

Apriority, Super-Rigidity, and Fregean Content

Robert Smithson

A thesis submitted to the faculty of the University of North Carolina at Chapel Hill in partial fulfillment of the requirements for the degree of Master of Arts in the Department of Philosophy.

Chapel Hill
2012

Approved by:

William Lycan

Keith Simmons

John Roberts

ABSTRACT

ROBERT SMITHSON: Apriority, Super-Rigidity, and Fregean Content
(Under the direction of William Lycan)

What is the difference between an utterance of the sentence ‘Hesperus is Hesperus’ and an utterance of the sentence ‘Hesperus is Phosphorus’? David Chalmers claims that these sentences are psychologically and epistemologically distinct because they are associated with distinct fine-grained Fregean propositions. In “The Nature of Epistemic Space” (2011a), Chalmers offers a rigorous analysis of epistemic possibility that he claims can be used to ground an account of Fregean content that is available to a wide variety of theorists. In this thesis, I will identify the assumptions required by Chalmers’ construction of epistemic space in order to evaluate whether it is indeed capable of grounding a widely acceptable account of Fregean content. I conclude the thesis by comparing Chalmers’ Fregean content to alternative Russellian accounts of cognitive significance.

ACKNOWLEDGMENTS

I want to thank my advisor Bill Lycan for his guidance throughout this project. I also want to thank my readers Keith Simmons and John Roberts for their valuable contributions to this thesis.

TABLE OF CONTENTS

LIST OF FIGURES.....	v
Section	
1. Overview: Epistemic possibility and Fregean content.....	1
2. Russellian guises.....	15
3. Apriority roles.....	27
4. Super-rigidity.....	37
5. The Fregean and the Russellian.....	48
Conclusion.....	54
REFERENCES.....	55

LIST OF FIGURES

Figure

- 1. A sample apriority web.....28
- 2. A symmetric apriority web.....33

1 Overview: Epistemic possibility and Fregean content

There is a vast literature on metaphysical possibility. The dominant contemporary approach to understanding metaphysical possibility is to think of possibility and necessity in terms of possible worlds. There are certainly disagreements about the nature of possible worlds (are they concrete spatiotemporally isolated universes? are they ersatz representations? etc.), but the different competing accounts of possible worlds are at least relatively well-understood.

But there is a second type of possibility that is not nearly as well-understood. Let's say that it is *epistemically* possible for a subject that p when it might be that p for all a subject knows. Here are some examples. For all I know, it might be that the stock market goes up tomorrow, and it might be that it goes down. It might be that the woman sitting at the bus stop is a skilled musician, or it might be that she is not. There are even more ways the world might be, for all I know with certainty. It might be that I am a brain in a vat, and it might be that I am not. It might be that water is H_2O , and it might be that water is XYZ. It might be that Hesperus is identical to Phosphorus, or it might be that Hesperus is not identical to Phosphorus.

From the last two examples in particular, it should be clear that epistemic possibility cannot be analyzed using the traditional framework of metaphysically possible worlds. After all, there are no possible worlds where water is XYZ, yet in

at least one sense 'water is XYZ' is epistemically possible. As an alternative, Chalmers introduces the notion of a scenario, which is a kind of epistemic analogue to a metaphysically possible world. Intuitively, a scenario is a maximally specific way things might turn out to be. Chalmers' hope is that, just as we use the space of possible worlds to analyze metaphysical possibility, we can use the space of scenarios to analyze epistemic possibility.

But giving a rigorous account of the space of scenarios is important for a second, independent reason: Chalmers hopes to use epistemic space to ground an account of Fregean propositions that reflect a sentence/thought's cognitive significance for its speaker/subject. Without discussing the finer details of Chalmers' account of Fregean content (see footnote 1), suffice to say that one of the crucial semantic values Chalmers assigns to a sentence/thought is its primary ("epistemic") intension.¹ Primary intensions are functions from scenarios to extension. So if Chalmers can give a widely acceptable account of epistemic space, it seems that he will also be able to give a widely acceptable account of Fregean content. The general aim of this paper is to identify exactly what assumptions are required by Chalmers' account. To this end, I will use section 1 to provide an overview of the essentials of Chalmers' preferred account of

¹ Chalmers (2011b) defines the primary intension for a linguistic expression as a function from scenarios to extension. He defines an enriched intension as the ordered pair of an expression's primary intension and its extension. Chalmers also defines a structured primary intension for a complex logical expression to be a structure consisting of the primary intensions of all the simple expressions in a sentence, structured according to the sentence's logical form. He then is able to define the enriched proposition of a sentence to be a structure consisting of the enriched intension of the sentence's logical parts, structured according to the sentence's logical form. Chalmers claims that *enriched propositions* and *enriched intensions* behave very similarly to Fregean thoughts and Fregean senses, respectively. So strictly speaking, a sentence's full cognitive significance is reflected in its enriched proposition, not its primary intension. But since these semantic values are closely related and since the differences between them are not usually relevant to the arguments in this thesis, I will use the expression "Fregean content" to refer to each of them interchangeably.

epistemic space, which he calls the *epistemic construction*.² I will conclude section 1 with a list of the assumptions that seem to be required by Chalmers' account.

1.1 Epistemic possibility

Chalmers claims that it is epistemically possible (for a subject) that p iff there exists a *scenario* (for that subject) that *verifies* p (2011a: 64) This definition presents at least four immediate questions. What is epistemic possibility? What type of object is p ? What is a scenario? What is the verification relation? I will consider these questions in the next several sections.

What exactly is epistemic possibility? Chalmers is a *pluralist* in his understanding of epistemic possibility: he allows that there are many notions of epistemic possibility which may be more or less useful for different projects. In other words, there isn't a single "correct" account of what sentences are epistemically possible; what counts as epistemically possible depends on one's present philosophical purposes. For example, we might distinguish between subject-relative and subject-independent epistemic possibility. We can define *strict epistemic possibility* as ways things might be, for all some particular subject knows. We could distinguish this from *deep epistemic possibility* as ways things might be, for all *any* subject knows (that is: p is deeply epistemically possible when there is *no* subject who knows that $\sim p$). Strict epistemic possibility will be more useful if we are interested in modeling the doxastic state of a single

² Chalmers also offers an alternative *metaphysical construction*, according to which scenarios are centered possible worlds. Because some of the assumptions required by the metaphysical construction are controversial, Chalmers prefers the epistemic construction and thinks that it is available to a wider variety of theorists.

individual. But deep epistemic possibility is more important if we are interested in grounding an account of subject-independent Fregean content. Throughout this thesis, I follow Chalmers in focusing on deep epistemic possibility.

Even when restricting our attention to deep epistemic possibility, we can define a spectrum of standards for what counts as epistemically possible. On the maximally liberal end of the spectrum, we might have a standard of epistemic possibility according to which even logical contradictions are epistemically possible. Perhaps this standard would be useful for making sense of the epistemic states of extremely non-ideal thinkers. But it certainly won't be help to ground an account of Fregean content.

On the other end of the spectrum, we can try to define a type of idealized epistemic possibility. On this idealized account, p is deeply epistemically necessary when p is a priori under idealization. The idealization employed in this account abstracts away from contingent cognitive limitations: if a hypothesis can be excluded only by a great amount of a priori reasoning, it is nonetheless ruled out a priori. For example, if Goldbach's Conjecture is true, then it is a priori; if it is false, then its negation is a priori.³ Between these two ends of the spectrum, there are a range of standards of epistemic possibility that one might define, perhaps in order to model the epistemic states of non-ideal rational subjects. But

³ Chalmers (forthcoming) offers many details on what is required by the idealization employed in his account. The idealization allows subjects to possess any concept that it is possible to possess, regardless of whether they actually possess it. The idealization abstracts away from limitations of complexity, allowing subjects to entertain thoughts whose complexity greatly exceeds ordinary human capacity. There is also an idealization in calculation: arbitrary numbers of steps of reasoning are allowed, so that (for example) enormously complex mathematical theorems come out a priori. The idealization also abstracts away from missteps in reasoning. Chalmers (forthcoming) devotes considerable time defending various ways to cash out the notion of idealized apriority, but assessing these arguments is outside the scope of this paper.

Chalmers is most interested in ideal deep epistemic space, and this is the only type of space I will consider in this thesis.

1.2 The objects of epistemic possibility

The objects of epistemic/metaphysical possibility are most commonly understood to be propositions, where propositions are some sort of language-independent entity. But the nature of propositions is contested. What type of propositions could serve as the objects of epistemic possibility? On a popular Russellian view, the proposition expressed by a sentence is a structure involving the objects/properties that are the extensions of the simple logical parts of the sentence. But this account of propositions encounters immediate difficulties. For example, we want an account of epistemic space on which it is epistemically possible that 'Hesperus \neq Phosphorus' even though it is not epistemically possible that 'Hesperus \neq Hesperus'. But on a Russellian view, 'H \neq P' and 'H \neq H' express the same proposition. So it is unclear how Russellian propositions could serve as the objects of epistemic possibility.⁴

On a Fregean view, the proposition expressed by a sentence is a structure of senses expressed by the simple logical parts of the sentence, where senses are fine-grained entities reflecting cognitive significance. So Fregean propositions seem like better candidates for serving as the objects of epistemic possibility. But Chalmers observes that there are two major problems with using Fregean

⁴ The same problem arises on a possible-worlds account of propositions, according to which the proposition expressed by a sentence is the set of possible worlds where the sentence is true. But for ease of presentation, I will frame the issues in this thesis as a dialectic involving the Russellian and the Fregean. This is mainly because Russellian theorists have developed a variety of accounts to explain the cognitive difference between 'H=H' and 'H=P', and one of the goals of this thesis is to compare these accounts of "propositional guises" to Chalmers' account of Fregean propositions.

propositions in his account of epistemic space. First, the existence of Fregean propositions is heavily disputed. It doesn't seem likely that many theorists will accept an account of epistemic space that takes Fregean propositions as primitive. But Chalmers wants his account of epistemic space to be available to a wide variety of theorists. Second, I've mentioned that Chalmers hopes to use epistemic space to ground an account of Fregean propositions. So he can't presuppose Fregean propositions without facing a circularity objection. So instead, Chalmers considers *sentences* to be the objects of epistemic possibility.⁵

Recall that p is deeply epistemically necessary when p is a priori under idealization. So if sentences are the objects of possibility, the notion of apriority must be something that can apply to sentences. But the notions of sentence knowledge and sentence apriority are unintuitive and unfamiliar. What does it mean to know a sentence a priori, if this doesn't just mean knowing the proposition associated with the sentence a priori?

Intuitively, there is a difference between the occurrent mental state of someone who sincerely asserts 'H=H' and someone who sincerely asserts 'H=P'. Even theorists who assign 'H=H' and 'H=P' the same proposition usually acknowledge *some* difference between the mental states associated with these utterances. For this reason, Chalmers (forthcoming) prefers to cash out the

⁵ Both sentence *tokens* and *sentence* types can serve as objects of epistemic possibility. A sentence token is a specific orthographic item uttered by a subject at a specific place and time. A sentence type is some kind of abstract structured entity composed out of simpler expression types, such as words types. Chalmers assumes that sentence types belong to their language essentially so that the same expression type cannot have tokens in different languages. So the sentence types that Chalmers has in mind are not purely orthographic types.

notion of sentence apriority directly in terms of mental states.⁶ Let's say that *entertaining* is a maximally general propositional attitude (occurrent or non-occurrent) with a mind-to-world direction of fit. We can then say that a *thought* is a specific state of entertaining. For example, whenever there is a belief or a knowing, there is a thought; in these cases, we can say that such a thought *constitutes* the corresponding belief or item of knowledge.⁷ Chalmers also assumes that thoughts are the kind of things that can stand in relations of negation, conjunction, and disjunction to each other. We next introduce the idea that utterances of truth-apt sentences typically *express* (occurrent) thoughts. For example, sincere utterances express beliefs, and thus express thoughts. In fact, even insincere utterances typically express thoughts.⁸

With these materials, Chalmers' defines a notion of *knowing a sentence*: S knows a sentence token *t* when S's utterance of *t* expresses a thought that constitutes an item of knowledge. To introduce apriority to the picture, Chalmers begins with an account of apriority for thoughts: a thought T is *a priori* when T constitutes an item of knowledge that is justified independently of experience. We

⁶ See Excurses 3 of *Constructing the World*. Chalmers also mentions that sentence apriority is to be cashed out in terms of token mental states in *The Nature of Epistemic Space*, but the treatment is much more detailed in *CtW*. So my presentation of the notion of sentence knowledge will draw mainly from *CtW*.

