

Frege Cases and Bad Psychological Laws

MAHRAD ALMOTAHARI

University of Edinburgh
malmotah@ed.ac.uk

AIDAN GRAY

University of Illinois, Chicago
adgray@uic.edu

We draw attention to a series of implicit assumptions that have structured the debate about Frege's Puzzle. Once these assumptions are made explicit, we rely on them to show that if one focuses exclusively on the issues raised by Frege cases, then one obtains a powerful consideration *against* a fine-grained conception of propositional-attitude content. In light of this consideration, a form of Russellianism about content becomes viable.

Two influential discussions of Frege's Puzzle cite the familiar proverb that hard cases make bad laws.¹ Out of context, the proverb's inclusion might seem like a mere rhetorical flourish. In context, it's obvious that the authors are making a contentious claim about the kinds of cases—namely, *Frege cases*—which give rise to the puzzle. They're suggesting that, in the relevant sense, Frege cases are 'hard', and that, therefore, they shouldn't influence our understanding of the generalizations, or 'laws', that govern propositional attitudes.

This paper is about whether Frege cases are hard in the relevant sense—whether the intentional laws that underwrite psychological explanation must make explicit accommodations for them, or whether they're just the sort of abnormality that the laws can justifiably ignore. It is, in other words, an investigation into the nature of Frege cases.

This paper is also about the individuation of propositional-attitude content. Our plan is to take some assumptions about Frege cases—assumptions that have been mostly implicit in the literature—and make them explicit. Doing so will allow us to rethink the plausibility of two familiar positions: *Russellianism* and *Fregeanism*. With these assumptions explicit, we'll formulate an argument that turns the tables on the widespread belief that Frege cases undermine

¹ See Saul Kripke (1979, p. 270) and Jerry Fodor (1994, p. 45).

Russellianism. The upshot is that Frege cases don't have the right structure to motivate a fine-graining of content.

1. Explanationism and content

Propositional attitudes have content. How should this content be individuated? It's hard to say, but one strategy seems promising. Consider the successful theories in which an appeal to mental content is supposed to be explanatory; now focus on the intentional generalizations that appear in those theories—content should be individuated in whatever way is best suited to play the role that's required by the assumption that those generalizations are true and explanatory. Call this strategy *explanationism*.

Call the way that an attitude represents the distribution of properties and relations over objects its *referential content*.² Contents representing the same distribution of the same properties or relations over the same objects are *referentially equivalent*. Russellianism is the thesis that attitude contents are referential contents. For the Russellian, if contents are referentially equivalent, they're equivalent *simpliciter*.

Now, it's widely thought that explanationism, in conjunction with one's favourite Frege case, requires a fine-graining of attitude content, and thus the abandonment of Russellianism. Exactly how this is supposed to go will be a central concern in what follows. But the basic idea is familiar. The belief that Superman is Clark Kent might play a role in an explanation of why one ran towards the *Daily Planet* when danger loomed. The belief that Clark Kent is Clark Kent didn't, and couldn't, play the same role in an explanation of that behaviour. It seems, then, that explanationism requires us to distinguish the contents of those beliefs. But they're referentially equivalent. So Russellianism appears false.

According to this prevalent line of thought, Frege cases show that attitude content is more fine-grained than Russellians allow: referentially equivalent beliefs can differ with respect to their explanatory role. There are various ways of developing this thought into a more detailed theory. The general strategy is to posit a layer of content that determines reference but that isn't itself determined by reference. Call this additional

² If one likes, one can think of referential contents as Russellian propositions, but we don't want to build this assumption into our official characterization of the notion. We want to cast as wide a net as possible: in principle, theories that represent propositions as cognitive-act types (*à la* Scott Soames 2015), or pairs consisting of truth-conditions and subject-matters (*à la* Stephen Yablo 2014), may qualify as referentialist in our sense. We won't prejudge the issue.

layer of content *sense*. Every representation is thought to express a sense from which its referential content can be recovered. Senses are, therefore, *modes of presentation* (or ‘MOPs’ for short). Call this strategy Fregeanism. Fregeans deny that referential content exhausts propositional-attitude content, but they don’t deny that propositional attitudes have referential content. Sense determines reference, after all.

The traditional path from explanationism to the falsity of Russellianism relies on Frege cases. Treating Frege cases in this way presupposes that a theory of content is answerable to them—that they aren’t dodgy or abnormal. In general, a theory isn’t falsified by dodgy or abnormal cases. This presupposition (that a theory of content is answerable to Frege cases) is so widespread and well-entrenched that its suppression in our summary of the argument probably wasn’t even registered. But some authors believe that we can simply invert the pattern of reasoning above to problematize the presupposition itself, thereby rescuing Russellianism.³ According to these authors, the generalizations in which content plays an explanatory role are *ceteris paribus* laws (or ‘*cp*-laws’ for short), which notoriously tolerate certain kinds of exceptions. These authors believe that Frege cases involve some sort of breakdown of psychologically normal conditions, and that, therefore, they’re *unthreatening* exceptions (not genuine counterexamples) to the intentional generalizations (construed in Russellian terms) which underwrite psychological explanation. Call this view *exceptionalism*.

Russellians and their rivals view the logical landscape differently. Russellians see it like this:

explanationism & Russellianism → exceptionalism

Their rivals view it like that:

explanationism & ~exceptionalism → ~Russellianism

Neither side is likely to get any traction with the other. The debate has stalled. Our aim is to advance it.

We begin in the next section by regimenting some common, though implicit, assumptions about the nature of Frege cases. In §3, we’ll use these assumptions to characterize Fregeanism more precisely. In §4, we’ll draw on our earlier assumptions, and on Stephen Yablo’s neglected discussion of the proportionality constraint on *cp*-laws,⁴ to

³ See David Braun (2000; 2001), Herman Cappelen and Josh Dever (2013, p. 53, fn. 16), Fodor (1994), Susan Schneider (2005), and Soames (1990).

formulate an argument to the effect that,⁵ if one's focus is exclusively on the proper treatment of Frege cases, Russellianism and exceptionism are corollaries of explanationism. In §5 we articulate a version of Fregeanism that withstands our argument, but in a surprising way: it draws no support from Frege cases. We believe this atypical form of Fregeanism is, ironically, the doctrine in its strongest form. Articulating it clearly, and distinguishing it from its weaker cousin, will illuminate how the debate between Fregeans and Russellians should proceed.

2. Russellianism and Frege cases

Recall the basic challenge to Russellianism: Frege cases show that attitudes which are referentially equivalent can differ in the way that they participate in psychological explanations; so content must be more fine-grained than Russellianism allows. It should be clear that evaluating the argument—and, more generally, understanding the import of Frege cases given explanationism—requires some substance about the kind of explanations at issue and how Frege cases disrupt them. Our strategy, in this section, is to extract a set of assumptions from the literature.

First, we'll be assuming that the relevant explanations are slightly regimented folk-psychological explanations. This assumption isn't mandatory. We might, alternatively, think that content should be individuated by its place in cognitive-scientific explanations. The relationship between these two approaches is tricky. But we'll follow the general trend in the literature on Frege's Puzzle and focus on philosophical reconstructions of our everyday practice of folk-psychological explanation.^{6,7}

⁴ The literature on proportionality focuses primarily on how it constrains the identification of causes, not the statement of laws.

⁵ We add this proviso to acknowledge that there may be issues about cognitive significance (most saliently, the *de se*) that pose difficulties for Russellianism but that don't take the form of a Frege case. See §5 for discussion.

