Press.

Kölbel, Max (2002). *Truth Without Objectivity*. London and New York: Routledge.

——— (2008). Introduction: "Motivations for Relativism". In Manuel García-Carpintero & Max Kölbel (eds.), *Relative Truth*. Oxford: Oxford University Press. 1-40.

——— (2009). The Evidence for Relativism. *Synthese* 166 (2): 375-395.

Kriegel, Uriah (2007). Intentional Inexistence and Phenomenal Intentionality. *Philosophical Perspectives* 21 (1): 307-340.

——— (2011). *The Sources of Intentionality*. Oxford: Oxford University Press.

Kripke, Saul A. (1963). Semantical Considerations on Modal Logic. *Acta Philosophica Fennica* 16 (1963): 83-94.

——— (1979). A Puzzle about Belief. In A. Margalit (ed.), *Meaning and Use*. Dordrecht: D. Reidel publishing company. 239-283.

——— (1980). *Naming and Necessity*. Cambridge, Mass: Harvard University Press.

Lewis, David (1973). Causation. *Journal of Philosophy* 70: 556-567.

———— (1979). Attitudes *De Dicto* and *De Se*. *Philosophical Review* 88 (4): 513-543.

———— (1986). *On the Plurality of Worlds*. Oxford: Blackwell Publishing.

———— (2004). Causation as Influence. In John Collins, Ned Hall and L. A. Paul. O (eds.), *Causation and Counterfactuals*. Cambridge, Mass.: MIT Press. 75–106.

Linsky, Bernard & Zalta, Edward N. (1996). In Defense of the Contingently Nonconcrete. *Philosophical Studies* 84 (2-3): 283-294.

———— (1994). In Defense of the Simplest Quantified Modal Logic. *Philosophical Perspectives* 8 (Logic and Language): 431-458.

Loar, Brian (1988). Social Content and Psychological Content. In Robert H. Grimm & D. D. Merrill (eds.), *Contents of Thought*. Tucson: University of Arizona Press.

———— (2003). Phenomenal Intentionality as the Basis of Mental Content. In Martin Hahn & B. Ramberg (eds.), *Reflections and Replies: Essays on the Philosophy of Tyler Burge*. Cambridge, Mass.: MIT Press. 229-258.

Lowe, E. J. (1998). *The Possibility of Metaphysics: Substance, Identity, and Time*. Oxford: Clarendon Press.

Lynch, Michael P. (2001). *Truth in Context: An Essay on Pluralism and Objectivity*. Cambridge, Mass.: MIT Press.

Marcus, Ruth Barcan (1986). Possibilia and Possible Worlds. *Grazer Philosophische Studien* 25/26 (1985/1986): 107-133.

MacFarlane, John (2003). Future Contingents and Relative Truth. *Philosophical Quarterly* 53 (212): 321-336.

———— (2005). The Assessment Sensitivity of Knowledge Attributions. In Tamar Szabo Gendler John Hawthorne (ed.), *Oxford Studies in Epistemology*. Oxford: Oxford University Press. 197-234.

McCulloch, Gregory (1992). The Spirit of Twin Earth. *Analysis* 52 (3): 168-174.

McDowell, John (1984). *De Re* Senses. *Philosophical Quarterly* 34 (136): 283-294.

McGinn, Colin (1977). Charity, Interpretation, and Belief. *Journal of Philosophy* 74 (9): 521-535.

———— (1982). The Structure of Content. In Andrew Woodfield (ed.), *Thought and Object*. Oxford: Oxford University Press. 207-258.

———— (1983). *The Subjective View: Secondary Qualities And Indexical Thoughts*. Oxford: Clarendon Press.

McGrath, Matthew (2014). Propositions. *The Stanford Encyclopedia of Philosophy* (Spring 2014 Edition), Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/spr2014/entries/propositions/>.