⁷ Is the thought state identical to the belief state or the knowledge state? On some views, these states are all identical to one another. On other views, these states are distinct, but there is a close relation between them. Chalmers wants to remain neutral on this question.

⁸ It is important to remember that, on the current usage, expression is a relation between utterances and mental states, *not* a relation between utterances and propositions. But intuitively, says Chalmers, an utterance and the thought it expresses share the same propositional content. Chalmers does not build sameness of propositional content into the notion of expression, because there may be views where thoughts and utterances have different types of content. But at the very least, an utterance should have the same truth conditions as the thought it expresses.

can then say that S *knows a sentence token t a priori* when S's utterance of *t* expresses an a priori thought.⁹

One might worry that, in moving to sentences, Chalmers hasn't escaped the original difficulties with propositions. After all, Chalmers cashes out sentence knowledge in terms of thoughts with propositional content. But importantly, Chalmers' analysis of a thought is *neutral* about what type of propositional content is associated with thoughts. What's important for the epistemic construction is that it be possible for a subject to know 'H=H' but not know 'H=P'. In order to allow for this possibility, the epistemic construction only requires that a theorist accept *some way or other* to distinguish the thoughts associated with these utterances. To be sure, since theorists disagree about the nature of thoughts, theorists will also disagree about what exactly is required for a subject to know 'H=P'. The Russellian may claim that S knows 'H=P' when S's utterance of 'H=P' expresses a thought with a certain proposition presented under a certain *guise*. In contrast, Chalmers himself will ultimately cash out sentence knowledge in terms of fine-grained Fregean propositions. But either one of these analyses will allow a theorist to assign different modal properties to 'H=H' and 'H=P'. Since Chalmers doesn't build Fregean content into the notion of a thought, he can get the notion of sentence knowledge off the ground while avoiding a charge of circularity.

⁹ What about a notion of sentence knowledge for sentence types? For ease of presentation, I will wait until section 2.1 to present the notion of sentence type knowledge.

1.3 Scenarios, the verification relation, and plenitude

When a sentence t is true at a possible world u , we can say that u *satisfies* t . I've already intuitively characterized scenarios as the epistemic analogs of possible worlds. Similarly, we can intuitively characterize the verification relation as the epistemic analog of the satisfaction relation: a scenario w *verifies* a sentence s when s is true at w . While these characterizations are useful, one shouldn't put *too* much weight on the analogy between epistemic and metaphysical possibility. Scenarios are analogous to possible worlds, but they may be very different from possible worlds. Similarly, while we may grant that sentences have something *analogous* to truth conditions across the space of scenarios, it would be too quick at this point to speak as if sentences *literally* have a second set of (epistemic) truth conditions.

Indeed, it is important to remember that there are many standards of epistemic possibility; as a result, we should not expect that there is one single "correct" account of scenarios/verification. Instead, our aim is to find some account of scenarios/verification that is appropriate for the type of epistemic possibility we are modeling. For example, since 'H≠P' is not ruled out a priori, it is a constraint on the adequacy of any proposed account of ideal epistemic space that it include a scenario that verifies 'H≠P'. More generally, Chalmers offers the following principle as a constraint on any adequate theory of epistemic possibility:

Plenitude: s is epistemically possible iff there is a scenario w such that w verifies s

I'll give an example to make this discussion concrete. Suppose we are interested in the maximally liberal account of deep epistemic possibility according to which every sentence (including contradictions) is epistemically possible. Chalmers suggests that an adequate space of scenarios for this type of epistemic possibility might be the power set of the set of sentences, such that every set of possible sentences corresponds to a scenario (2011a: 65). He doesn't suggest a corresponding verification relation, but a natural suggestion is the set membership relation: a scenario w verifies a sentence s when s is a member of w . It seems like this construction will satisfy Plenitude.¹⁰

1.4 Compositionality

Obviously, the example construction just given is inadequate for modeling ideal epistemic space. This is because it verifies sentences that can be ruled out a priori (ie, '2+2=0', 'John is a bachelor and John is a married', etc.), thus violating Plenitude. A natural suggestion is to restrict the sets of sentences that count as scenarios to the sets that only include sentences that cannot be ruled out a priori (we can continue to interpret the verification relation as a set membership relation).

With this adjustment, Plenitude is satisfied. But because Chalmers wants to use scenarios to ground an account of Fregean content, it turns out that an additional constraint is needed. I briefly mentioned that one of the key semantic

¹⁰ Incidentally, this construction shows why the analogy between scenarios/verification and possible worlds/satisfaction should not be taken too seriously. I intuitively characterized the verification relation as follows: a scenario w verifies a sentence s when s is true at w . But it is unnatural to speak of sentences (including contradictions) being *true* (in some sense) just because they happen to be members of certain arbitrary sets. Similarly, one can see that the power set of the set of sentences is nothing like the space of possible worlds.

values that Chalmers associates with a sentence is its primary intension. Having now encountered the verification relation, I can say more precisely that a sentence's primary intension is a function from scenarios to the truth value of $ver(w,s)$.¹¹ To fill the role that Chalmers intends for them, primary intensions must satisfy certain principles of compositionality. So, for example, if a given scenario verifies s and verifies t , it should also verify the conjunction $s\&t$. But it is easy to see that the current suggestion will not allow for compositionality. Consider the scenario $w = \{s = \text{'Water is H}_2\text{O'}, t = \text{'Lead is a compound'}\}$. w verifies s and verifies t (given that s is a member of w and t is a member of w), but w doesn't verify $s\&t$ (given that w only contains two atomic sentences). This example motivates adopting the following additional constraint of adequacy for an account of epistemic space:

Compositionality: When a complex sentence s is composed from simpler sentences s_i and truth-functional connectives, $ver(w,s)$ is determined by $ver(w,s_i)$ in the corresponding truth-functional way. For example, $ver(w,\sim s)$ iff $\sim ver(w,s)$, and $ver(w,s\&t)$ iff $ver(w,s)\&ver(w,t)$.

1.5 The epistemic construction

In this section, I'll present the specifics of Chalmers' account of ideal epistemic space: the *epistemic construction*. Chalmers begins by using the notion of apriority to define three derivative notions (2011a: 67,76). (1) A sentence type d *implies* a sentence token/type s when the sentence $\sim dvs$ is a priori. (2) A sentence d is *epistemically complete* iff for all sentences s , d implies

¹¹ Chalmers actually says that a sentence's epistemic intension is a function from scenarios to $ver(w,s)$ (2011a: 65). But it is clear that Chalmers intends for epistemic intensions to be functions from scenarios to truth values.

s or d implies $\sim s$. (3) Two sentences s and t are members of an equivalence class when s implies t and t implies s .

To help convey the intuitive idea behind epistemic completeness, I'll give an example of how we might go about constructing an epistemically complete sentence. Suppose we start with the sentence 'John is a bachelor'. 'John is a bachelor' will imply many sentences, such as 'John is a bachelor or John is tall', 'John is male', 'John exists', etc. It will also rule out many sentences a priori, such as the sentences 'John is married', 'It is not the case that John exists', etc. But many sentences are neither implied nor ruled out, such as 'Water is identical to H₂O', 'Mozart composed 41 symphonies', 'There are 12 oceans on planet earth', etc. Since it is still an open question whether 'Water is identical to H₂O', etc. are true when given 'John is a bachelor', 'John is a bachelor' is not epistemically complete. So we have to adjoin more sentences to 'John is a bachelor' until we have a sentence that "settles everything". Obviously, many epistemically complete sentences will be quite long (in fact, Chalmers' account requires infinitary sentences).¹² For example, any epistemically complete sentence corresponding to the actual world will imply all physical truths of the actual world, all truths about the conscious experience of conscious subjects of

¹² One might wonder: isn't it always possible to add further sentences that aren't ruled out a priori? In one sense, the answer is yes. 'John is a bachelor' implies an infinite number of sentences (including its infinite logical consequences). So, for example, one can always adjoin additional logical consequences of 'John is a bachelor' to a sentence that includes 'John is a bachelor'. But when we do so, we don't move any step closer to building an epistemically complete sentence: all of these logical consequences were already implied by the original sentence. One might think that there is a stronger sense in which it is possible to add further sentences that aren't ruled out a priori. Can't we, for example, always add the description of some additional hydrogen atom to our growing sentence? While Chalmers allows that some epistemically complete sentences are infinitary, they need not be. This is because we can add a totality sentence that, intuitively, states that the world is no bigger than is necessary to accommodate the truth of all of the other sentences in the conjunction. Chalmers (forthcoming) discusses various ways to formalize totality clauses in (ch. 3).

the actual world, all truths about Joan of Arc, all truths about the stock market, all truths about bicycles, etc.

Intuitively, an epistemically complete sentence is a long conjunction that represents a maximally specific way the world might be, for all we know a priori. So on Chalmers account, scenarios are defined to be equivalence classes of epistemically complete sentences.¹³ It is then natural to say that a scenario w verifies a sentence s when d implies s , where d is some member of the equivalence class of w .

It is easy to see the advantages of the epistemic construction. The construction satisfies compositionality because the set-membership verification relation has been replaced by the implication relation, which is governed by an analogous principle of compositionality. For example, if a scenario w verifies s and t , then w will also verify their conjunction $s \& t$ (since any scenario that implies s and t will also imply their conjunction $s \& t$). Chalmers gives various arguments for why the epistemic construction also satisfies Plenitude. Another attractive feature of the epistemic construction is that each scenario corresponds to a *maximally specific* epistemic possibility (for example, equivalence classes won't include sets like $\{s = \text{'Water is H}_2\text{O'}$, $t = \text{'Lead is a compound'}\}$). At the very least, the inclusion of less specific scenarios seems redundant and less parsimonious.

¹³ Why equivalence classes? Intuitively, a single maximally specific epistemic possibility can be described by many different epistemically complete sentences. For example, rearranging the conjuncts in an epistemically complete sentence will yield another epistemically complete sentence that describes the same possibility. So it is more parsimonious to group together sentences that imply one another.

1.7 The materials required for the epistemic construction

In the last several sections, I have presented a basic outline of Chalmers' account of epistemic space. We are now in a position to clearly state what seems to be required by the epistemic construction. To accept the epistemic construction, a theorist must first accept:

(1) Some account or other of the psychological/epistemological difference between the thoughts $H=H$ and $H=P$. Chalmers wants to assign different modal properties to sentences like 'H=H' and 'H=P' without presupposing Fregean content. So he cashes out the notion of sentence knowledge in terms of token mental states, appealing to the intuitive idea that 'H=H' and 'H=P' are associated with cognitively distinct thoughts. Importantly, the epistemic construction only requires *some* account of the difference between $H=H$ and $H=P$; it doesn't specifically require Fregean content. Of course, some theorists (e.g. Soames (2005)) *deny* a difference between $H=H$ and $H=P$. But, says Chalmers, if a theorist can't distinguish the mental states associated with sincere utterances of 'H=H' and 'H=P', then so much the worse for her theory. Since most theorists share this intuition, it seems like a wide variety of theorists can accept (1).