⁶ See Ned Block (1986), Brain Loar (1988), Michael Devitt (1989), Fodor (1991; 1994), Gabriel Segal (2000), Braun (2000; 2001), Schneider (2005), and Richard Kimberly Heck (2012; 2014).

⁷ The discussion in Fodor (1991; 1994) focuses on explanations of this kind. But it's clear that Fodor thinks this kind of explanation is relevant to the foundations of cognitive science. This, of course, is because he's optimistic that mature cognitive science will legitimate folk-psychological categories. We take no stand on the matter. If you're disinclined to think that

Second, we'll assume that the relevant form of explanation operates by subsumption of individual actions under law-like intentional generalizations. Again, this is disputable. But, again, to the extent that authors in the literature indicate how they understand the nature of psychological explanation, this often seems to be what they have in mind.⁸ In any case, the idea that explanation, in general, involves subsumption under a law is relatively widespread. So it's natural that it would be the default assumption in discussions of psychological explanation.⁹ And given that psychological explanations aren't fundamental, we'll assume that they involve *cp*-laws.

With these initial assumptions in place, we can ask: how is it that Frege cases are supposed to interact with psychological explanations? In particular, how do Frege cases threaten the adequacy of psychological explanations framed in Russellian terms?

Here's the sort of case one tends to see in the literature: Fred wants to be in the good favour of his elderly neighbour, Samuel Clemens.¹⁰ One morning, at his local café, a mutual friend persuades Fred that Clemens is immensely vain. So Fred sits down to write Clemens an obsequious letter. Call this Case 1.

Given the situation—given what Fred believes and desires—his behaviour is perfectly explicable. The Russellian will claim that Fred's situation is subsumable under a generalization like (R).

- (R). *Ceteris paribus*, if x has a desire with the referential content that x be in y 's good favour, and x has a belief with the referential content that y is vain, x will flatter y .¹¹

We get an instance of (R) by fixing values for the variables.

Let's take a step back to clarify (R), and Russellianism more generally. First, Russellianism, as we understand it, is not a thesis about

cognitive science will vindicate folk psychology, you can treat our discussion as an examination of the interaction between Frege's Puzzle and folk psychology.

⁸ Fodor (1994) and Braun (2000; 2001) are explicit about this assumption.

⁹ For an alternative perspective, see Peter Godfrey-Smith (2005). For some reasons in favour of a nomic conception of folk-psychological explanation, see Fodor (1975, p. 23; 1991).

¹⁰ We adapt this series of examples from Heck (2012).

¹¹ One might worry that (R) illicitly builds in the assumption that the relevant desire is self-directed (because ' x ' occurs in subject position and in the content clause). This might appear like a way of smuggling in something akin to *de se* content for the desire. We address this worry head on in §5, where we argue that if this is the motivation for Fregeanism, Frege cases are orthogonal to the issue.

the semantics of natural-language attitude verbs. It's a thesis about the content of propositional attitudes. We've formulated (R) in an awkward way—'x has a desire' rather than 'x desires'—to bring out the fact we're asking the reader to interpret it in a stipulated quasi-technical way: the antecedent does nothing more than specify the attitude-type (belief, desire, and so on) of a subject's mental state and its referential content. If you're inclined to think that English 'that'-clauses do more than specify referential content, we ask that you treat (R) as involving technical vocabulary (on a par with the standard angle bracket notation for Russellian propositions). Recall that, given explanationism, if generalizations of this sort suffice for underwriting intentional explanation, then Russellianism is true.

Second, (R) does not express Russellianism's full picture of the nature of the attitudes. Standardly, Russellians hold that, at some level of analysis, attitudes are three-place relations between subjects, referential contents, and some third thing: a *guise*, *notion*, *concept*, *Mentalese term*, or whatever. But it's crucial to Russellianism, as we understand it, that none of these things are referenced in intentional *cp*-laws. The question we've set for ourselves is whether the law-like generalizations of folk psychology must be formulated in terms of anything other than referential content to subserve our ordinary explanations of action. The kind of Russellianism for which we will argue sees the three-place analysis of attitudes as abstractly characterizing the psychological machinery that *implements* the fully intentional law-like generalizations of folk psychology. To put it picturesquely, the generalizations of folk psychology don't 'see' any of the implementing details; they only see referential content.

This shouldn't be a surprise. This kind of distinction—between the features of a state that are relevant to a certain class of explanatory generalizations, and the full story about its underlying nature—is required in many forms of inquiry. The explanatory generalizations of biology, for example, don't mention the details of the implementation of biological properties in the underlying physical substances that instantiate them. The Fregean, just as much as the Russellian, will have to make some distinction of this kind. The Fregean thinks that psychological laws are framed in terms of intentional features that go beyond referential content. But, presumably, they don't hold that the laws make explicit reference to the full underlying nature of attitude states. (To put this another way: though Russellians and Fregeans

draw the line between semantics and metasemantics in different places, they each draw the line somewhere.)

To return to the main thread, (R) seems to make good sense of Case 1. The relevant instance looks like this:

- (R₁) *Ceteris paribus*, if Fred has a desire with the referential content that Fred be in Twain's/Clemens's good favour, and Fred has a belief with the referential content that Twain/Clemens is vain, Fred will flatter Twain/Clemens.¹²

But here is where Frege cases pose a problem. Suppose Fred is told, and comes to believe, that *Twain* is vain. As before, he desires to be in Clemens's good favour. He knows that Twain is a famous author. But he wouldn't assent to 'Twain is Clemens'; nor would he assent to 'I desire to be in Twain's favour'.¹³ The news that Twain is vain doesn't lead Fred to write a letter, and he makes no effort to flatter anyone. Call this Case 2.

Applied to Case 2, (R) again yields (R₁), which now has a true antecedent and false consequent. In Case 2, Fred has attitudes with appropriate referential content, and yet he doesn't flatter anyone. So it looks as though Fred is a counterexample to (R) in Case 2. We *could* hold that other things aren't equal between Case 1 and Case 2. This would be to treat Fred in Case 2 as an unthreatening exception to (R) rather than a counterexample. But on what basis should we do that? We can stipulate that everything else—that is, everything except for the 'guise' under which Fred forms the belief—is held fixed. (Fred isn't drunk or tired; he doesn't have any other pressing projects; and so on.) The Russellian seems forced to say that what's abnormal about Case 2 is that Fred is in a Frege case.

Going forward, we'll need a careful expression of this idea. To that end, we need to introduce the notion of *coordination*. Two representations of the same object are coordinated if and only if the identity of their referents is 'directly' encoded in the way they're represented. 'Directness' here is supposed to contrast with an explicit representation of identity (*cf.* the distinction between *representation as the same*

¹² A note about the slash notation: we use it to emphasize that a referential position is to be understood as transparent. It doesn't indicate that the subject of the attribution is aware of the relevant identity.

¹³ We talk about assent to sentences to make the description of the case consistent with Russellianism about the semantics of ascriptions. In other parts of the paper, we're less careful. We ask that readers make appropriate substitutions if necessary.

and representation as being the same in Kit Fine (2007), and the notion of trading on identity in John Campbell (1987)). It's the relation we capture by using the same schematic letter when we formalize predications—for example, in the inference from Fa and Ga to $\exists x(Fx \ \& \ Gx)$. The re-occurrence of 'a' in these formalizations directly indicates that we're representing the same individual twice over. We don't assume that coordination in thought or natural language is constituted by re-occurrence of symbols (although, as we note below, this is a possible view). We use the familiar device of singular-term recurrence in a formal language only to fix attention on the coordination relation.