Meinong, Alexius (1904/1960). The Theory of Objects (translation of `Über Gegenstandstheorie', 1904). In Roderick Chisholm (ed.), *Realism and the Background of Phenomenology*. Glencoe, Illinois: Free Press. 76-117.

Menzel, Christopher (1991). The True Modal Logic. *Journal of Philosophical Logic* 20 (4): 331 - 374.

———— (2016). Actualism. *The Stanford Encyclopedia of Philosophy* (Summer 2016 Edition), Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/sum2016/entries/actualism/>.

Menzies, Peter (2014). Counterfactual Theories of Causation, *The Stanford Encyclopedia of Philosophy* (Spring 2014 Edition), Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/spr2014/entries/causation-counterfactual/>.

Merricks, Trenton (2001). *Objects and Persons*. Oxford University Press.

———— (2012). Singular Propositions. In Kelly James Clark and Michael C. Rea (ed.), *Reasons, Metaphysics, and Mind: New Essays on the Philosophy of Alvin Plantinga*. Oxford: Oxford University Press. 61-81.

———— (2015). *Propositions*. Oxford University Press: Oxford.

Moore, Joseph G. (1999). Propositions, Numbers, and the Problem of Arbitrary Identification. *Synthese* 120 (2): 229-263.

Paul, L. A. & Hall, Ned (2013). *Causation: A User's Guide*. Oxford: Oxford University Press.

Perry, John (1977). Frege on Demonstratives. *Philosophical Review* 86 (4): 474-497.

———— (1979). The Problem of the Essential Indexical. *Noûs* 13 (December): 3-21.

Plantinga, Alvin (1974). *The Nature of Necessity*. Oxford: Oxford University Press.

———— (1979). De Essentia. *Grazer Philosophische Studien* 7: 101-121.

———— (1983). On Existentialism. *Philosophical Studies* 44 (1): 1-20.

———— (1985). Replies. In Tomberlin, James & van Inwagen, Peter (eds.), *Alvin Plantinga (Profiles, Vol. 5)*. Dordrecht: D. Reidel Publishing Company. 313-398.

Pollock, John L. (1985). Plantinga on Possible Worlds. In Tomberlin, James & van Inwagen, Peter (eds.), *Alvin Plantinga (Profiles, Vol. 5)*. Dordrecht: D. Reidel Publishing Company. 121-144.

Priest, Graham (2005). *Towards Non-Being: The Logic and Metaphysics of Intentionality*. Oxford: Oxford University Press.

Prior, A. N. (1957). *Time and Modality*. Oxford: Oxford University Press.

——— (1967). *Past, Present and Future*. Oxford: Clarendon Press.

Putnam, Hilary (1973). Meaning and Reference. *Journal of Philosophy* 70 (19): 699-711.

——— (1975). The Meaning of 'Meaning'. *Minnesota Studies in the Philosophy of Science* 7: 131-193.

Quine, W. v. O. (1961). On What There Is. In *From a Logical Point of View*. Cambridge, Mass.: Harvard University Press. 1-19.

——— (1969). Propositional Objects. In *Critica*. Columbia University Press. 139-160.

Rowe, William L. (2007). *Philosophy of Religion: An Introduction*. 4th Edition. Belmont, CA: Wadsworth/Thomson Learning.

Recanati, François (1993). *Direct Reference: From Language to Thought*. Oxford: Blackwell.

Russell, Bertrand (1903). *Principles of Mathematics*. Cambridge: Cambridge University Press.

——— (1905). On Denoting. *Mind* 14 (56): 479-493.

Salmon, Nathan U. (1986). *Frege's Puzzle*. Ridgeview.

——— (1987). Existence. *Philosophical Perspectives* 1: 49-108.

——— (1998). Nonexistence. *Noûs* 32 (3): 277-319.

Schaffer, Jonathan (2016). The Metaphysics of Causation. *The Stanford Encyclopedia of Philosophy* (Fall 2016 Edition), Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/fall2016/entries/causation-metaphysics/>.