(2) A notion of *idealized apriority*. (2) is much more controversial than (1). Some theorists will be troubled by the notion of apriority itself, while others will be troubled by the idealization employed in the notion of apriority. For what it is worth, Chalmers (forthcoming) extensively defends both the general notion of apriority (ch. 5) and the idealization (ch. 2) required for the account. While I will generally take the notion of idealized apriority for granted throughout this paper,

suffice to say that anyone who rejects this notion will also reject the epistemic construction.¹⁴

Prima facie, it seems like any theorist who accepts (1) and (2) can accept Chalmers' account of epistemic space.¹⁵ And because primary intensions are nothing more than functions defined across the space of scenarios, it seems that anyone who accepts the epistemic construction is also in a position to accept Chalmers' Fregean content. Is it really true that any theorist who accepts (1) and (2) can also accept Fregean content? Of these two requirements, (2) seems to be the more controversial. But as I've mentioned, I'm going to take the notion of idealized apriority for granted. Instead, I will use section 2 to consider (1) more carefully.

2 Russellian guises

Russellian theorists claim that $H=H$ and $H=P$ are associated with the same proposition. But Russellians can still distinguish the thoughts $H=H$ and $H=P$ by appealing to the notion of a *propositional guise*.¹⁶ So any Russellian who accepts an account of propositional guises can accept (1). But in this section, I will argue that it is possible to acknowledge a cognitive difference between $H=H$ and $H=P$ while stopping short of accepting Fregean content.

¹⁴ The one place where I will discuss idealized apriority further is in section 5, where I will argue that this notion is incompatible with certain Russellian accounts of propositional guises.

¹⁵ Chalmers' account also requires the assumption that there are token mental states that can stand in relations of negation, disjunction, etc. to each other. Because the vast majority of theorists will grant this assumption and because I will assume that thoughts stand in logical relations throughout this thesis, I haven't listed it separately. But some philosophers (such as Lewis (1994)) reject the view that there are token mental states. Chalmers acknowledges that if someone rejects token mental states, they will need to understand sentence knowledge in a different way.

¹⁶ Of course, there is disagreement among Russellians over what explains the cognitive difference between $H=H$ and $H=P$. But in general, I will use the term "guise" as a blanket term covering all Russellian accounts.

Specifically, I will show that different accounts of the $H=H/H=P$ distinction will generate different spaces of scenarios. But not all spaces of scenarios are capable of grounding primary intensions with suitable Fregean properties. I will argue that, if it is to generate a suitable space of scenarios, any account of the difference between $H=H$ and $H=P$ must meet three constraints. As shall be seen, various Russellian accounts of guises are unable to meet these constraints. This threatens the claim that any theorist who accepts (1) and (2) is in a position to accept the epistemic construction and the derivative Fregean content.

2.1 Constraint 1: The publicity constraint

One important constraint on any adequate account of linguistic/mental content is the publicity requirement: it must be possible at least in principle for different thinkers/speakers to have thoughts/sentences that share content. It follows that, if Chalmers is to use scenarios to ground an adequate account of Fregean propositions, the space of scenarios must be subject-independent. For example, Chalmers wants to develop an account where Oscar's utterance of 'Water is wet' and Twin-Oscar's utterance of 'Water is wet' are associated with the same Fregean content (2011a: 105). But Oscar and Twin-Oscar can only grasp the same Fregean content if the space of scenarios for Oscar is the same as the space of scenarios for Twin-Oscar.¹⁷

¹⁷ Chalmers notes that the picture of a subject-independent space of scenarios has to be modified slightly in order to handle certain types of sentences involving phenomenal demonstratives. But if we restrict the language for scenario construction to eliminate these types of expressions, we can use the resulting scenarios to analyze sentences that do not contain these types of demonstratives. Chalmers adds that, for subjects using phenomenal demonstratives, scenario spaces are still isomorphic. Throughout this paper, I will ignore all complications arising from demonstratives and indexicals.

Chalmers recognizes that, if scenarios need to be subject-independent, then it must be that scenarios (ie, epistemically complete sentences) are built up out of sentence *types* rather than sentence *tokens*. To see why, suppose that scenarios are composed of sentence tokens. One can see from the account of sentence token apriority that any sentence token is paired 1:1 with a thought token. But it doesn't seem like thought tokens are the kind of thing that can be shared by different thinkers. So similarly, a given sentence token is essentially tied to the speaker uttering that token. So we cannot generate a subject-independent account of epistemic space if we compose scenarios out of sentence tokens.¹⁸

So Chalmers (forthcoming) offers the following account of sentence type knowledge: a subject *S* knows a sentence type *t* iff *S* has a thought apt to be expressed by *t* that constitutes an item of knowledge.¹⁹ Similarly we can say that a sentence type *t* is a priori iff there is an a priori thought token apt to be expressed by *t*. Chalmers claims that a thought is only apt to be expressed by a sentence type *t* if it could be expressed by a fully competent and non-deferential utterance of *t*.²⁰ Presumably, a sentence token *t_i* only counts as a competent and

¹⁸ One can observe that there is no similar reason why the *objects* of epistemic possibility must be considered sentence types. In fact, Chalmers wants assign primary intensions to tokens as well as types. So the objects of epistemic possibility are allowed to be sentence types or sentence tokens.

¹⁹ See Excurses 3 of *Constructing the World*. Once again, I am following Chalmers' discussion in CtW because it offers more details than the discussion in *The Nature of Epistemic Space*.

²⁰ In this footnote, I'll explain the restriction to competent and non-deferential utterances. Consider the sentence type 'Water is H₂O'. Most of the time, English speakers use this sentence to express a certain thought about a certain chemical kind. But it seems like the same orthographic expression could be used to express any number of other thoughts. For example, it is surely possible to use 'Water is H₂O' to express the thought that we normally express with the sentence 'All bachelors are unmarried'. But then, according to the definition of sentence type apriority, 'Water is H₂O' counts as an a priori sentence type. By parity of reasoning, it seems like every sentence type will qualify as a priori, which would be a clear violation of

non-deferential utterance of a sentence type t if it expresses a thought of some appropriate type T . So we see that, just as there is a 1:1 pairing of sentence tokens with thought tokens, there is analogous 1:1 pairing of sentence types with thought types. The picture is: scenarios are built up out of sentence types (' $H=H$ ', ' $H=P$ ', etc.) that are distinguished by their corresponding thought types. For an even simpler picture, one can think of building scenarios directly out of thought types.²¹ For the rest of this paper, it will sometimes be more convenient to speak in terms of thought types and sometimes in terms of sentence types, but one should keep in mind the close relation between them.

So it now seems that, if a theorist is to accept Chalmers' construction of scenarios, she must first accept (1'): some account of thoughts that distinguishes $H=H$ and $H=P$ as subject-independent thought *types*. But (1') is stronger than (1): not every theorist who distinguishes the thought tokens $H=H$ and $H=P$ also accepts that $H=H$ and $H=P$ are distinct thought *types*. For example, Field (1977) claims that ' $H=H$ ' and ' $H=P$ ' are distinct because they have different inferential roles, but Field's inferential roles are not subject-independent.²² In other words, Field will accept that that it is possible to know ' $H=H$ ' while failing to know ' $H=P$ ', but Field won't accept that there is a space of epistemically complete sentence

Plenitude (and also Compositionality). Chalmers avoids this worry with the restriction to competent and nondeferential utterances. It is clear that because of this restriction, the sentences of the epistemic construction are *not* typed merely according to orthography. They are interpreted sentences in a language.

²¹ Chalmers acknowledges that scenarios could be built out of thought types instead of sentence types. But he claims that working with sentences is more convenient because the identity conditions for sentences are more straightforward than the identity conditions for thoughts.

²² More specifically, Field cashes out inferential roles in terms of a subjective probability function defined over all the sentences of a speaker's language. The probability function represents how the speaker will update her credence in a certain sentence conditional on the credence she assigns to other sentences in the language.

types. And while it may be safe to reject any account that doesn't distinguish a subject's sincere utterances of 'H=H' and 'H=P', it is not so obviously obligatory that an account must separate these utterances into subject-independent types. I conclude that the epistemic construction is unavailable to any theorist who accepts an account of Russellian guises not satisfying the *publicity constraint*. In some sense, it shouldn't be too surprising that one cannot ground subject-independent Fregean content if one is starting with an account of sentence knowledge that is ultimately cashed out in terms of subject-dependent guises.

Still, many Russellians acknowledge that $H=H$ and $H=P$ fall under distinct subject-independent types. These theorists have a method of typing thoughts that will allow for the construction of subject-independent scenarios. But in 2.2 and 2.3, I will argue that the scenarios generated when typing thoughts according are nonetheless unfit for grounding an account of Fregean content.

2.2 Constraint 2: The holism constraint

Many Russellians have sought to differentiate $H=H$ and $H=P$ as types by appealing to the distinct causal/computational roles associated with these thoughts (see, for example, Block (1986), Harman (1987), and Lycan (1994)). Simplifying, we might say that a token thought A is causally connected to a token thought B if a subject in state A is causally disposed in certain actual and counterfactual circumstances to transition into state B. A thought's causal role captures all of its causal connections to other thoughts.

Causal/computational roles are "intrasubjectively holistic": in order to specify the causal/computational role of a given subject's thought T, one must

specify the subject's entire causal thought "web". T's identity conditions are determined by the causal web in which it is embedded; if the subject's thought web was different, T would no longer be one of the subject's thoughts. To make the same point another way: if two subjects have different thought webs, then those subjects share none of the same thoughts.²³

Per the discussion in 2.1, scenarios are constructed out of sentence types, each of which corresponds to a thought type. Now suppose we type thoughts according to causal role. On this proposal, the sentences of the language will be associated (derivatively) with a causal web, each node of which represents a corresponding thought type. Now presumably, no matter what the causal network associated with the language looks like, there will be some subject B whose thoughts form a different causal network. The problem for the causal role proposal arises when we think about what it means for a scenario to verify one of B's token thoughts.²⁴ A scenario w verifies a token thought T when an epistemically complete thought U associated with w implies t , that is, w verifies T when $\sim U \vee T$ is a priori. So the idea that w verifies T is only coherent if it is possible to entertain the thought $\sim U \vee T$. But if U and T belong to different causal webs, it doesn't make any sense to speak of entertaining the thought $\sim U \vee T$. If we tried to somehow put U and T in the same thought web, they would no longer be

²³ This feature of causal inferential role semantics (CRS) has been often criticized (see Fodor and Lepore (1992), Lepore (1994)). CRS theorists have given a variety of responses. For example, many theorists have claimed that only a subset of a thought's causal transitions are actually meaning-constitutive. While I will briefly return to this issue in section 5, the argument I make in this section is independent of these concerns. The present argument is not an argument against CRS itself; the argument only shows that typing thoughts according to causal/computational role cannot generate a space of scenarios that can ground an account of Fregean content.

²⁴ Remember that Chalmers want to assign epistemic intensions to token thoughts, and the epistemic intension for a thought T is a function from scenarios to $\text{ver}(T,w)$, where w is a scenario.

the same thoughts U and T, but some other thoughts U' and T' that belong to a different causal web.²⁵

The problem is that even though causal/computational webs *can* be shared between subjects (thus satisfying the publicity constraint), they need not be shared. It is easy to see that similar problems will be generated for any proposal on which a subject's thoughts are individuated based only on how they are related to the other thoughts of that particular subject. Methods of typing thoughts that are "intrasubjectively holistic" generate scenarios that are unsuitable for grounding an account of Fregean content. This shouldn't be too surprising, since Chalmers' Fregean propositions are not supposed to be intrasubjectively holistic. I'll call this the *holism constraint*.