The nature of coordination is central to the debate between Russellians and Fregeans. Put briefly, Fregeans treat coordination as a feature of intentional content; Russellians don't. We'll discuss the Fregean picture of coordination below. Russellians treat coordination as a feature of the *implementation* of the attitudes. The exact nature of coordination will depend on the details of the implementational picture. One simple idea is that two representations of the same object are coordinated when the mental names responsible for them are tokens of the same Mentalese type. We won't be assuming anything that determinate here; we'll only assume that the implementing states for the attitudes, whatever they are, have enough structure to determine a pattern of coordination and its absence.

A Frege case is, then, a situation in which a subject has uncoordinated attitudes about the same object. In Case 2, for example, the belief that Fred would express with 'Twain is vain' and the desire he would report with 'I want to be in Clemens's good favour' are uncoordinated (or, more carefully, the relevant representations of Twain/Clemens are uncoordinated).

Returning to what the Russellian must say about Case 2, she must embrace:

(*Exceptionalism*) For all subjects, S , and intentional generalizations, G , if (i) G involves coreference at referential positions $r_1 \dots r_m$ (ii) S satisfies the antecedent of G in virtue of attitudes that are not coordinated at $r_1 \dots r_m$ and (iii) S fails to satisfy the consequent of G , then S isn't (thereby) a counterexample to G .

Three preliminary remarks about *Exceptionalism*: First, the ‘positions’ in an intentional generalization are just the particular aspects of referential content that they attribute to subjects that satisfy them. A generalization involves coreference at two positions when they’re introduced by occurrences of the same variable (or, equivalently, when any subject who instantiates the generalization will have attitudes that are coreferential at those positions).

Second, *Exceptionalism* appeals to coordination, which isn’t, for the Russellian, a feature of intentional content. But it shouldn’t be surprising that to specify the nature of exceptions to a *cp*-law, we must characterize them in terms of the vocabulary of a lower level. In fact, this is exactly what we should expect. We find unthreatening exceptions to a *cp*-law when lower-level implementation of a higher-level property is non-standard in some way. For example, the existence of albino ravens doesn’t falsify

(B) *Ceteris paribus*, if x is a raven, then x is black.

Albino ravens are unthreatening exceptions to (B) because, though they instantiate *ravenhood*, they do so in a genetically non-standard way (Bernhard Nickel 2010). Similarly, *Exceptionalism* claims that although Fred instantiates the antecedent of (R) in Case 2, he does so in a psychologically non-standard way.

Finally, note that *Exceptionalism* is *not* the claim that being in a Frege case, *as such*, is abnormal. It is a much more targeted claim: having attitudes about an object o that would rationalize a particular o -directed behaviour, but having them in an uncoordinated way, is abnormal with respect to intentional explanation. It is consistent with *Exceptionalism*, for example, that in Case 2 Fred might non-deviantly satisfy many other intentional generalizations in virtue of his Clemens/Twain directed attitudes (just as albino ravens non-deviantly satisfy many other raven-involving generalizations).

However, in the context of a choice between Russellian and Fregean theories of content, reliance on *Exceptionalism* looks to be an egregious form of special pleading. A commitment to explanationism only amplifies the worry, since it requires that we be guided *to* a theory of content by an unbiased assessment of our most promising intentional explanations, not tailor our assessment of these explanations to fit a particular theory of content we happen to antecedently endorse!

Russellianism is in trouble. Unless we can answer the charge of special pleading, we will need to replace (R) with a generalization

that presupposes a more fine-grained conception of content. We will need a conception of content that distinguishes the content of Fred's attitudes in Case 1 and Case 2.

Before characterizing how Fregeans would handle Cases 1 and 2, we need to make one more assumption explicit. We'll assume that the way (R), Case 1, and Case 2 interact is characteristic of the way Frege cases bear on explanationism. That is to say, we'll assume that the intentional laws to which Frege cases pose a potential problem have a form that's broadly analogous to (R) and that Frege cases pose a potential problem for them in the same way that Case 2 poses a potential problem for (R).

To be more explicit: we'll assume (*i*) that the sorts of intentional generalizations to which Frege cases pose a potential problem involve antecedents which cite distinct attitudes about the same object, and (*ii*) that Frege cases interrupt the applicability of an intentional generalization by breaking the coordination between the attitudes cited in their antecedents. For the sake of convenience, we'll stipulatively introduce a shorthand slogan for this package of assumptions: Frege cases are relational.

This is the most 'extrapolative' assumption we've made so far. You won't find it explicitly discussed in the literature.¹⁴ But when you see discussions of Frege's Puzzle and psychological explanation, you'll typically find examples relevantly like Case 2, and discussions of them which implicitly satisfy assumptions (*i*) and (*ii*). This assumption will play an important role in our argument that Frege cases pose no threat to Russellianism.

3. Fregeanism

We saw above that Case 2 poses an apparent problem for Russellian interpretations of intentional generalizations. If we frame intentional generalizations in Russellian terms, in order to account for Case 1, our generalization, (R), will be in danger of falsification by Case 2. The Fregean takes this as grist for her mill.

¹⁴ There are some exceptions, of course. For example, [Braun's \(2001\)](#) focus on believing propositions in matching ways is clearly relevant here. That idea is taken up in [Schneider \(2005\)](#). The recent literature on Relationism is also relevant. See [Fine \(2007\)](#) and [Heck \(2012\)](#), among others. The connection between Russellian Exceptionalism and Relationism is complicated enough to require separate treatment, which it receives in [Mahrad Almotahari and Aidan Gray \(manuscript\)](#).

Recall that Fregeans take situations like Case 2 to show that we need to individuate content more finely than referential equivalence: they hold that intentional contents contain MOPs for the objects they're about. Notoriously, what MOPs are is a vexed question. We'll avoid wading into these dark waters. Given the assumptions we've made, there's a more straightforward path to problematizing the Fregean response to Frege cases. We'll present our argument in §4.

One thing that all Fregeans have in common is that, whatever they think MOPs are, they understand coordination as *sameness* of MOP.¹⁵ That is, Fregeans hold that two representations of the same object are coordinated if and only if the object is presented via the same sense in both. This assumption will guide how the Fregean revises (R).

What's required to explain Fred's failure to act in Case 2 is *not* that Fred thinks of Twain/Clemens in a *specific* way—say, in the *Twain* way or the *Clemens* way. Rather, it's that the two attitudes involve thinking of Twain/Clemens in *different* ways. This is to say that if the Fregean is motivated by Case 2 to revise (R), their new generalization won't refer to specific MOPs, but will generalize over them:

- (F) *Ceteris paribus*, if x has a desire with the referential content that x be in y 's good favour in which y is presented via MOP m , and x has a belief with the referential content that y is vain in which y is presented via MOP m , x will flatter y .

The Fregean's account relies on one more assumption: that in Cases 1 and 2 the attitudes that Fred would express with 'Twain' involve a different MOP from the attitudes that Fred would express with 'Clemens'. Let's call the former 'TWIN' and the latter 'CLEMENS'. We make no assumptions about their nature except that they're distinct.

Given these assumptions, the Fregean's account of Case 1 and Case 2 is straightforward. In Case 1, the relevant instance of (F) is

- (F₁) *Ceteris paribus*, if Fred has a desire with the referential content that Fred be in Twain/Clemens's good favour in which Twain/Clemens is presented via CLEMENS, and Fred has a belief with the referential content that Twain/Clemens is vain in which Twain/Clemens is presented via CLEMENS, Fred will flatter Twain/Clemens.

