Referring When Push-Comes-to-Shove[1]

Kevan Edwards
Syracuse University

1. Direct Reference, Neo-Fregeanism, and the Threat of Collapse

The anchoring focus of this paper is a cluster of complaints that have been raised against reference-based approaches to semantics, in particular against the view defended by Scott Soames (2002). I am going to lump the complaints that I have in mind under the heading of the *Threat of Collapse* (or the *Threat*, for short). At the heart of the Threat of Collapse is the accusation that various moves referentialists make in dealing with well-known problems end up undercutting the motivations for a reference-based semantics in the first place.

I am going to propose a way of disarming the Threat. The proposed solution will involve stepping back and saying something about features of a broader framework in which (I believe) a reference-based semantics ought to be embedded. In effect, I am going to use the Threat of Collapse as a jumping off point—an excuse, really—for making some suggestions about how the Direct Reference theorist ought to construe the relationship between language and thought.[2] Thus, my aim is two-fold: (i) to make plausible the claim that the Direct Reference theorist has more than sufficient resources to avoid the Threat of Collapse, (ii) to lay some suggestive foundations for future work that embeds Direct Reference in a, roughly speaking, *pragmatist* conception of the relationship between language and thought.[3]

Here is the plan for the paper: The rest of this first section will introduce some features of the dialectical geography, including an abstract outline of how the Referentialist can hope to steer clear of the Threat. Section 2 will gesture towards a way of thinking about the relationship between language and thought. Section 3 will build on the previous section, focussing on a more specific proposal regarding the nature and theoretical motivations for a notion of *literal meaning*, or, more precisely, *literal usage*. Section 4 will return to a specific case of the Threat of Collapse, and show how it can be resolved.

─────────────────

[1] It will be obvious in what follows that I owe an intellectual debt to Stephen Neale and Jerry Fodor. I would also like to thank colleagues at Syracuse University for conversation and helpful feedback, in particular Kris McDaniel, Ben Bradley, Melissa Frankel, and Kara Richardson.

[2] Here and throughout, I use the term 'thought' to pick out a generic class (or classes) of psychological states. Hence my use contrasts with that of, for example, Frege, who uses the term to pick out abstract platonic entities.

[3] It is worth mentioning a point that will become clear in what follows, namely the extent to which the 'pragmatism' that I will advocate is carefully delimited. Not only will the approach be invoked in the service of a robust reference-based approach to semantics, but it will be consistent with a strong *anti*-pragmatism about the nature of the mental.

An important caveat: In recent years, there has been an explosion of work on issues at the intersection of philosophy of mind and philosophy of language. However, where possible, I am going to side-step potentially distracting details and focus on issues that are germane to the specific topic at hand. I beg the reader's forgiveness for various cut corners and ignored details.

Direct Reference

As indicated in the opening line of the paper, Scott Soames is going to be my chief protagonist. Soames is an arch defender of a *Millian* or *Direct Reference* account of the semantics of proper names, demonstratives, and more recently (ms) natural kind terms. Briefly, I'll consider an expression to be *Directly Referential* iff its semantic content is exhausted by reference.[4] Not surprisingly, Direct Reference finds its roots in the work of J.S. Mill (1858) and, arguably, Russell (1905).[5]

The considerations that have done the most to pave the way for Direct Reference are a cluster of so-called anti-descriptivist arguments typically credited to Kripke (1980).[6] A standard taxonomy divides these arguments into three categories: the *semantic*, *epistemic*, and *modal* arguments.[7] Very roughly, the heart of these arguments is a cluster of intuitions to the effect that the semantic, epistemic and modal profiles of (e.g.) proper names fail to pattern after those of descriptions. To give a bit more detail, the intuitions suggest that (i) there are no particular descriptions or clusters of descriptions that a speaker *must* associate with familiar referring expressions in order to count as a competent user of those expressions; (ii) various statements that a descriptivist semantics is

---

[4] I am overlooking details here that have the potential to make a difference. For example, one might distinguish between an expression's semantic *content* and its semantic *contribution* to larger expressions of which it is a proper part (See Crimmins, 1992). Worse still, once one distinguishes between an expression's semantic content and its semantic contribution, room opens up for accepting (for example) context-dependence vis-à-vis the latter while denying it vis-à-vis the former. I myself doubt there is much to be gained by splitting this particular hair, but my thinking so turns on issues—for example, the pointlessness of preserving the letter of Direct Reference by giving up on a strict notion of compositionality—that I will not argue for here. At any rate, Caplan (2007), who is going to be my main antagonist, does not distinguish semantic content from semantic contribution; indeed, he treats the semantic content of a proper name to be its contribution to propositions expressed by sentences of which it is a constituent.

[5] Russell is more obviously a Millian when viewed in abstraction from (i) his particular views about how much underlying structure is masked by the surface grammar of natural language, (ii) various epistemological commitments that serve as restrictions on the class of genuine referring expressions (e.g. the commitment to a principle of acquaintance).

[6] Also see (at least!) Donnellan (1997) and Kaplan (1989a, 1989b). Davidson (2007) is a recent and very helpful collection of seminal contributions to the Direct Reference tradition.

[7] See, for example, Caplan (2007), Soames (2002), and Salmon (2005).

apt to classify as *necessary* are in fact *contingent,* and (iii) various statements that a descriptivist semantics is apt to classify as *a priori* are in fact *a posteriori.* For example, "was the teacher of Alexander" is a plausible candidate for (part of) the descriptive content of "Aristotle", but "Aristotle was the teacher of Alexander" is neither necessary nor knowable a priori, and someone can be a perfectly competent user of "Aristotle" while failing to associate this particular piece of descriptive information with the bearer of the name. Call the intuitions underlying these arguments the *Kripkean Intuitions*.