2.3 Constraint 3: the invariance constraint

So far, I've talked as if English sentence types might be used in scenarios. But the proposal to use English sentence types runs into immediate difficulties. In fact, Chalmers identifies three constraints that a language must satisfy if it is to be suitable for the construction of scenarios; any language satisfying these constraints is called an *ideal language* (2011a: 75). First, an ideal language must allow infinitary sentences; this is because some epistemically possible sentence tokens can only be verified by scenarios with infinite extent. Second, an ideal language must have a sufficiently broad lexicon in order to have adequate expressive power to describe all maximally specific epistemic possibilities.

²⁵ Of course, if we abstract away from causal roles and consider just the Russellian content of U and T, then it is clearly possible that $\sim U$ and T can stand in a relation of disjunction to one another. But at present, we are considering the causal/computational roles associated with U and T.

Finally, the sentence types of an ideal language must be *invariant* (for an explanation of why the sentences of the ideal language must be invariant, see footnote 28).

A sentence type *s* is *variant* iff it supports potential differences in apriority among fully competent users. A sentence type is *invariant* iff it is not variant. For example, suppose that Leverrier uses the name 'Neptune' synonymously with the definite description 'the body that disturbs Pluto's orbit'. Then it is a priori for Leverrier that 'Neptune disturbs Pluto's orbit'. We can further suppose Leverrier's wife uses the name the name 'Neptune' synonymously with the definite description 'the planet that is typically second furthest from the sun in the Solar System'. Then it is not a priori for Leverrier's wife that 'Neptune disturbs Pluto's orbit'. So, assuming both Leverrier and his wife are competent speakers, 'Neptune disturbs Pluto's orbit' is a variant type.

A sub-sentential expression *e* is invariant if all sentences involving it and only other invariant terms are invariant. For the ideal language to include only invariant sentence types, the ideal language's vocabulary must include only invariant terms. But it is immediately clear that almost all terms of a natural language such as English are variant. Chalmers suggests that certain indexical terms ('I' and 'now') and certain logical and mathematical terms may qualify as invariant as they stand; almost every other term ('Bill', 'know', 'water', 'hydrogen', etc.) must be made more precise if they are to be included in the ideal language. To give an example of just how precise a term must be to be invariant, I will consider sentences involving the term 'know'.

There are many sentences involving the term ‘know’ that are a priori to any competent speaker. For example, suppose G_m is ‘John believes that grass is red & Grass is not red’ and K_m is ‘John does not know that grass is red’. Solely on the basis of the description in G_m and our grasp of the concept of knowledge, most of us judge that G_m implies K_m . So for most of us, $G_m \rightarrow K_m$ is a priori.²⁶ In fact, it is plausible that *any* competent user of the term ‘know’ will judge that G_m implies K_m . Presumably, anyone who doesn’t judge that G_m implies K_m isn’t actually a competent user of the term ‘know’. So the sentence $G_m \rightarrow K_m$ is invariant. But other sentences involving the term ‘know’ provoke disagreement even among competent users. Let G_n be the conjunction of the following sentences in the following passage:

‘Smith believes with justification that Jones owns a Ford. Smith initially has no beliefs about Brown’s whereabouts. Smith forms a belief that Jones owns a Ford or Brown is in Barcelona, based solely on a valid inference from his belief that Jones owns a Ford. Jones does not own a Ford, but as it happens, Brown is in Barcelona’. (Chalmers and Jackson, 2001)

Let $\sim K_n$ be the statement ‘John does not know that Jones owns a Ford or Brown is in Barcelona’. Upon accepting the description in G_n , most English speakers will judge that $\sim K_n$. For these speakers, $G_n \rightarrow \sim K_n$ is a priori. But, as is widely documented, many seemingly competent users disagree about Gettier cases.²⁷ For these users, $\sim(G_n \rightarrow \sim K_n)$ is a priori. If we agree that the disagreeing parties each competently use the term ‘know’, then it is clear that the sentence

²⁶ Some philosophers claim that if $A \rightarrow B$ is a priori, the terms of B must be definable using the terms of A. Chalmers/Jackson describe this view as follows: ‘On this view, a priori [implication] requires definitions, or explicit conceptual analyses: that is, finite expressions in the relevant language that are a priori equivalent to the original terms, yielding counterexample-free analyses of those terms’ (Chalmers and Jackson, 2000). This is not the view of apriority Chalmers is employing. Chalmers/Jackson explicitly deny that ‘knowledge’ can be given an explicit, finite conceptual analysis: this is one of the main ways in which they resist the objections raised by Block/Stalnaker (1999).

²⁷ See Weinberg, Stich and Nichols (2003).

$G_n \rightarrow \sim K_n$ supports differences in apriority between competent users. It follows that the term ‘know’ itself is variant (on the assumption that the rest of the terms in ‘ $G_n \rightarrow \sim K_n$ ’ are invariant). Of course, there are seemingly endless iterations of controversial cases in the Gettier spirit: there would have to be uniform agreement about these cases as well. From this example we can see just how strong the restriction to invariance terms really is: if ‘know’ is to be admitted to the ideal language, *all* competent users have to agree on the a priori status of *every* sentence involving “know”.²⁸

But not every Russellian account has a method of typing sentences that is able to meet the invariance constraint.²⁹ For example, some theorists have suggested that “stereotypes” (Putnam, 1975) can help explain the cognitive significance of certain terms. The stereotype associated with a term C is a set of descriptions that any competent user associates with C. For example, any competent user of the term ‘water’ might associate ‘water’ with the description

²⁸ The motivation for the requirement that the language contain only invariant terms is not immediately clear from Chalmers quick remarks in *The Nature of Epistemic Space*. But here is an argument for why invariance is required. Consider again the example of Leverrier, who uses the name ‘Neptune’ synonymously with the definite description: ‘the body that disturbs Pluto’s orbit’. Leverrier’s wife uses the name ‘Neptune’ synonymously with the definite description: ‘the planet that is typically second furthest from the sun in the Solar System’. Suppose we are constructing a scenario, where, intuitively, there is no body that disturbs Pluto’s orbit but there is a planet that is typically second furthest from the sun. Suppose this scenario can be described with the conjunction D. Let N be the sentence ‘Neptune does not exist’. Then for Leverrier, the sentence $D \rightarrow N$ is a priori, while for Leverrier’s wife, the sentence $\sim(D \rightarrow N)$ is a priori. But now recall the definition of sentence type apriority: a sentence type t is a priori iff there is an a priori thought that is apt to be expressed by t . So both $D \rightarrow N$ and $\sim(D \rightarrow N)$ count as a priori sentence types. The easiest way to see that this is a disastrous consequence is that since both $D \rightarrow N$ and $\sim(D \rightarrow N)$ are a priori, *both* will be implied by any scenario. So every scenario will imply a contradiction. Since every sentence is implied by a contradiction, every scenario will verify every sentence. Most basically, we can see that an inconsistency will be generated *whenever* a sentence s and its negation are both apt to express a thought constituting a priori knowledge. This is why the ideal language must consist of only invariant sentences: any sentence that supports a difference in apriority between competent users will generate inconsistent scenarios.

²⁹ Of course, since sentence types are cashed out in terms of thought types, the invariance constraint on the sentences of the ideal language directly translates into an *invariance constraint* on the thoughts associated with the sentences of the ideal language. But for convenience, I will continue to cast the dialectic in terms of sentences for the rest of 2.3.

'the odorless, colorless liquid... etc.' Some have thought that stereotypes can be used to explain the cognitive similarity between Oscar's utterance of 'Water is wet' and Twin-Oscar's utterance of 'Water is wet'.

Putting aside the many objections that can be made against this type of account of cognitive significance, it is clear that typing sentences according to stereotype is not going to satisfy the invariance constraint. For example, consider again the term 'Neptune'. 'Neptune disturbs Pluto's orbit' is a priori for Leverrier, but it isn't a priori for many competent users who grasp the stereotype associated with 'Neptune'. The basic problem with Putnam's stereotypes is that they are not nearly fine-grained enough to capture the full cognitive significance of an expression. But then it is no surprise that they can't be used to generate scenarios that are able to ground an account of fine-grained Fregean content. We can see that the same problem will arise for any other account of propositional guises that doesn't fix the a priori status of every token sentence falling under a give type.

2.4 Summary

At the end of section 1, we saw that Chalmers' account of Fregean content seems to be available to any theorist who accepts (1) some account or other of the difference between the thoughts $H=H$ and $H=P$ and (2) the notion of idealized apriority. My purpose in section 2 has been to investigate whether it is really true that *any* theorist who acknowledges a cognitive difference between $H=H$ and $H=P$ is in a position to accept Chalmers' Fregean content. To this end, I offered three constraints that any method of distinguishing $H=H$ and $H=P$ must

satisfy if it to be used to generate scenarios that can ground Fregean content. I showed that a variety of well-known Russellian accounts of proposition guises are unable to meet these constraints. (Of course, this “failure” is not necessarily a strike against these accounts; the failure only shows that they are unavailable for Chalmers’ construction).

I don’t think it should be very surprising that the scenario spaces generated by the various Russellian accounts are not suitable for grounding Fregean content. After all, there is a very tight connection between scenarios and primary intensions, and presumably primary intensions have different properties than Russellian guises. Field’s inferential roles can be used to generate a subject-dependent space of scenarios, but this space can’t be used to ground subject-independent primary intensions. Causal/computational roles can be used to generate different spaces of scenarios for subjects with different causal networks, but they can’t be used to ground Fregean content that *isn’t* intrasubjectively holistic. Similarly, coarse-grained stereotypes can’t be used to ground fine-grained Fregean propositions.

On the one hand, the failure of Russellian methods of typing shows that the dispute between the Fregean and the Russellian is not merely verbal. Primary intensions are distinct from Russellian guises in at least one of the following ways: they are publicly sharable, they are fine-grained, and they are not intrasubjectively holistic. But on the other hand, it is no longer clear that Chalmers’ account of Fregean content is available to as wide a variety of theorists as was initially supposed.

Since it is not the case that *any* account of the cognitive difference between $H=H$ and $H=P$ can generate suitable scenarios, what kind of account is required? It would obviously be unsatisfactory to take Fregean content itself as primitive; while this primitive would allow us to satisfy the above constraints, it would open Chalmers' account to a circularity objection. Instead, I will consider in section 3 whether we can use the notion of apriority *itself* to type thoughts in a way that generates an appropriate space of scenarios. If we can, we will thereby show that Chalmers' account only requires (2): a notion of idealized apriority.

3 Apriority roles

A thought token T_i stands in the implication relation to another thought token T_j when $\sim T_i \vee T_j$ is a priori. For any thought token T_i , there will be a set C_i of thoughts implied by T_i and a set D_i of thought tokens that imply T_i . At the very least, C_i will include all logical consequences of T_i and D_i will include all sentences logically implying T_i . Often C_i (or D_i) will include more than the set of logical consequences. For example, the thought *Grass is green* implies the thought *Grass is colored*. Of course, each member T_j of the sets C_i and D_i associated with T_i will itself be associated with a set C_j of thoughts implied by T_j and a set D_j of thoughts implying T_j . When we consider the set of all thoughts taken together, the picture that emerges is of a web of thoughts interconnected by implication relations. We can represent this web with the directed graph $\langle T, A \rangle$, where T is the set of thoughts (the nodes) and A is a set of ordered pairs of nodes such that $\langle T_i, T_j \rangle \in A$ when T_i implies T_j . A sample apriority web is

depicted pictorially in Figure 1, where the nodes are thoughts and the arrows represent implication relations.