¹⁵ At least when it comes to intrapersonal, synchronic coordination, which is the only kind of coordination that will be at issue here.

Relative to Case 1, (F_1) has a true antecedent and a true consequent. The Fregean maintains that it underwrites the explanation of Fred's behaviour in that case. In Case 2, however, (F_1) has a false antecedent. And no relevant instance of (F) has a true antecedent, because instances of (F) require sameness of MOP across the belief and desire. So Case 2 is no threat to (F) or any of its instances.

4. Proportionality

We're now in a position to turn the tables. Given the assumptions we've made, not only is Russellian Exceptionalism innocent of the charge of special pleading (at least with respect to its handling of Frege cases), but it now looks as if Fregeanism is guilty of it. At any rate, that's what we hope to show here.

Our argument is inspired by Yablo (1997). But we'll need to expand on his discussion substantially. And, at one pivotal moment, the assumption that Frege cases are relational will play a role.

We've assumed that intentional explanation is *nomioic*: that it explains an agent's behaviour by reference to a *cp*-law. Since a generalization qualifies as a law only if it exemplifies the appropriate form of generality, we require some reliable method of assessing whether the proposed Fregean generalizations are appropriately general. Here is where we draw inspiration from Yablo's discussion of the *proportionality* constraint on *cp*-laws.

Yablo observes that there are two opposite ways in which our formulation of a law might go wrong. Consider two candidate laws in materials science:

(M) *Ceteris paribus*, matter conducts electricity.

(P) *Ceteris paribus*, pennies conduct electricity.

(M) goes wrong because lots of matter doesn't conduct electricity. More importantly, lots of non-conductive bits of matter are paradigmatic enough, *qua* matter, to falsify (M). (P) goes wrong for a different reason. Pennies do, admittedly, conduct electricity; but not *because* they're pennies. (P) incorporates nomically irrelevant detail given its explanatory aims. Yablo doesn't say this explicitly, but we take it that, on his view, (P) isn't false. It's simply not a law of materials science, and thus doesn't characterize explanatory relationships in that domain.

When a generalization is free from the defects that Yablo (1997, p. 285) identifies in (M) and (P)—that is, when it's neither 'underspecific' nor 'overspecific'—we'll say that it's *proportional*. Going forward, we need a more precise way of thinking about nomic proportionality. Unfortunately for us, Yablo doesn't offer one. His focus is on proportionality in causation.

With respect to causation, Yablo says that for a condition that's *causally relevant* to an outcome to be *the cause* of it, it mustn't contain a substantial amount of detail that's irrelevant to the occurrence of its effect. But it also mustn't be so weak as to require substantial help from other independently contributing conditions. If we have two conditions, C_1 and C_2 , that are both causally relevant to a third condition, E , we'll say that C_1 *screens off* C_2 if and only if, had C_1 obtained without C_2 , E would still have obtained (Yablo 1997, p. 266). Even conditions that fail to be proportional (and thus fail to be *the cause*) may still be causally relevant. Failing to be proportional means incorporating too much detail relative to the effect in question. Conditions that incorporate too much detail *eo ipso* incorporate *enough* detail to be *causally sufficient* for the effect. Causal sufficiency implies causal relevance. In a moment, we'll clarify all of this with a familiar example. What we want to emphasize here is that the clearest instance of the proportionality requirement is this: a cause must screen off its determinates and not be screened off by its determinables.

Suppose I've trained my pet pigeon, Sophie, to peck at a button when it's illuminated with a red light, but not when it's illuminated with a green light.¹⁶ On one occasion, the light is scarlet and Sophie pecks. Consider three candidate causes: the button's being *illuminated*; the button's being illuminated *with red light*; the button's being illuminated with *scarlet* light. Each of these conditions is more determinate than the last, but only the second condition is proportional to the pecking. If the button had been illuminated without being red, Sophie wouldn't have pecked. If the button had been illuminated red without being scarlet, Sophie would still have pecked. So being illuminated red screens off being illuminated scarlet, and it isn't itself screened off by being illuminated. Finally, note that although being scarlet isn't the cause of the pecking, it is causally relevant. After all, it was causally sufficient for Sophie to peck.

¹⁶ This example is from Yablo (1992).

We'd like to acknowledge one potentially awkward feature of the framework here. Borrowing ideas from Yablo, we've spoken of how to identify the cause of an event. To some ears, it might sound odd to single out *the* cause of some event from among the events that are causally relevant to it. Or it might seem that any such singling out will be context-sensitive or interest-relative. Though we sense the attraction of this sort of view, we're convinced by Yablo's discussion that the proportionality constraint is a genuine tool for identifying the causal/explanatory structure of a domain. And we take it that reflection on Sophie and the coloured lights, or on pennies, copper, and conductivity, supports our position. We do not assume that there will *always* be a matter of fact about which of various causally relevant conditions is *the* cause of an event. We only claim that it's sometimes possible to use considerations of proportionality to privilege some conditions over others. So proportionality should be taken seriously by anyone adopting the explanationist perspective on attitude content. It isn't clear how one could be an explanationist without thinking that we can isolate, from among the various causal antecedents of behaviour, a level of description that is proportional to the behaviour.

So far, we've summarized the essential components of Yablo's conception of *causal* proportionality. In order to obtain a useful notion of *nomical* proportionality, we need to do some work. First, we need to characterize the circumstances in which one generalization, G , is a determinable of another generalization, G' . We'll assume this occurs when the antecedent condition of G is a determinable of the antecedent condition of G' , and the consequent of G is the same as the consequent of G' .¹⁷ Next, we'll characterize *nomical screening off* in terms of Yablo's definition of screening off *simpliciter*. To do this, we must assume that each domain of explanation comes along with a range of conditions that are normal given its explanatory demands. These assumptions allow us to say that G *nomically* screens off G' if and only if every normal instance of G screens off (*simpliciter*) the corresponding instance of G' .¹⁸

¹⁷ Like Yablo, we acknowledge that maybe not every aspect of the determinable-determinate distinction is relevant here: '[...] all that "Y is a determinate of X" needs to mean in this paper is that Y necessitates X (not because it has a metaphysically infallible way of bringing X about but) because X is immanent in or included in Y' (1997, p. 275, fn. 22).

¹⁸ You might worry that applying this test in the way we have amounts to the assumption that psychological explanation is causal explanation. We're comfortable with that assumption, but we're actually assuming something substantially weaker: that Yablo's counterfactual test for screening off (*simpliciter*) is relevant to the proportionality of intentional *cp*-laws. That test,

Finally, inspired by Yablo's discussion of (M) and (P), we'll assume that a generalization qualifies as a *cp*-law only if it satisfies the nomic proportionality constraint: the generalization must nomically screen off its determinates and not be nomically screened off by its determinables. So, for example, (E) nomically screens off (P).

(E) *Ceteris paribus*, copper conducts electricity.