Schiffer, Stephen R. (2003). *The Things We Mean*. Oxford: Oxford University Press.

Schnieder, Benjamin Sebastian (2005). A Certain Kind of Trinity: Dependence, Substance, Explanation. *Philosophical Studies* 129 (2): 393-419.

Schroeter, Laura. Two-Dimensional Semantics. *The Stanford Encyclopedia of Philosophy* (Winter 2012 Edition), Edward N. Zalta (ed.), URL = < http://plato.stanford.edu/archives/win2012/entries/two-dimensional-semantics/ >.

Searle, John R. (1983). *Intentionality: An Essay in the Philosophy of Mind*. Cambridge: Cambridge University Press.

Soames, Scott (1987). Direct Reference, Propositional Attitudes, and Semantic Content. *Philosophical Topics* 15 (1): 47-87.

———— (2008). Truthmakers? *Philosophical Books* 49 (4): 317-327.

———— (2015). *Rethinking Language, Mind, and Meaning*. Princeton, NJ: Princeton University Press.

Stalnaker, Robert (1976). Propositions. In Alfred F. MacKay and Daniel D. Merrill (eds.) *Issues in the Philosophy of Language*. New Haven, CN: Yale University Press.

———— (1978). Assertion. *Syntax and Semantics* 9: 315-332.

———— (1981). Indexical belief. *Synthese* 49 (1): 129-151.

———— (1984). *Inquiry*. Cambridge: Cambridge University Press.

———— (1989). On What's in the Head. *Philosophical Perspectives* 3: 287-319.

———— (1990/1999). *Context and Content: Essays on Intentionality in Speech and Thought*. Oxford: Oxford University Press.

———— (2003). *Ways a World Might Be: Metaphysical and Anti-Metaphysical Essays*. Oxford: Oxford University Press.

———— (2011). *Mere Possibilities: Metaphysical Foundations of Modal Semantics*. Princeton, NJ: Princeton University Press.

Sullivan, Meghan (2012). The Minimal A-theory. *Philosophical Studies* 158 (2): 149-174.

———— (2014). Modal Logic as Methodology. *Philosophy and Phenomenological Research* 88 (3): 734-743.

Sylvan (*né* Routley), Richard (1980). *Exploring Meinong's Jungle and Beyond: An Investigation of Noneism and the Theory of Items*. Research School of Social Sciences, Australian National University.

Tye, Michael (2000). *Consciousness, Color, and Content*. Cambridge, Mass.: MIT Press.

——— (2007). Intentionalism and the Argument from No Common Content. *Philosophical Perspectives* 21 (1): 589-613.

——— (2009). *Consciousness Revisited: Materialism Without Phenomenal Concepts*. Cambridge, Mass.: MIT Press.

van Inwagen, Peter (1990). *Material Beings*. Ithaca, NY: Cornell University Press.

Weatherson, Brian and Marshall, Dan (2014). "Intrinsic vs. Extrinsic Properties", *The Stanford Encyclopedia of Philosophy* (Fall 2014 Edition), Edward N. Zalta (ed.), URL = <https://plato.stanford.edu/archives/fall2014/entries/intrinsic-extrinsic/>.

Williamson, Timothy (1998). Bare Possibilia. *Erkenntnis* 48 (2/3): 257-273.

——— (2002). Necessary Existents. In A. O'Hear (ed.), *Royal Institute of Philosophy Supplement*. Cambridge: Cambridge University Press. 269-287.

——— (2013). *Modal Logic as Metaphysics*. Oxford: Oxford University Press.

Zalta, Edward N. (1983). *Abstract Objects: An Introduction to Axiomatic Metaphysics*. Dordrecht: D. Reidel Publishing Company.

——— (1988). *Intensional Logic and the Metaphysics of Intentionality*. Cambridge, Mass.: MIT Press.

——— (2006). Essence and Modality. *Mind* 115 (459): 659-693.