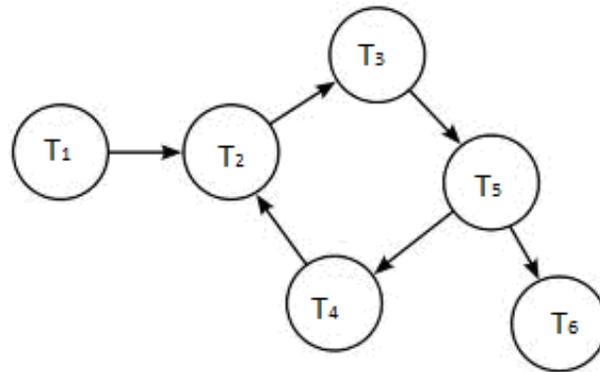


Figure 1: A sample apriority web

We can characterize a thought T_i 's location in a given apriority web with the ordered pair $\langle C_i, D_i \rangle$. We can say that this ordered pair specifies the *apriority role* associated with T_i . In specifying a thought's apriority role, we capture how it is connected to other thoughts via implication relations.

Now consider the proposal to type thoughts according to apriority role. Apriority roles meet the publicity constraint: presumably every subject can in principle have a thought that stands in certain implication relations to other thoughts. Apriority roles also meet the holism constraint: while apriority roles are holistic, they are not intra-subjectively holistic. This can be seen from the fact that, on the idealized notion of apriority, a given thought can imply thoughts outside of its subject's causal inferential web. For example, my thought *Water is a liquid* implies every true mathematical thought even though many of these mathematical thoughts are found nowhere in my causal web. If we type thoughts according to apriority role, there will be a *single* apriority web associated with the

language used to construct scenarios. Apriority roles also meet the invariance constraint, since all thoughts with the same apriority role will have the same a priori connections to other thoughts. In fact, it is easy to see that, if *any* method of typing thoughts is to satisfy the invariance constraint, then all tokens falling under a given type generated by the method must share the same apriority role. Suppose two thought tokens A and B were connected by an implication relation in one subject's thought web but not in another subject's thought web. Then $A \rightarrow B$ will be a priori for the first subject but not the second subject, thus violating the invariance constraint.

3.1 Apriority potentials

Before moving on, I want to address an obvious problem with apriority roles as they now stand. One can observe that all a priori thoughts have the same apriority role. This is because any given a priori thought is implied by all thoughts. Similarly, the conjunction of any thought T with an a priori thought will have the same apriority role as T. We don't want every a priori thought subsumed under the same type; apriority roles as they stand are too coarse-grained. But this is a familiar problem; the familiar solution is to differentiate the thought $2+2=4$ from the thought $3+7=10$ insofar as they are built out of *concepts*. Just as a sentence can ultimately be decomposed into sub-sentential expressions, so too we can say that a thought can ultimately be decomposed into concepts. Of course, just as we required some account of separating $H=H$ from $H=P$, we also need some way to distinguish the concept *Hesperus* from the concept *Phosphorus*. The natural suggestion is to distinguish concepts via their

distinct contributions to the compositional structure of various thoughts (which would themselves be characterized in terms of their a priori connections to other thoughts).³⁰ We can then distinguish between two a priori thoughts by specifying how they are built out of their component concepts. For the rest of the paper, I will assume that the notion of apriority role is fine-grained enough to distinguish between a priori thoughts with different compositional structures.

3.2 Apriority roles vs. Fregean propositions

I'll now explain how my account of apriority roles figures in the dialectic. Chalmers' strategy for motivating his account of Fregean content can be roughly summarized as follows: first, he asks that we acknowledge that there is *some* difference or other in the psychological/epistemological properties of $H=H$ and $H=P$. Any theorist who acknowledges this difference can accept a notion of sentence knowledge that allows for 'H=H' and 'H=P' to have distinct modal properties. With everyone on board with some notion of sentence knowledge, Chalmers uses the notion of *apriority* to define scenarios and the verification relation, which he then uses to ground an account of Fregean content. I think the problem with this strategy is that not every account of sentence knowledge can generate scenarios that are appropriate for grounding Fregean content. The problem seen in section 2 was that, while most Russellians acknowledge a difference between $H=H$ and $H=P$, these accounts generate scenarios that are unsuitable for grounding Fregean content.

³⁰ Roughly, the concept *Hesperus* is the concept which: when combined with *is a planet*, generates a thought with apriority role A_i , when combined with *is Phosphorus*, generates a thought with apriority role A_j , when combined with ..., etc.

My proposal in this section has been to use the notion of apriority *itself* to develop a way of typing thoughts that will allow for the construction of suitable scenarios. I think this proposal is still very much in the spirit of Chalmers' account of Fregean content, given that the core notion of apriority is clearly the distinctive feature of his account. When we tried to use Russellian guises to individuate thoughts for scenario construction, there was a mismatch: the generated scenarios didn't have the right features (subject-independence, fine-grainedness, etc.). But with apriority roles, there is no mismatch, which is not surprising given that both apriority roles and Fregean content are built from the core notion of apriority. The other obvious advantage of the apriority role proposal is that it doesn't require that a theorist accept any new controversial primitives or assumptions; Chalmers' account *already* required the notion of apriority.

Apriority roles avoid the pitfalls mentioned in section 2. Does this mean that any theorist who accepts a notion of idealized apriority can accept Chalmers' primary intensions? Or might there be some *other* reason why the scenarios generated by apriority role typing are unsuitable for grounding Fregean content? Suppose the notion of apriority *is* the only requirement of Chalmers' account. Then we can use the apriority web associated with the ideal language to directly read off the thought types that count as scenarios (a node w will be a member of a scenario when, for every other thought type t , either w implies t or w implies $\sim t$). We can also read off which thoughts are verified by a scenario including w simply by checking which thoughts are implied by w . So if the notion of apriority is all that is required for Chalmers' account, it seems plausible that there will be 1:1

relation between apriority roles and primary intensions.³¹ But if there is a 1:1 pairing of apriority roles with primary intensions, we could equally well use apriority roles themselves in our account of Fregean content.

In fact, I will argue in 3.3 that apriority roles are too coarse-grained to generate an appropriate space of scenarios. To this end, I will give examples that show that there can be a cognitive difference between two thoughts with the same apriority role. This result suggests that an adequate account of Fregean content requires *more* than just the notion of apriority.

3.3 Symmetric apriority roles

The first example is a standard “spectrum inversion” case. Most philosophers agree that it is at least epistemically possible for there to be symmetric color spaces (if one does not accept that symmetric color spaces are epistemically possible, one can skip to the second example). We can suppose that there are pure phenomenal concepts associated with the different qualitative experiences that make up this space.³² It is very plausible that symmetry in a color space will generate “symmetry” in the apriority web (for example: there might be symmetry between the nodes corresponding to the thoughts *This is white* and *This is black*).³³ We can represent the situation with the following

³¹ This is because primary intensions are just functions from scenarios to $\text{ver}(t,w)$. Strictly speaking, I said that our account of apriority roles should somehow distinguish between thoughts with the same a priori connections to other thoughts but different compositional structure. So apriority roles are actually more fine-grained than primary intensions. But Chalmers’ enriched intensions (see footnote 1) take the compositional structure of a thought into account. So if we wanted to be more precise, we would ask: is there a 1:1 pairing between (fine-grained) apriority roles and enriched intensions? But for ease of presentation, I will continue to speak of primary intensions.

³² For a detailed discussion of pure phenomenal concepts, see Chalmers (2003).

³³ For example, the thought *This is white* implies thoughts such as *This is not black* and *This is white or John is tall*. But the thought *This is black* implies thoughts such as *This is not white* and *This is black or John is*

graph (Fig. 2), where T_1 is the thought *This is **white*** and T_2 is the thought *This is **black***:

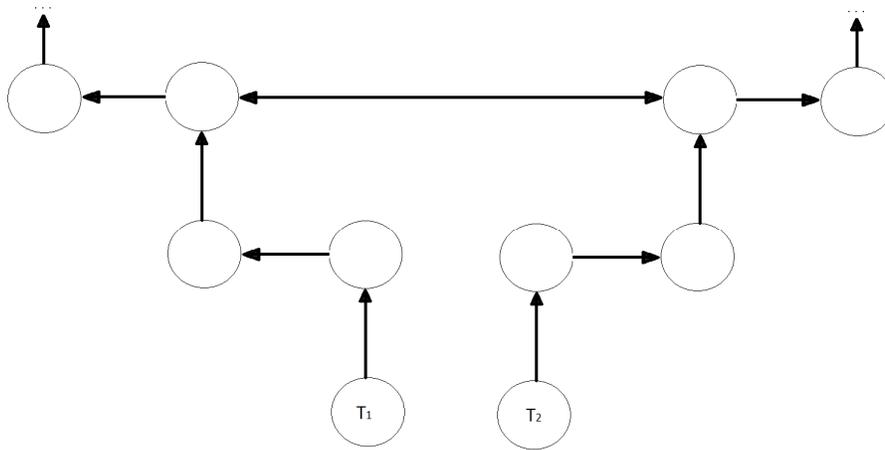


Figure 2: A symmetric apriority web

It is easy to see from the above graph that there is nothing to separate T_1 from T_2 if we characterize T_1 and T_2 only in terms of their a priori connections to other thoughts. So T_1 and T_2 will both fall under the same apriority role type. One might try to appeal to the fact that *This is **white*** and *This is **black*** involve distinct concepts. But this is no help because, for the same reasoning as above, there is symmetry between the concepts **black** and **white**. So apriority roles, as they are currently defined, are too coarse-grained to generate a space of scenarios that can ground a suitable account of Fregean content.³⁴ (Similarly, this shows that apriority roles, as they currently stand, cannot themselves serve as our account of Fregean content.)

tall. This argument echoes a criticism often leveled at functionalist theories of mental states, which, it is sometimes claimed, are unable to separate phenomenal mental states according to functional role alone.

³⁴ This can be seen in the fact that, if we construct scenarios using apriority role typing, T_1 and T_2 will be included in all of the same scenarios (given that they have the same apriority role). This in turn ensures that every scenario that verifies T_1 will also verify T_2 and vice versa, giving both token thoughts the exact same primary intension despite the fact that they are cognitively distinct.

While the above example shows that apriority roles are inadequate as they currently stand, it would be premature to abandon them at this point. One thing to note is that nothing in the above discussion directly hinged on taking the relations between thoughts in the web to be a priori implication relations. Indeed, more familiar causal inferential role theories have also had to confront the problems posed by symmetric color spaces.³⁵ Since this is a problem everyone faces, we should consider whether the solutions employed by other theorists might also help with apriority roles.

Many inferential role accounts eliminate symmetries between thoughts by building thought-world relations directly into the inferential web. For example, some inferential role accounts allow for language “entry points” and language “exit points”.³⁶ The exit points are behavioral states that are caused by a subject’s beliefs and desires, and the language entry points are states that cause a subject to token a certain mental state. While language exit points seem irrelevant to apriority roles, language entry points seem more useful. Suppose we take phenomenal states themselves as entry points and connect these entry points to pure phenomenal thoughts (such as *This is **white***, etc.). Then we can distinguish T_1 from T_2 insofar as T_1 is inferentially connected to a white phenomenal state and T_2 is inferentially connected to a black phenomenal state. The relation between a sensory state and a thought like *This is **white*** isn’t

³⁵ For a discussion of this problem from a causal inferential role theorist, see Harman (1987). Functionalist theories in philosophy of mind have had to confront a closely related problem: how to individuate phenomenal mental states. For a functionalist solution to this problem, see Lycan (1987).

³⁶ This terminology is borrowed from Sellars (2007).

technically an a priori implication relation, but introducing these language entry points doesn't seem to do any great violence to the notion of an apriority role.