In any normal case of a copper penny conducting electricity, if the piece of copper hadn't been a penny, it would still have conducted electricity.¹⁹

Our criticism of (F) will closely parallel Yablo's take on (P). We won't challenge the truth of (F), but rather its status as a law of folk psychology. Our conclusion will be that (F) doesn't identify the explanatory structure of an agent's state of mind. Strictly speaking, we neither affirm nor deny the truth of (F). (But at various points we will speak *as if* (F) is true because doing so will facilitate a clearer presentation of our central argument.) Officially, then, our position is this: even if there *were* MOPs, and even if (F) *were* true, (F) wouldn't be an explanatory generalization. Given explanationism, and given the choice between (R) and (F), reflection on Frege cases puts pressure on us to opt for (R). To adequately address our challenge, it isn't enough for the Fregean to demonstrate (F)'s truth. The Fregean must show that (F) is a law: either it satisfies nomic proportionality, or its failure to satisfy nomic proportionality is somehow outweighed.

With all of this in place, let's return to psychological explanation. Note that because Fregean contents have a finer granularity than referential contents (indeed, each Fregean content *determines*, by design, a corresponding referential content), Fregean intentional generalizations will be determinates of Russellian intentional generalizations.²⁰

itself, makes no reference to causation. To deny this assumption one must either deny that intentional *cp*-laws are subject to a proportionality requirement, or offer an account of proportionality that doesn't appeal to counterfactuals of this kind.

¹⁹ (E) itself may be screened off by one of its determinables. Perhaps it's screened off by '*Ceteris paribus*, metal conducts electricity'.

²⁰ According to the conception of Fregean content presented in David Chalmers (2002), the entities best suited to play the explanatory role characteristic of propositions are *primary intensions*: roughly speaking, functions from possibilities considered as actual to extensions. But a primary intension doesn't determine referential content all by itself. Chalmers (2011a, p. 601) acknowledges this as one way in which his earlier view falls short of fully vindicating a Fregean conception of sense. According to his official position, Fregean contents are *enriched intensions*: ordered pairs consisting of a representation's primary intension and its actual

So Fregean generalizations are proportional only if they aren't nomically screened off by Russellian generalizations. To determine whether Fregean generalizations satisfy this condition, consider (R) and (F) again.

- (R) *Ceteris paribus*, if x has a desire with the referential content that x be in y 's good favour, and x has a belief with the referential content that y is vain, x will flatter y .
- (F) *Ceteris paribus*, if x has a desire with the referential content that x be in y 's good favour in which y is presented via MOP m , and x has a belief with the referential content that y is vain in which y is presented via MOP m , x will flatter y .

Is (F) nomically screened off by (R)? The answer depends on whether, for each normal instance of these generalizations, the instance of (R) screens off the corresponding instance of (F). Assuming that Case 1, and the instances of (R) and (F) which apply to it, are paradigmatic, we can focus on them and then generalize accordingly. Is (F₁) screened off by (R₁) *vis-à-vis* Case 1?

- (R₁) *Ceteris paribus*, if Fred has a desire with the referential content that Fred be in Twain's/Clemens's good favour, and Fred has a belief with the referential content that Twain/Clemens is vain, Fred will flatter Twain/Clemens.
- (F₁) *Ceteris paribus*, if Fred has a desire with the referential content that Fred be in Twain's/Clemens's good favour in which Twain/Clemens is presented via CLEMENS, and Fred has a belief with the referential content that Twain/Clemens is vain in which Twain/Clemens is presented via CLEMENS, Fred will flatter Twain/Clemens.

Take Case 1 as a description of the actual world. Now ask: if Fred had satisfied the antecedent of (R₁) (that is, if he had had a desire with the referential content that he be in Twain's/Clemens's good favour and a belief with the referential content that Twain/Clemens is vain), but had failed to satisfy the antecedent of (F₁) (that is, if he hadn't

extension. The enriched intension of a belief is an *enriched proposition*, and 'enriched propositions are the best single candidate in the two-dimensional framework for playing the core explanatory roles associated with propositions' (p. 603). Enriched propositions determine referential contents (p. 602).

deployed CLEMENS in both of those attitudes), would he still have flattered Twain/Clemens?

To evaluate this counterfactual we must consider *minimal departures* from Case 1. That is, we need to consider the minimal departure from actuality required to make Fred instantiate the antecedent of (R) otherwise than by deploying CLEMENS in both attitudes. After all, it's widely known that counterfactuals are evaluated with respect to the nearest possibility where the antecedent holds.

This interpretive requirement raises some difficult questions. First, we haven't said much about the details of Case 1. (How did Fred come to have the relevant attitudes? What other projects does he have? And so on.) So it's difficult to know what the nearest possibilities in which Fred employs different MOPs are like. More than that, we haven't said anything about MOPs or how to individuate them. So it isn't clear what, exactly, we're asking. If MOPs are very coarsely individuated, then the nearest possibility in which Fred employs different MOPs might involve a very different cognitive situation.²¹

At this point, one might despair. Without a principled Fregean account of the individuation of senses—something that's never been easily forthcoming—we'll be unable to settle the question. In fact, though, we think we can bring out the implausibility of the Fregean position without saying very much about the nature of sense.

Suppose that in the nearest possible world in which Fred satisfies the antecedent of (R₁) but not that of (F₁), he deploys the *same* MOP for Twain/Clemens in the relevant belief and desire. Given the assumption that Frege cases are relational—that Frege cases disrupt the applicability of Russellian generalizations by breaking the coordination across some relevant attitudes—it follows that the change in MOPs would *not* alter Fred's behaviour, because (by supposition) sameness of MOP, and thus coordination, is preserved. The relevant world would be one in which Fred had coordinated attitudes of the right sort—perhaps he would believe, as we might naturally put it, that *Twain* is vain and desire to be in *Twain's* good favour—so we are entitled to expect him to behave in the same way.²² This would show that (F) was nomically screened off by (R).

²¹ In the discussions from which we draw inspiration, both Fodor and Yablo assume that MOPs are very finely individuated. So it appears to be part of their arguments that an agent's cognitive situation will *not* be very different in the nearest possibilities. One nice feature of our argument is that we make no such assumption.

So the Fregean *must* hold that the nearest world in which Fred has attitudes with the same referential content, presented with different MOPs, is a world in which the belief and the desire involve *distinct* MOPs. And it might seem that this is easy for the Fregean to maintain: after all, wouldn't the minimal change in Fred's MOPs involve replacing only *one* of the occurrences of CLEMENS by a different MOP?

There are very good reasons to be sceptical of this claim, and, therefore, very good reasons to doubt that the Fregean can ultimately reconcile her intentional generalization with nomic proportionality. The mistake here is to think of the distribution of MOPs *atomistically*. In normal situations in which generalizations like (F) apply—that is, normal situations in which an agent possesses attitudes about an object that would rationalize a specific behaviour directed at it—it's not an *accident* that the subject deploys the same MOPs in the two attitudes. Rather, the sameness of the MOPs is the upshot of a rational subject's sensitivity to evidence of the identity of the objects about which they are forming attitudes. And this fact about their aetiology means that the *sameness* of the MOPs will be held fixed in nearby possibilities in which *either* is changed.

Before arguing for this claim, we'll offer an analogy to clarify it. Suppose that two American pennies, p_1 and p_2 , are minted consecutively from the same mint: p_1 was minted on 4 June 1985 at 13:23:35; p_2 was minted on the same day at 13:23:36. They were each minted with the Lincoln Memorial Design (which was used from 1959–2008). And, for the sake of argument, suppose the mint was operating smoothly: no mistakes could easily have been made. We can now ask: is the nearest world in which at least one of either p_1 or p_2 has a different design a world in which they *share* a design or a world in which they have *different* designs?

It seems that, given normal ways of filling in the details here, the answer will clearly be that the relevant world is one in which they share a design. It's easy to imagine a slight difference in the 1959 design (or a decision to change the design some time in the intervening years). But we would have to depart substantially further from actuality to find a world in which the mint decides to stop the presses

²² If you're sceptical of this point, we urge you to suspend judgment until §5. You might be convinced that different pairs of coordinated attitudes—say, TWAIN-TWAIN pairs and CLEMENS-CLEMENS pairs—can differ in the way they rationalize actions. That is consistent with the point we're making here. We only hope to establish that *Frege cases* don't warrant any such view.

on a particular Tuesday afternoon at precisely 13:23:36 in order to initiate a new design. There's no change to the presses big enough to disrupt the use of the Lincoln Memorial Design between 13:23:35 and 13:23:36 but still small enough to make the imagined situation more similar to the actual situation than *any* change that involves the same design between those two periods of time.

We're making an analogous claim about normal instances of rationalizing explanation. In a normal instance where two attitudes jointly rationalize an action, and share a MOP, it's no accident that those attitudes share a MOP; that they share a MOP is the upshot of the way that the subject acquired the attitudes that rationalize the action. So we would have to depart further from actuality to find a world in which the same referential content is presented in an uncoordinated way than to find a world in which coordination is established via different MOPs.

As a first step towards defending this claim, we'll note that the connection between sameness of MOP and a certain kind of shared cognitive aetiology is built into many contemporary Fregean accounts. It's a recurring theme of post-Kripke theorizing about Fregean sense that senses are constitutively connected to characteristic modes of epistemic access to objects. See, for example, Gareth [Evans \(1982; 1985\)](#), John [Campbell \(2002\)](#), François [Recanati \(2012\)](#), and Imogen [Dickie \(2015\)](#), among others. According to such views, two attitudes share a sense when each is the upshot of a particular exercise of an epistemic capacity (to track an object in perception, say, or to use a name for the object).

But we don't need to assume the details of these theories to make our point. Our claim is weaker. We only have to hold that when an agent finds themselves in a situation where they have attitudes about *o* which rationalize a specific *o*-directed behaviour, it's typically not an accident that the attitudes are coordinated. Why not? Because it's crucial to the rationalizing character of these explanations that the deployment of the same MOP in two attitudes be sensitive to the agent's evidence that the two attitudes are coreferential. We make no claim here about the form this sensitivity might take—it might be sensitivity that's built into the deployment of certain tracking capacities, or that's reflected in explicit reasoning about identity, or some mix of these things, or perhaps something else. If this fact about redeployment of MOPs were *not* the norm, it would be mysterious why rational action tends to be successful: why it tends to achieve the agent's desires.

Suppose that the cognitive mechanisms that assigned MOPs to distinct attitudes were *not* normally sensitive to evidence of coreference. Then the fact that two attitudes were coordinated would no longer be a reliable indication that they were carrying information about the same object. And unreliability of this sort is hard to square with the observation that, normally, rational behaviour satisfies the desires that generate it; for part of the story about why rational behaviour normally succeeds is that coordination reliably guarantees that one's desires and instrumental beliefs are homing in on the same object. And if this sort of referential convergence were *unreliable*, then it would be a matter of luck that rational behaviour tends to satisfy the desires that generate it. But it's *not* a matter of luck that rational behaviour tends to succeed; indeed, it's constitutive of rationality that it not be reliant on luck in this way (Yablo 1997, p. 272). Fodor (1994) puts the point more succinctly: '[S]urely, no serious belief/desire psychology could treat the routine success of prudent behaviour as *accidental*' (p. 40, emphasis in the original).

So the matching of the MOPs deployed in rational action—like the matching designs of authentic pennies—is modally held in place by the process that generates it. Small-scale variations in the conditions under which the attitudes came about don't normally disrupt that sameness. We can see, then, how implausibly the Fregean must view the modal space around Case 1 to maintain that (F) is proportional. They must hold that the smallest changes that alter MOPs, but that leave referential contents fixed, also happen to be exactly those changes that would short-circuit the process that generated coordination.

We don't deny that such situations are *possible*. One can easily imagine them. Suppose the facts in Case 1 are as follows. The friend who told Fred that Clemens is vain did so because he flipped a coin and the coin landed heads; if it had landed tails, the friend would have informed Fred that *Twain* is vain. If the friend had done so, Fred's belief and his desire wouldn't have been coordinated. So there's a nearby world—one in which the coin came up tails—in which the MOPs in the belief and desire are distinct.

We only claim that such situations are not *normal*. The coordination of the attitudes upon which we act doesn't *normally* depend on bodies of attitudes which could have been cleaved apart, without change in referential content, had things been only slightly different. We have forced the Fregean into an *extremely* specific claim about the modal space around Case 1—a claim that isn't motivated by any other

part of her theory. With respect to the charge of special pleading, the shoe is now snugly on the other foot.²³

Our argument crucially relied on the claim that it's typically not an accident that attitudes of the sort mentioned in (R) are coordinated. Now, admittedly, this claim by itself is enough to cast doubt on the allegation that *Exceptionalism* is guilty of special pleading; for if it's not an accident that such attitudes are typically coordinated, then in the normal course of events, they will be. So cases in which they're *not* coordinated will be *abnormal*, and abnormal cases are unthreatening exceptions, not genuine counterexamples.²⁴ Consequently, the non-accidentality of coordination goes some distance on its own toward defending the truth of Russellian generalizations like (R). But, then, why do we bother with nomic proportionality at all? Isn't it just unnecessary complexity? No. It's not enough, for our purposes, to defend the *truth* of Russellian generalizations. Given a commitment to explanationism, defending a Russellian conception of content requires demonstrating that the corresponding generalizations are *explanatory* (lawlike). And this requires testing them for their relative proportionality. To put the point another way, it's the Yablovia argument that allows us to 'turn the tables' on the Fregean.

Let's pause and reassess. From the beginning, we advertised an argument that turns the tables on the standard challenge from Frege cases. The starting point of our argument is explanationism, which says that our most promising explanatory generalizations ought to guide us toward an adequate principle of content individuation. We've assumed, along with many others, that these generalizations take the form of *cp*-laws. The question, then, is whether the relevant *cp*-laws should be given a Russellian understanding, as in (R), or a Fregean understanding, as in (F). If the former, then Case 2 should be treated as an unthreatening exception in the manner that *Exceptionalism* dictates; if the latter, then Case 2 motivates a fine-graining of content. In this section, we've argued that a Fregean understanding of the relevant *cp*-law can't plausibly be made to satisfy an independently motivated constraint (nomic proportionality) that bears on the question at issue. So, as between a Russellian and a

²³ It's an interesting question—one that we won't explore—whether a *cp*-law analogous to those we considered above, but in which content is conceived even more coarsely (say, in terms of possible worlds), would nomically screen off (R₁).

²⁴ Thanks to a referee for pointing this out.