I think introducing sensory states as thought entry points can successfully eliminate symmetry between phenomenal thoughts. But I'll now argue that this suggestion is unable to eliminate *all* symmetries in the apriority web. In the philosophy of science, structural realists claim that physics cannot reveal the intrinsic natures of the entities it studies. For example, some structural realists claim that while we can know the structural/causal properties of an electron, we cannot know anything about the fundamental "nature" or "essence" of an electron. We can refer to the potentially hidden intrinsic properties of microphysical entities as *quiddities*.

There are a variety of stances one might take towards quiddities. For example, Chalmers favors a "panpsychist" view on which various causal roles are filled by distinct phenomenal properties. One could also have a view on which quiddities are non-phenomenal intrinsic properties. Whatever stance on quiddities is correct, it doesn't seem like their existence can be ruled out a priori. If this is the case, we will need to introduce sentences/thoughts to scenarios that describe the quiddistic features of those scenarios.³⁷ I will show that certain thoughts needed to specify these quiddistic features have the same apriority roles. For simplicity, I'll give an example of symmetric thoughts involving *phenomenal quiddities* (ie, the quiddities on the panpsychist picture). In other

³⁷ If one *could* rule out the existence of quiddities a priori, one wouldn't need to introduce thoughts describing the quiddistic features of scenarios. The example I'm about to present assumes no such anti-quiddity argument is available.

words, the thoughts I describe below are needed to describe the quiddistic features of “panpsychist scenarios”.

Chalmers claims that, in order to describe the quiddistic features of a panpsychist scenario, we need to introduce a variety of quiddistic concepts p_1, p_2, \dots , etc., where each p_i refers to a distinct phenomenal quiddity (forthcoming: ch. 7). By including these concepts, we can distinguish two scenarios w_1 and w_2 that differ only at the quiddistic level. For example, scenario w_1 might include the thought p_1 fills causal role R while scenario w_2 might include the thought p_2 fills causal role R .

Let T_1 be a thought token p_1 fills causal role R . Let T_2 be the thought token p_2 fills causal role R . It should be immediately clear that there can be symmetry between the concepts p_1 and p_2 , just as there was symmetry between the pure phenomenal concepts **black** and **white**.³⁸ As a result, it does not seem like apriority roles can be used to distinguish the quiddistic thoughts T_1 and T_2 . I also observe that introducing sensory states as language entry points does nothing to eliminate the symmetry between T_1 and T_2 . While we can certainly entertain the thoughts T_1 and T_2 , these thoughts aren’t normatively connected to our sensory states.³⁹ (Incidentally, the quiddity example doesn’t seem to provide a similar challenge to causal inferential role theories. This is because, while it is

³⁸ Indeed, the reason I’ve given an example specifically involving *phenomenal* quiddities is because I think it is *particularly* straightforward to see that phenomenal quiddity thoughts could have the same apriority role. For example, I think that anyone who accepts the possible symmetry between *This is white* and *This is black* should also accept that there can be symmetry between T_1 and T_2 . But I see no reason why symmetric thoughts involving *non-phenomenal* quiddistic concepts could not be generated as well.

³⁹ It is true that T_1 and T_2 are inferentially linked to certain thoughts that are themselves linked to phenomenal states, which might be thought to help matters (for example: T_1 is linked to the thought T_1 v *This is white*). But it is easy to see that for any such link involving T_1 , there will be an analogous link involving T_2 .

tremendously plausible that T_1 and T_2 have symmetric normative roles, there is no reason in particular to think that my thoughts T_1 and T_2 have symmetric *causal* roles.)

3.4 Summary

At the end of section 1, I observed that Chalmers' account of Fregean content seems to be available to any theorist who accepts (1) some account of the cognitive difference between $H=H$ and $H=P$ and (2) the notion of idealized apriority. But in section 2, I argued that various Russellian accounts of the $H=H/H=P$ distinction do not generate a space of scenarios that can be used to ground Fregean content. In this section, I argued that apriority roles can avoid the problems encountered in section 2. Another plus is that accepting apriority roles only requires that a theorist accept (2). But I also argued that apriority role typing encounters its own problem. Since apriority roles are not fine-grained enough to reflect a thought's full cognitive significance, they *also* are unable to generate a space of scenarios that can ground Fregean content. The significance of the quiddity case in particular is that it doesn't seem like there is *any* way to distinguish thoughts like T_1 and T_2 using just the notion of apriority. This suggests that an adequate account of Fregean content will require *more* than just the notion of apriority. In section 4, I will argue that Fregean content also requires the notion of *super-rigidity*.

4 Super-rigidity

What *is* different about the thought tokens T_1 and T_2 ? Here is one obvious difference that the Russellian will probably acknowledge: T_1 and T_2 are

associated with different Russellian propositions. So here is one tempting suggestion for how to avoid collapsing T_1 and T_2 together: type thoughts according to apriority role *and* Russellian content.⁴⁰ While this suggestion may seem to satisfactorily handle the quiddity case, it runs into immediate difficulties when applied to other thoughts. Consider Oscar's thought token *Water is a liquid* and Twin Oscar's thought token *Water is a liquid*. Chalmers wants an account of Fregean content on which Oscar and Twin Oscars' thoughts are associated with the same Fregean content. Oscar and Twin-Oscar plausibly express thoughts with the same apriority role, but they obviously do not express thoughts with the same Russellian proposition. So Oscar and Twin Oscars' thoughts will fall under different types. It looks like this suggestion will introduce two thought types (*Water_{XYZ} is a liquid* and *Water_{H₂O} is a liquid*) where we need one (*Water is a liquid*).⁴¹

⁴⁰ According to a "dual aspect theory", thoughts are associated with Russellian propositions and causal/computational roles (see Block (1986)). This suggestion resembles a "dual aspect theory" with apriority roles replacing the causal/computational role component. One might worry about the fact that the current suggestion presupposes *Russellian* propositions. But Chalmers is a "semantic pluralist": he allows that there could be a variety of semantic values associated with sentences/thoughts. Indeed, he fully endorses the claim that one of the semantic values associated with a thought is its Russellian proposition.

⁴¹ It might have been more elegant to have single type encompassing both of their thoughts, but does the fact that the typing is more fine-grained generate any insurmountable difficulties? One might worry that the fine-grained typing threatens the subject-independence of scenarios; after all, all of Oscar's thought tokens fall under one type and all of Twin-Oscar's thought tokens fall under a second type. But I don't think this is a serious problem: it is at least possible in principle for Oscar to entertain thoughts of the XYZ type (perhaps Oscar moves to Twin-Earth stays there long enough for the extension of his term 'water' changes). But when we think more about what introducing two thought types would look like, we can see that the suggestion to include *Water_{XYZ} is a liquid* and *Water_{H₂O} is a liquid* as separate types for scenario construction is incoherent. I observe that, if *Water_{XYZ} is a liquid* and *Water_{H₂O} is a liquid* are separated as thought types in scenario construction, it must be that *Water_{XYZ} is a liquid* and *Water_{H₂O} is a liquid* tokens occupy distinct nodes in the apriority web associated with the ideal language. But this doesn't seem to make any sense, given that the difference between *water_{XYZ}* and *water_{H₂O}* is not cognitively accessible to the subject. (Of course, Oscar *could* entertain a thought involving the concept *water_{H₂O}**, where Oscar uses *water_{H₂O}** as synonymous with the definite description *the clear, tasteless liquid ... that is identical to H₂O*. But Oscar's thought that *water_{H₂O} is wet* is different from his thought that *water_{H₂O}* is wet*.)

4.1 Super-rigid thoughts

The problem with typing according to apriority role *and* Russellian content stems from the fact that, for example, different tokens of the thought type *Water is a liquid* can be associated with different Russellian propositions.⁴² We can re-describe the situation in epistemic terms: we don't know a priori which Russellian proposition is associated with a given *Water is a liquid* token. As a result, typing according to Russellian content/apriority role is too fine-grained. What's interesting is that this problem doesn't seem to arise when we apply this method of typing to the quiddity thoughts T_1 and T_2 . The same goes for the pure phenomenal thoughts. In these cases, it seems like we *can* know a priori what Russellian proposition is associated with a given thought token.

In a different context, Chalmers discusses the notion of *super-rigidity* (forthcoming: excurses 3). We say that a concept is super-rigid when (i) its extension can be known a priori and (ii) one can know a priori that the concept is metaphysically rigid.⁴³ Intuitively, a super-rigid concept is a concept that takes the same extension across all scenarios and all possible worlds. Chalmers suggests that the following concepts are plausibly super-rigid: pure phenomenal concepts, *and, zero, law, fundamental, some, plus, believes, cause, friend*, etc. We can say that a thought is super-rigid (SR) when it is logically composed of all and only

⁴² In this section, whenever I refer to, say, the type *Water is a liquid*, I mean to refer to a thought type that includes both Oscar's thought token and Twin Oscar's thought token.

⁴³ There are complications with cashing out what it means to know the extension of a concept a priori, but the examples given below should suffice to get across the intuition behind the notion.

super-rigid concepts.⁴⁴ By construction, all tokens of a given SR thought type will be associated with the same Russellian proposition (equivalently: we can know a priori which proposition is associated with a token of a SR thought type). It follows that, when Oscar and Twin-Oscar express thought tokens falling under SR thought types, these tokens are guaranteed to be associated with the same Russellian proposition. So there is no barrier to typing SR thoughts according to (cognitively accessible) Russellian proposition and apriority role (henceforth: “Russellian/apriority typing”). This provides a nice explanation of why Russellian/apriority typing did not present any obvious difficulties in the quiddity or phenomenal examples.

Is the notion of super-rigidity acceptable to a wide variety of theorists? By this I mean: will most theorists accept the claim that certain concepts/thoughts have (a priori) cognitively accessible Russellian content? While I don't think there is any reason why the Russellian qua Russellian should reject super-rigidity, in section 5 I will explain why certain theorists may be inclined to reject the claim that we *ever* have a priori cognitive access to the Russellian proposition associated with our token thoughts. But for now, let's grant the notion of super-rigidity (just as we granted the notion of apriority) and instead consider the question: is a theorist who accepts the notion of apriority *and* the notion of super-rigidity in a position to accept Fregean content? We saw that with just the notion of apriority, we weren't able to construct a space of scenarios appropriate for grounding Fregean content. Can the notion of super-rigidity solve this problem?

⁴⁴ We can also allow for a super-rigid thought to include indexical terms. In this paper, I am ignoring various complications involving indexical terms. So from now on I will often forego explicitly mentioning indexicals.

First, we can consider *super-rigid* thoughts themselves. It is tremendously plausible that Russellian/apriority typing is fine-grained enough to reflect the cognitive significance of SR thoughts: it is very difficult to see how two thoughts could have the same cognitively accessible Russellian content *and* the same a priori connections to other thoughts while still being cognitively distinct. But it isn't as clear that the notions of apriority and super-rigidity allow us to capture the full cognitive significance of *non-super-rigid* (NSR) thoughts, since NSR thoughts don't have cognitively accessible Russellian content. I will consider this question in the next two sections.

4.2 Implied NSR thoughts

Here is an example of NSR symmetry. From the quiddity case, p_1 and p_2 are concepts for two symmetric but distinct phenomenal quiddities. I'll now introduce two further quiddistic concepts p_3 and p_4 for two additional symmetric but distinct phenomenal quiddities. We can let T_3 and T_4 be the thoughts p_3 fills causal role R and p_4 fills causal role R . I note that, while T_1 - T_4 all have the same apriority role, they can be distinguished via Russellian/apriority role typing, per the above discussion. Now consider the following two stipulated concepts:

Let p_5 be a concept such that:

if p_1 fills causal role R , then p_5 's extension is p_1 ; else p_5 's extension is p_3

Let p_6 be a concept such that:

if p_2 fills causal role R , then p_6 's extension is p_2 ; else p_6 's extension is p_4

Let T_5 be the thought that p_5 fills causal role R and let T_6 be the thought that p_6 fills causal role R . It should be clear from their description that T_5 and T_6 have the same apriority role, even though neither is a super-rigid thought (for example, T_5 is associated with a different Russellian proposition when p_1 fills

causal role R vs. when p_1 does not fill causal role R). Because T_5 and T_6 are NSR thoughts, we can't use Russellian/apriority typing in order to separate them.