Fregean understanding of the relevant *cp*-laws, it looks as if Frege cases guide us toward the former.²⁵

All of this is compatible with the truth of some Fregean generalizations, just as the truth of (P)—‘Pennies conduct electricity’—is compatible with (E)—‘Copper conducts electricity’—being the relevant explanatory principle. Assuming there are MOPs, and that (F) is true, we want to understand the theoretical shortcoming of (F) by analogy with the theoretical shortcoming of (P). (P) is theoretically defective, to the extent that it is, because it’s not a law of materials science. Similarly, we’re suggesting that (F) is theoretically defective, to the extent that *it* is, because it’s not a law of folk psychology. By ‘law’ we just mean the sort of generalization in virtue of which certain ‘because’-statements are true. Assuming it’s true that this piece of metal conducts electricity, it’s true not because the item under discussion is a penny (though we may well suppose that it is a penny), but because it’s made of copper. These explanations are underwritten by (E). Thus (E) is the relevant law. Similarly, on the assumption that Fred flatters Twain/Clemens, he does so not because he has attitudes involving certain MOPs (though we may well admit that his attitudes do indeed involve such MOPs), but because he has attitudes with appropriate referential contents. In both cases, the failure of nomic proportionality tells us that the relevant generalization isn’t a law: it doesn’t identify the relevant explanatory structure. So, given explanationism, Frege cases provide us with reason, surprisingly, to prefer Russellianism over Fregeanism.

5. Fregeanism Again

In the previous section we argued that a proper treatment of Frege cases provides considerable support for Russellian Exceptionalism. The Fregean, not the Russellian, is in the unenviable position of having to make special allowances for Frege cases. We’d like to take a step

²⁵ As advertised, the upshot of our argument is about the nature of propositional attitudes. But the resulting position, Russellian Exceptionalism, does bear on how one ought to think about the semantics and pragmatics of attitude ascription. If our argument is sound, then a robust Fregeanism about attitude reports—that is, a view according to which the truth of an attitude ascription requires a certain Fregean content to be encoded in the ascriber’s attitude—is implausible. If, as we argue, attitudes don’t have Fregean content, then Fregeanism about attitude ascriptions would involve one in an error theory about them that isn’t very palatable. Having said that, Russellianism about attitude content is consistent with a wide range of views about the semantics and pragmatics of ascriptions. But because space is limited, and the issues are highly complicated, we won’t explore the matter here.

back, now, to situate these results in a broader understanding of the theoretical options—one that appreciates the full range of issues about cognitive significance. From this new perspective, we'll be able to see how Fregeans might bypass our argument.

Let's return to a crucial moment in our argument from nomic proportionality. We claimed that, in the context of responding to the challenge from Frege cases, the Russellian is entitled to assume the following: it makes no difference, with respect to the applicability of (F), whether the relevant belief and desire are both presented *via* CLEMENS or both presented *via* TWAIN. Generalizations like (F) treat pairs of referentially equivalent, isomorphically coordinated attitudes indistinguishably. This assumption was justified by the kind of challenge that Frege cases were supposed to pose for Russellianism. Returning to the assumption now, with our new background in place, we're in a position to characterize an important choice-point for Fregeans.

Recall that when we introduced (F) in response to Case 2, we noted that we were only motivated to *generalize over* MOPs; nothing beyond sensitivity to sameness/difference of MOP was motivated by Frege cases. If the fine-graining of content that Fregeans posit is motivated by Frege cases, then *all* Fregean generalizations will have this form. Call this position *Relational* Fregeanism. Relational Fregeanism holds that differences in the way referentially equivalent attitudes can participate in psychological explanation bottom out in the presence or absence of coordination between different attitudes.²⁶

To be a Fregean without being a Relational Fregean would be to hold that the way referential content is encoded in a particular attitude can have a 'direct' impact on its relevance to intentional explanation; that is, an impact that isn't mediated by coordination with other attitudes.²⁷ It would be to hold that there could be psychological generalizations whose applicability is sensitive to the difference, for example, between the pair CLEMENS-CLEMENS and the pair TWAIN-TWAIN. We'll look at two prominent ways in which

²⁶ An obvious question that arises here is the relation between Relational Fregeanism and Relationism. We leave discussion of this to Almotahari and Gray (manuscript).

²⁷ The significance of this distinction (between MOPs that have a 'direct' impact on cognitive significance and those whose impact is mediated by their effect on coordination) for the plausibility of Fregeanism is discussed in Aidan Gray (2020; manuscript).

this idea has been developed, though not *exactly* in these terms. Both ways of developing the idea are controversial, involving claims that go beyond the familiar case for Fregeanism that we rehearsed in §3, but each would enable the Fregean to bypass our argument in §4.

To begin, consider an example from John Perry (1977). Suppose while hiking in the woods Perry has a close encounter with an angry bear. Believing the bear is about to attack him, Perry rolls into a ball. We might try to explain Perry's behaviour by appealing to a Russellian generalization like

- (A) *Ceteris paribus*, if x has a belief with the referential content that x is about to be attacked by a bear, and x has a desire with the referential content that x not be injured, x will roll into a ball.

But now imagine that, due to a cleverly placed mirror in the forest, Perry sees a man about to be attacked by a bear. Unbeknownst to him, he's looking at his own reflection. He doesn't roll into a ball. Instead, he yells out a warning.

In the second situation, Perry is in a Frege case (with respect to himself) and he appears to be an exception to (A). But it seems less plausible that we can locate the reason why he fails to conform to (A) in facts about coordination. After all, we can assume that Perry doesn't want the man he sees to be injured. So he satisfies the antecedent of (A), and does so in virtue of the possession of coordinated attitudes. (He might even say to himself, 'I believe *that* man is about to be attacked by a bear and I don't want *him* to be injured'.)

We might, at this point, decide that we need to appeal to non-relational differences between the attitudes that Perry would express with 'I' and the attitudes he would express with 'that man'. There might be a substantive difference between thinking of Perry as 'that man' and thinking of Perry as 'I', and this substantive difference might be responsible for the behavioural difference between the two cases. In both cases, Perry has a belief with the referential content that Perry is about to be attacked by a bear. Ultimately, the difference between their explanatory role might *not*—if the present line of thought is correct—be explained in terms of the *other attitudes* with which the beliefs are coordinated in the two cases. This would motivate a fine-graining of intentional content in a way that would bypass the argument we gave above.

If we were led to revise (A) on the basis of this sort of consideration, we might come up with something like

- (A_{de se}) *Ceteris paribus*, if x has a belief with the referential content that x is about to be attacked by a bear in which x is presented via a first-person MOP, and x has a desire with the referential content that x not to be injured in which x is presented via a first-person MOP, x will roll into a ball.

Note that (A_{de se}) is importantly different from (F). (A_{de se}) doesn't simply require sameness of MOP across the two attitudes. It requires that the relevant MOP be of a certain type. (Of course, to make this idea work we'd have to offer a characterization of *de se* MOPs.) This is the crucial departure from Relational Fregeanism.

Essential indexicality isn't the only way of motivating a departure from Relational Fregeanism. Another live option is to affirm a broadly rationalist picture of the *a priori* and its relation to content—a picture that receives detailed defence in the work of Frank Jackson (1998) and David J. Chalmers (2011b). On this sort of view, differences in sense (or 'enriched intensions') aren't merely supposed to track patterns of coordination; they're also supposed to explain (or model) epistemically significant differences between referentially equivalent contents. For example, suppose it's knowable *a priori* for Fred that if something is the clear, colourless, odourless, drinkable liquid that fills the oceans, lakes, and rivers around here, then it's water. Suppose, further, that it's not knowable *a priori* for Fred that if something is the clear, colourless, odourless, drinkable liquid that . . . , then it's H₂O. These two conditional propositions differ in epistemic status for Fred; so they differ in cognitive significance. But, plausibly, this epistemic-cum-cognitive difference can't be explained by patterns of coordination, because we might suppose that each conditional proposition belongs to a distinct body of beliefs that is, with respect to the other body, referentially equivalent and isomorphically coordinated. These suppositions describe a cognitive difference without a corresponding difference in either reference or coordination. If their mutual satisfaction is nomically possible, then the psychological impact of sense would extend beyond coordination in a way that the laws of psychology should accommodate.

Call the kind of Fregeanism that posits differences in sense that are relevant to psychological explanation beyond the fact that they establish coordination *Substantive Fregeanism*.²⁸

We won't comment here on the plausibility of Substantive Fregeanism. Philosophical orthodoxies about essential indexicality have recently come under attack.²⁹ It might be that the explanation we offered of the behavioural difference between the two versions of Perry's case is incorrect. And perhaps 'water'-thoughts and corresponding 'H₂O'-thoughts don't differ in epistemic status. Maybe neither of the conditional propositions above are knowable *a priori* for Fred (Ned Block and Robert Stalnaker 1999; Derek Ball and Bryan Pickel 2014). Maybe both are (Nathan Salmon 1986). We only want to point out where the Fregean explanationist might look to argue for fine-grained psychological generalizations in a way that bypasses the argument in §4. Frege cases involve a lack of coordination. But, if you're a Substantive Fregean, what motivates fine-graining is visible in cognitive differences that don't rely on differences in coordination. So Frege cases as such are a distraction.

To see this, note that the presence of a Frege case was incidental to the problem that Perry's case posed for Russellianism. It's true that in the case as we described it, Perry was in a Frege case with respect to himself. (He had 'I'-thoughts and 'that man'-thoughts; they were co-referential but not coordinated; and they played different explanatory roles.) But we could have posed the same problem for Russellianism by imagining a passing hiker watching Perry's encounter with a bear. The passing hiker might have attitudes with the same referential content as Perry's by thinking of Perry as 'that man', and would then fail to roll into a ball (indeed, this is often how the case is imagined in the *de se* literature). If there's an explanatory difference between first-person and third-person MOPs, we needn't look at Frege cases to find it. The same is true for Jacksonian/Chalmersian rationalism. In fact, Jackson and Chalmers often motivate their picture by appealing to Gettier cases and Kripke's Gödel-Schmidt case (Chalmers and Jackson 2001).

This is all to say that if Fregeans want to use explanationism to motivate a fine-graining of content, they're wrong to focus on the role

²⁸ This position is, perhaps confusingly, just called 'Fregeanism' in Gray (manuscript).

²⁹ See Cappelen and Dever (2013) and Ofra Magidor (2015).

that sense plays in establishing coordination. Fregeans shouldn't care about Frege cases.³⁰

References

- Almotahari, Mahrad and Gray, Aidan. Manuscript. 'Coordination, *De Se* Thought, and Russellian Exceptionalism'.
- Ball, Derek and Pickel, Bryan. 2014. 'One Dogma of Millianism', *Philosophy and Phenomenological Research* 88: 70–92.
- Block, Ned. 1986. 'Advertisement for a Semantics for Psychology', *Midwest Studies in Philosophy* 10: 615–678.
- Block, Ned and Stalnaker, Robert. 1999. 'Conceptual Analysis, Dualism, and the Explanatory Gap', *Philosophical Review* 108: 1–46.
- Braun, David. 2000. 'Russellianism and Psychological Generalizations', *Noûs* 34: 203–236.
- . 2001. 'Russellianism and Explanation', *NOÛS* 35: 253–289.
- Campbell, John. 1987. 'Is Sense Transparent?', *Proceedings of the Aristotelian Society* 88: 273–292.
- . 2002. *Reference and Consciousness*. Clarendon.
- Cappelen, Herman, & Dever, Josh. 2013. *The Inessential Indexical: On the Philosophical Insignificance of Perspective and the First Person*. Oxford University Press.
- Chalmers, David J. 2002. 'On Sense and Intension', *Philosophical Perspectives* 16: 135–182.
- . 2011a. 'Propositions and Attitude Ascriptions: A Fregean Account', *NOÛS* 45: 595–639.
- . 2011b. *Constructing the World*. Oxford University Press.
- Chalmers, David J. and Jackson, Frank. 2001. 'Conceptual Analysis and Reductive Explanation', *Philosophical Review* 110: 315–361.
- Devitt, Michael. 1989. 'Against Direct Reference', *Midwest Studies in Philosophy* 14: 206–240.
- Dickie, Imogen. 2015. *Fixing Reference*. Oxford University Press.
- Evans, Gareth. 1982. *The Varieties of Reference*. Clarendon.
- . 1985. 'Understanding Demonstratives', *reprinted in his Collected Papers* (Oxford: Clarendon Press)
- Fine, Kit. 2007. *Semantic Relationism*. Wiley-Blackwell.

³⁰ The authors contributed equally to this paper. Comments from Rachel Goodman, David Hilbert, Eliot Michaelson, Bernhard Nickel, and a reviewer for this journal helped improve the discussion.

- Fodor, Jerry A. 1975. *The Language of Thought*. Harvard University Press.
- 1991. ‘You Can Fool Some of the People All of the Time, Everything Else Being Equal: Hedged Laws and Psychological Explanation’, *Mind* 100: 19–34.
- 1994. *The Elm and the Expert: Mentalese and its Semantics*. MIT Press.
- Godfrey-Smith, Peter. 2005. ‘Folk Psychology as a Model’, *Philosophers’ Imprint* 5.
- Gray, Aidan. 2020. ‘Indistinguishable Senses’, *NOÛS* 54: 78–104.
- Manuscript. ‘Minimal Fregeanism’.
- Heck, Richard Kimberly. 2012. ‘Solving Frege’s Puzzle’, *Journal of Philosophy*, 109: 132–174. (Originally published under the name ‘Richard G. Heck, Jr.’)
- 2014. ‘Intuition and the Substitution Argument’, *Analytic Philosophy* 55: 1–30. (Originally published under the name ‘Richard G. Heck, Jr.’)
- Jackson, Frank. 1998. *From Metaphysics to Ethics*. Oxford University Press.
- Kripke, Saul. 1979. ‘A Puzzle about Belief’, in A. Margalit (ed.), *Meaning and Use* (D. Reidel)
- Loar, Brian. 1988. ‘Social Content and Psychological Content’, in R. H. Grimm and D.D. Merrill (eds.), *Contents of Thought*. Arizona University Press.
- Magidor, Ofra. 2015. ‘The Myth of the De Se’, *Philosophical Perspectives* 29: 249–283.
- Nickel, Bernhard. 2010. ‘*Ceteris Paribus* Laws: Generics and Natural Kinds’, *Philosophers’ Imprint* 10.
- Perry, John. 1977. ‘Frege on Demonstratives’, *Philosophical Review* 86: 474–497.
- Récanati, François. 2012. *Mental Files*. Oxford University Press.
- Salmon, Nathan. 1986. *Frege’s Puzzle*. Ridgeview Publishing Company.
- Schneider, Susan. 2005. ‘Direct Reference, Psychological Explanation, and Frege Cases’, *Mind & Language* 20: 423–447.
- Segal, Gabriel. 2000. *A Slim Book about Narrow Content*. MIT Press.
- Soames, Scott. 1990. ‘Belief and Mental Representation’, in P. P. Hanson (ed.), *Information, Language and Cognition*. University of British Columbia Press.
- 2015. *Rethinking Language, Mind, and Meaning*. Princeton University Press.

- Yablo, Stephen. 1992. 'Mental Causation', *Philosophical Review* 101: 245–280.
- 1997. 'Wide Causation', *Philosophical Perspectives* 11: 251–281.
- 2014. *Aboutness*. Princeton University Press.