Nonetheless, it doesn't seem like we need to introduce any new resources to capture the fact that T_5 and T_6 are cognitively distinct. One can observe that T_5 and T_6 are not linked to the SR base in the same way. For example, T_1 implies T_5 but not T_6 , while T_2 implies T_6 but not T_5 . But then, insofar as we are able to avoid collapsing the SR thoughts T_1 - T_4 by using Russellian/apriority typing, we can also avoid collapsing the NSR thoughts T_5 and T_6 insofar as they are linked to T_1 - T_4 in different ways. We see from this example how the notion of super-rigidity can eliminate symmetry worries even among NSR thoughts.

This example suggests that we should distinguish *implied* NSR thoughts (NSR thoughts that, when they are included in a scenario, are always implied by some SR thought or other) from *non-implied* NSR thoughts (all other NSR thoughts). It is very plausible that we capture the full cognitive significance of implied thoughts just by specifying their apriority role (since this apriority role will always include connections to certain "designated" nodes with cognitively accessible Russellian content). If two thoughts (i) are symmetrically entailed by SR thoughts and (ii) have the exact same a priori connections to other thoughts, then it is difficult to see how they could be cognitively distinct.

Although he discusses this issue in a very different context, I mention that Chalmers claims that *all* NSR thoughts are implied by SR thoughts (forthcoming: ch. 8). If this thesis is true, then we can confidently say (on the basis of the above discussion) that any theorist who accepts the notions of apriority and super-

rigidity can accept Fregean propositions. But this thesis is extremely controversial (see footnote 45).⁴⁵ So in the next section, I will assume this thesis is false and that there are non-implied NSR thoughts. So the question remaining is whether we can capture the full cognitive significance of *non-implied* thoughts using only the notions of apriority and super-rigidity.

4.3 Non-implied NSR thoughts

While it was easy to generate an example of cognitively distinct implied NSR thoughts with the same apriority role, it is much more difficult to generate an example for non-implied thoughts. One might interpret the difficulty in generating such an example as evidence that there are no symmetries of this kind. If this were true, then we could just type non-implied thoughts according to apriority role. On the other hand, one might think that the previously observed cases of symmetry count as *prima facie* evidence that it is at least metaphysically possible for there to be symmetry between non-implied thoughts.

It is difficult to assess this issue, but I will make a significant observation. In the previous examples, it seems like apriority roles were too coarse-grained because they failed to capture something cognitively significant about the *extension of the concepts* in the relevant thoughts. For example, in the SR

⁴⁵ More precisely, Chalmers claims that all NSR thoughts are implied by SR thoughts *and* indexical thoughts (for simplicity, I am ignoring complications arising from indexicals in this thesis). Chalmers calls this thesis *Generalized Super-Rigid Scrutability* (GSRs) (forthcoming: ch.7, p.71). To get the intuition behind GSRs, I'll describe how it applies to the actual world. In the actual world, there are many true thoughts involving NSR concepts: *Water is identical to H₂O*, *Electrons and protons attract one another*, *The population of the US is over 10,000*, *Earth has one moon*, etc. There are also, presumably, many true thoughts that involve only super-rigid concepts (thoughts describing the distribution of fundamental properties, true mathematical thoughts, thoughts describing conscious experience with pure phenomenal concepts, etc.). If GSRs is true, then there is some conjunction C of super-rigid/indexical thoughts that are true of the actual world such that C implies *Water is identical to H₂O*, *Earth has one moon*, etc. GSRs is an even stronger thesis than Chalmers and Jackson's well-known *A Priori Entailment* thesis, which states that all truths of our world are implied by PQT1 (the set of microphysical truths, phenomenal truths, indexical truths, and a 'that's all' clause). See Block and Stalnaker (1999) for a series of objections to the A Priori Entailment thesis.

symmetry cases, apriority roles couldn't account for the cognitively accessible difference in extension between the concepts p_1 and p_2 . In the NSR symmetry example considered earlier, apriority roles couldn't account for the cognitively accessible difference in how p_5 's extension and p_6 's extension depend on the super-rigid base. Now consider the case of non-implied NSR thoughts involving, say, the concept *water*. We don't have cognitive access to the extension of *water* (as we did for p_1 and p_2). Nor do we have cognitive access to the way in which *water*'s extension depends on the super-rigid base (as was the case for p_5 and p_6); if we did, then the super-rigid base would always imply a fact of the form *Water is identical to X*, where X is a super-rigid concept. But since, in the previous examples, apriority roles failed to reflect something cognitively significant about the extension of the relevant concepts, we can consider whether there might be something *else* cognitively significant about the extension of the concept *water* that isn't reflected in *water*'s inferential potential.

Here's one possibility of what I have in mind. Suppose we are given a super-rigid base C . Per the above discussion, we can't uniquely identify the extension of *water* on the supposition that C obtains. But let's suppose that, because of something cognitively significant about the concept of *water*, we are able to assign *water* a set of possible extensions (conditional on C). For example, suppose that if we are given C , we know that *water* either takes H_2O as its extension or XYZ as its extension (I'm assuming H_2O and XYZ are super-rigid terms, see footnote 46). Now could there be another concept with the same apriority potential as *water* but with a different class of possible extensions (given

C)? If there could be such an example, then *water's* apriority potential does not reflect *water's* full cognitive significance.⁴⁶

But there is no worry that we could ever generate such an example. This is because any information about the possible extension of a concept is captured in the concept's apriority potential. For example, if it is only possible for water to take H₂O or XYZ as its extension, then every thought type of the form $\sim(C \ \& \ \textit{Water is } ___)$ will be a priori except when the blank is filled with *Water* or *XYZ*. As a result, all of this "information" will be reflected in *water's* apriority potential. But then if some concept *X* were to have the same apriority potential as *water*, all of this information would be equally reflected in *X's* apriority potential. We can refer to thoughts like *Water is H₂O* as *mixed super-rigid identities*.⁴⁷ I think the inclusion of MSR identities in the apriority web shows that, if the apriority roles of non-implied thoughts are too coarse-grained to reflect cognitive significance, it won't be because they fail to capture something cognitively significant about the extension of the concepts composing those thoughts.

⁴⁶ Here are some clarificatory notes on the present example. (1) For ease of presentation, I am adopting the dubious assumption that *H₂O* and *XYZ* are super-rigid concepts. Nothing in the argument turns on this assumption. (2) For simplicity, I am assuming that *water* only has two possible extensions when given C. But nothing in the argument hinges on their only being two possible extensions. Maybe there are 15 possible extensions, or an infinite number of possible extensions. Maybe any *physical* substance described in C (ie, C₆H₁₂, H₂O₂, etc.) is a possible extension for *water*. (3) To be clear, I am not arguing that we actually *can* assign a set of possible extensions to *water* based on some super-rigid base C. I'm only assuming this is possible for the sake of argument, since if it was possible, it could be a way in which two NSR terms with the same apriority potential could be cognitively distinct.

⁴⁷ Up to this point in the thesis, I hadn't acknowledged the possibility of "mixed-super-rigid" (MSR) thoughts. Since MSR thoughts are composed of SR and NSR concepts, it seems like we have cognitive access to at least a part of the Russellian proposition associated with an MSR thought. So typing MSR thoughts according to apriority role alone may not be able to capture the full cognitive significance of such thoughts. Instead, we might type MSR thoughts according to apriority role/"partial proposition", where a partial proposition is some structured entity that reflects what we know a priori about an MSR thought's Russellian content. Do we capture the full cognitive significance of MSR thoughts when typing them in this way? It seems to me that the issues arising for NSR thoughts are also likely to arise for MSR thoughts.

Could there be some *other* way in which two thoughts with symmetric apriority roles could nonetheless be cognitively distinct? The difficulty in coming up with a convincing positive example is prima facie evidence that the answer is no. Consider also the fact that there didn't seem to be any other way for SR and implied NSR thoughts to be cognitively distinct apart from their apriority roles and whatever was cognitively accessible about their Russellian content. On the basis of these considerations, I think the burden of proof resides with any theorist who claims that two cognitively distinct non-implied thoughts can have the same apriority role.

I conclude that, using only the notion of apriority and super-rigidity, we can capture the full cognitive significance of SR thoughts and implied NSR thoughts. The situation isn't as clear for the case of non-implied NSR thoughts, but the overall evidence seems to suggest that we can capture the full cognitive significance for these types of thoughts as well.

4.4 Summary

In section 1, we saw that Chalmers' epistemic construction seems to require (1) some account or other of the difference between $H=H$ and $H=P$ and (2) a notion of idealized apriority. In section 2, I argued that many Russellian accounts of the difference between $H=H$ and $H=P$ are in fact *not* suitable for grounding an account of Fregean content. This presented the question: if it isn't the case that *any* account of the $H=H/H=P$ distinction is suitable to ground Fregean content, is there *some* suitable account of the distinction that is available to a wide variety of theorists? In section 3, I introduced the notion of an

apriority role in order to see whether we might use (2) itself to develop an appropriate account of the $H=H/H=P$ distinction. But it was seen that apriority roles are too coarse-grained to ground Fregean content. In section 4, I sought to identify what *e/se* is required. I argued that there is prima facie reason to believe that any theorist who accepts (1') a notion of super-rigidity and (2) a notion of idealized apriority is thereby in a position to accept Fregean content. By typing thoughts according to apriority role *and* cognitively accessible content, it seems like we can separate cognitively distinct thoughts with the same apriority role.

While I've argued that an adequate account of Fregean content requires (1') and (2) rather than (1) and (2), I think the account of Fregean content built with (1') and (2) is still very much in the spirit of Chalmers' original proposal. Scenarios are still equivalence classes of epistemically complete sentences and the verification relation is still the implication relation. Similarly, the notion of apriority fills a crucial role in each account. The significant difference is that Chalmers' original proposal sought to remain neutral on how to cash out the notion of sentence knowledge, while I've argued that this line is untenable. Instead, I've developed a specific way of cashing out sentence knowledge (in terms of super-rigidity and idealized apriority) that does not presuppose Fregean content. But even though the notion of super-rigidity isn't built into Chalmers' original proposal, I note that Chalmers himself accepts and defends the notion of super-rigidity in other work.

5 The Fregean and the Russellian

One of the puzzling features of Chalmers' original proposal was that, in order to ultimately accept Fregean content, one had to *already* accept some account or other of the cognitive difference between $H=H$ and $H=P$. This was puzzling in some sense because, if one already has an explanation of the $H=H/H=P$ distinction, why should one also need an account of Fregean content? But in section 2, we saw that Russellian guises cannot generate scenario spaces suitable for grounding Fregean content. This was evidence that Russellian guises and Fregean content have different properties. On the basis of the conclusions from sections 2-4, we are now in a good position to compare the relative advantages of the Fregean and Russellian accounts. In 5.1, I will present some potential explanatory advantages of Chalmers' Fregean content, drawing from the discussion in section 2. In 5.2, I will consider why certain theorists may nonetheless have reason to resist Chalmers' account.

5.1 Advantages of Fregean content

The first advantage of Fregean content is that it is very *fine-grained*. In this respect it has a clear advantage over, say, accounts that distinguish $H=H$ and $H=P$ by appealing to descriptive content/stereotypes.⁴⁸ Fregean content is also *subject-independent*. In this respect, they seem to have a clear advantage over Field's subjective inferential roles. For example, with Fregean content, we can

⁴⁸ Chalmers (2011b) argues that Fregean propositions can even account for hyperintensionality. For simplicity, I have usually used the phrase "Fregean content" to refer to a thought's primary intension. With primary intensions alone, we cannot distinguish $2+2=4$ and $3+3=6$. But, as was mentioned in footnote 1, Chalmers distinguishes the thoughts $2+2=4$ and $3+3=6$ by first assigning primary intensions to the concepts, 2, 3, etc. He then defines *enriched intensions* to be structured entities built out of the primary intensions of component concepts. So $2+2=4$ and $3+3=6$ will be distinguished by their enriched intensions.

accommodate the intuition that there is something similar between Oscar and Twin Oscars' utterances of 'Water is a liquid'.

We've also seen that, in contrast to causal/computational roles, Fregean propositions are *not intrasubjectively holistic*. This seems like an advantage, since many of the most persistent criticisms lodged against causal inferential role semantics (CRS) are related to this aspect of causal role accounts. For example, it seems that no two subjects will ever completely agree in all of their counterfactual causal transitions from certain thought/sentences to other thoughts/sentences, etc. But since a thought's causal role is determined by the thought's location in a holistic causal network, this suggests that no two subjects will ever have beliefs with the same meaning (Lepore, 1994). There are a variety of proposed solutions to this problem. Some theorists suggest that CRS can at least provide an account of similarity in meaning by appealing to the notion of similarity in causal role (Block, 1995). But others argue that the notion of similarity in causal role cannot be cashed out without circularity (Fodor and Lepore, 1992). Some proponents have tried to restrict the set of inferences that contribute to a thought's inferential role to a certain subset that are "meaning-constitutive". But such a move seems to require an analytic/synthetic distinction, and many philosophers reject such a distinction. Even if one accepts such a distinction, it is unclear how such a notion can be cashed out in terms of causal dispositions (Fodor and Lepore, 1992). My intent is not to argue against CRS; it may be that proponents of this view can give satisfactory responses to these objections. My claim is only that, insofar as Fregean propositions are not

intrasubjectively holistic, they may avoid certain problems traditionally associated with causal inferential roles.

5.2 Costs of the Fregean account

We see that Chalmers' Fregean content has a variety of attractive explanatory features. Indeed, I don't see any obvious *explanatory* shortcoming of his account. For this reason, I think that if a theorist rejects Chalmers' account, it will instead be because it relies on the controversial notions of *super-rigidity* and *idealized apriority*. In this respect, Russellian accounts of cognitive significance seem to have an advantage. These accounts don't require notions that are as controversial as the notions of super-rigidity and idealized apriority.

It is outside the scope of this paper to assess the general arguments for and against super-rigidity and idealized apriority. Instead, I will conclude this paper by discussing why certain Russellian theorists may have *specific* reason to reject these notions. In this section, I'll show why any Russellian who cashes out propositional guises in terms of causal inferential roles is likely to reject the notions of idealized apriority and super-rigidity.

In section 2.2, I argued that typing thoughts according to causal role cannot generate scenarios suitable for grounding Fregean content. The same argument can be used to show that causal inferential roles are in tension with the notion of idealized apriority itself. For example, it is safe to assume that there is at least one complicated mathematical thought *T* that is not a part of my causal network. But under idealization, *T* counts as a priori. Thus, *T* is implied by any given thought *U* in my causal web (recall that *U* implies *T* when $\sim U \vee T$ is a priori).

For U to imply T, it must be possible to entertain the thought $\sim U \vee T$. But if U and T are members of different causal networks, it doesn't make any sense to speak of entertaining the thought $\sim U \vee T$. More generally, we see that the notion of idealized apriority is only coherent on the assumption that all thoughts are members of a single apriority web. But to be members of a single apriority web, two thoughts U and T need to be logically related in a certain way. But according to CRS, two thoughts can only be logically related in a certain way if they are members of the same causal network. So the CRS proponent can only accept a single apriority web if every possible subject has the same *causal* web. But, as a matter of fact, subjects have different causal webs. So idealized apriority and CRS are incompatible.

Does the incompatibility stem from the idealization or the notion of apriority *simpliciter*? I see no barrier to accepting CRS and accepting a *non*-idealized notion of apriority according to which certain thoughts a priori imply certain other thoughts within the *same* causal network. The incompatibility stems from the idealization's requirement that every thought is a member of the same apriority web.

I'll now argue that there is also tension between CRS and the notion of *super-rigidity*. First, I'll show how the notion of super-rigidity potentially sharpens an argument against CRS made by Fodor and Lepore (1992). Fodor and Lepore argue that CRS proponents need to provide an account of what glues a thought's causal inferential role to its Russellian content. For example, what rules out the possibility that T has the inferential role of the thought *Water is a liquid* but is true

iff 16 is a prime number? Block (1998) claims that the CRS proponent need not be committed to *any* restrictions regarding what causal role can be paired with a given set of truth conditions. According to Block, there is no reason why truth conditions and causal roles could not be completely independent. But if a theorist acknowledges super-rigid content, Block's response is unsatisfactory. In the case of a super-rigid thought type, all tokens have the *same* truth conditions. So if the thought $2+2=4$ is associated with a particular causal role, there must be *something* gluing that causal role to the thought's truth conditions. But then it seems like Fodor and Lepore's glue argument kicks in again: how does the proponent of CRS explain the "glue" between truth conditions and causal role?

There is a second, more basic tension between CRS and the notion of super-rigidity: if one accepts that a certain thought T is super-rigid, T's causal role seems irrelevant to explaining T's cognitive significance. For example, suppose my thought $2+2=4$ is super-rigid. Now further suppose that an alien with a completely different causal network has a super-rigid thought with the *same* cognitively accessible Russellian content.⁴⁹ If we accept that both of these thoughts have the same cognitively accessible Russellian content, then it seems like we would be inclined to judge that they have the same cognitive significance despite the fact that they have different causal roles. This is because the thoughts' cognitively accessible Russellian content seems to provide the best

⁴⁹ This argument obviously requires that it be possible for two thoughts with the same cognitively accessible Russellian content to have different causal roles. But if we grant super-rigid content, it does not seem like it is the type of thing that depends on a thought's causal connections to other thoughts.

explanation of their cognitive significance. When we can compare super-rigid content, difference in causal role doesn't seem to matter very much.⁵⁰

I've argued that the CRS theorist has reason to reject the notions of idealized apriority and super-rigidity. As a consequence, a theorist cannot both accept CRS and Fregean propositions as explanations of the cognitive significance of thoughts. It is worth noting that there does not seem to be similar pressure for proponents of other accounts of Russellian guises to reject Fregean content. I see no reason why a theorist who accepts Fregean content might not also accept that sentences are associated with stereotypes (see section 2.3). While stereotypes are inferior to Fregean propositions when it comes to explaining cognitive significance, it seems like they might be useful for other purposes (perhaps stereotypes are useful for explaining what is required for a speaker to be competent with a public language expression). Similarly, I see no reason why a theorist who accepts Fregean content might not also accept that sentences are associated with Field's subjective inferential roles (see section 2.1) While Field's inferential roles may not explain the similarity between Oscar and Twin Oscars' thoughts, they may be useful for other purposes (for example: modeling the thought processes of non-ideal reasoners).

⁵⁰ One can consider the problem super-rigidity raises for CRS as analogous to a problem that has sometimes been raised against functionalist theories of mind. Some theorists argue that what makes a given pain state S the state that it is just is S's phenomenal character. Accordingly, S's causal connections to other states seem irrelevant to determining whether S is a pain state. For this reason, some theorists have found it plausible that a simple organism with a very different functional organization than our own (ie, a small fish) could nonetheless experience a phenomenal state like our pain states (Block, 1996).

Conclusion

Chalmers wants to develop an account of epistemic space that can be used to ground a widely available account of Fregean content. My purpose in this thesis has been to identify the assumptions required by Chalmers' account. My conclusion is that any theorist who accepts the notion of *super-rigidity* and the notion of *idealized apriority* is thereby in a position to accept a version of Fregean content. So on the model of Fregean content considered in this paper, there are at least two ways in which a thought may be cognitively significant. One aspect of a thought's cognitive significance is reflected in its a priori connections to other thoughts. But certain thoughts are also cognitively significant insofar as they have cognitively accessible Russellian content.

I concluded the thesis by comparing this account of Fregean content to alternative Russellian accounts of cognitive significance. The main disadvantage of the Fregean account is that it relies on the notions of super-rigidity and idealized apriority, both of which are likely to be controversial. I argued that CRS theorists in particular have reason to reject any account appealing to these notions. But if these notions can be adequately defended, Fregean content seems to have various explanatory advantages over Russellian guises.

REFERENCES

- Ned Block (1986). Advertisement for a Semantics for Psychology. *Midwest Studies in Philosophy* 10 (1):615-78.
- Ned Block (1995). An argument for holism. *Proceedings of the Aristotelian Society* 95: 151-69.
- Ned Block (1996). What is functionalism? In D. Borchert (ed.), *The Encyclopedia of Philosophy Supplement*. Macmillan.
- Ned Block (1998). Conceptual Role Semantics. In Edward Craig (ed.), *Routledge Encyclopedia of Philosophy*. Routledge.
- Ned Block & Robert Stalnaker (1999). Conceptual Analysis, Dualism, and the Explanatory Gap. *Philosophical Review* 108 (1):1-46.
- David J. Chalmers (2002). On Sense and Intension. *Philosophical Perspectives* 16 (s16):135-82.
- David J. Chalmers (2003). The Content and Epistemology of Phenomenal Belief. In Quentin Smith & Aleksandar Jokic (eds.), *Consciousness: New Philosophical Perspectives*. Oxford University Press.
- David J. Chalmers (2011a). The Nature of Epistemic Space. In Andy Egan & Brian Weatherson (eds.), *Epistemic Modality*. Oxford University Press.
- David J. Chalmers (2011b). Propositions and Attitude Ascriptions: A Fregean Account. *Noûs* 45 (4):595-639.
- David J. Chalmers (forthcoming). *Constructing the World*. Oxford University Press. Manuscript available online: < <http://consc.net/constructing/>>.
- David J. Chalmers & Frank Jackson (2001). Conceptual Analysis and Reductive Explanation. *Philosophical Review* 110 (3):315-61.
- Hartry Field (1977). Logic, Meaning and Conceptual Role. *Journal of Philosophy* 69, 379-408.
- Jerry A. Fodor & Ernest Lepore (1992). *Holism: A Shopper's Guide*. Blackwell.
- Gilbert Harman (1987). (Non-solipsistic) Conceptual Role Semantics. In Ernest Lepore (ed.), *New Directions in Semantics*, Academic Press.
- Ernest Lepore (1994). Conceptual role semantics. In S. Guttenplan (ed.) *A Companion to the Philosophy of Mind*. Blackwell.

David Lewis (1994). Reduction of mind. In S. Guttenplan (ed.), *Companion to the Philosophy of Mind*. Blackwell.

William Lycan (1994). *Modality and Meaning*. Kluwer Academic Publishing, Studies in Linguistics and Philosophy series.

William Lycan (1987). *Consciousness*. Bradford Books / MIT Press.

Shaun Nichols, Stephen Stich & Jonathan M. Weinberg (2003). Metaskepticism: Meditations in Ethnoepistemology. In S. Luper (ed.), *The Sceptics*. Ashgate.

Hilary Putnam (1975). The meaning of 'meaning'. *Minnesota Studies in the Philosophy of Science* 7:131-193.

Wilfrid Sellars (2007). *In the Space of Reasons*, ed. K. Scharp and R. Brandom. Harvard University Press.

Scott Soames (2005). *Reference and Description: The Case Against Two-Dimensionalism*. Princeton University Press.