Kripke's conclusion in *Naming and Necessity* is that names are *rigid designators*; that is, that a name refers to the same entity across possible worlds (or, at least, across worlds in which the entity exists). The claim that prima facie referring expressions are rigid falls short of a full commitment to Direct Reference, since rigidity can be squeezed out of the right kind of complex semantic structure. For example, descriptions can be *rigidified* by an 'actually' operator or by (somehow) being forced to take wide-scope over modal operators. Enter so-called *Direct Reference* theorists, such as Braun (2006, 2005, 2008), Kaplan (1989), Salmon (1986, 1989b, 1989a), Soames (1997, 2002), and Thau (2002). Soames (2002), in particular, argues at length and in detail that Kripke-style arguments can be extended to support a full-blown Direct Reference thesis.

Direct Reference theorists have done much to swing the tide in favor of reference-based approaches to semantics. Nevertheless, it is important to acknowledge that the standard Kripkean arguments, extended or otherwise, point not to positive features of the view, but to the apparent failure of the alternatives.[8] Moreover, the alternatives—approaches that dominated the fashion trends in philosophy of language in the interval between (roughly) Mill and Kripke—are not without their own motivations. At the core of these motivations is none other than the accusation that a significant class of intuitions militate *against* a reference-based semantics. Here, then, is the first hint of a brewing tension.

Neo-Fregeanism

For the purposes of setting up a simplified geography of competing views, call an approach *Neo-Fregean* if it contrasts with Direct Reference by claiming that *something* other than reference is part of the semantic content of the relevant class of expressions.[9] The most obvious motivations for Neo-Fregeanism stem from the uncontroversially impoverished nature of reference. Famously, while sentences of the form **a = a** are trivially true, uninformative, and knowable a

---

[8] Sider & Braun make this point. They say, "The data of Kripke et al. is genuinely puzzling. It in no way *undermines* the old Fregean arguments against Millianism; it simply is new, conflicting data." (2006, p. 669, original italics).

[9] The suffix 'Neo' indicates that Neo-Fregeans need not embrace Frege's anti-psychologism. Unlike Frege, a Neo-Fregean can, and I think should, claim that 'modes of presentation'—or whatever else are called on to do some or all of the work for which Frege posited *senses*—are most plausibly understood in terms of some or another psychological entities.

priori, sentences of the form **a** = **b** have the potential to be non-trivial, informative, cognitively significant, and only knowable a posteriori, including cases where **a**, **b** are replaced with co-referring names.[10] Things get even worse when one considers well known intensional contexts, notably belief ascriptions. Again, in the right context, the intuition can be very powerful that someone believes *that Hesperus is a planet* while also believing *that Phosphorus is not a planet*.[11] Arguably even worse, uncorrupted intuition suggests that someone can believe *that Hesperus is a planet* while simultaneously failing to believe *that Phosphorus is a planet*.[12]

_____

[10] The famous source of this 'puzzle' is Frege (1952). A few details: First, following a convention employed by Soames in various places, I use bold-face font in place of corner quotes. Second, I should be explicit that I am merely gesturing towards a standard presentation of Frege's Puzzle, and glossing over various subtleties. One subtlety has to do with the potential non-existence of the referent of (substitution instances of) **a** and **b**. A deeper issue turns on the relevant notion of 'form' employed in the statement of the puzzle. In the right context, someone can wonder whether "Fred = Fred" (e.g. when they are worried they might be occurrences of names of two different people), or even whether "that is that" (e.g. when pointing to the same thing twice without realizing it). It isn't obvious what notion of form can divide these cases from similar cases that are uncontroversially trivial. One might respond that slipperiness in the relevant notion of 'form' only makes things more difficult for the referentialist; but this isn't right. Insofar as the Neo-Fregean draws on instances of Frege's Puzzle to criticize the alternatives, they need to ensure that their own view does not fall prey to similar problems. Insofar as the alternatives to Millianism do face similar problems, this invites the conclusion that a *semantic* solution to Frege's Puzzle was never on the cards and that (Neo-)Fregeans were wrong to wield the problem against Millians. I take it this is the moral that Kripke (1979) draws from the famous case of Puzzling Pierre. (I should note, however, that Kripke's conclusions have provoked varying responses. Even among those broadly sympathetic to Millianism there is substantive disagreement; for example, McKay (1981) argues for a straightforwardly Millian resolution to the puzzle, whereas Loar (1987) suggests that the puzzle motivates an *impure* version of Millianism.)

[11] I assume a standard story where "Hesperus" and "Phosphorus" are two names for the same heavenly body, namely Venus, even though competent speakers (and belief attributers) may not realize they co-refer. Sometimes I will use italics to highlight that-clauses, in particular when they are used in belief attributions. One reason for doing this is that it serves as a reminder that belief reports being reported *by the theorist* occupy a potentially uneasy middle ground between pre- and post-theoretic usage. For example, a committed Millian might offer a *pragmatic* explanation for why speakers utter various sentences that *semantically express* contradictory propositions. If so, the Millian is as free as anyone to utter such sentences in standard conversational contexts, but they might be reluctant to do so in contexts where attention is focussed on questions about literal semantic content.

[12] "Arguably even worse" because according to the Direct Reference theorist the last intuition involves not just a report of inconsistent beliefs but a contradiction in the mouth of the belief reporter. See Kripke (1979).

Advocates of Direct Reference have some explaining to do. To summarize, they need to explain at least the following: (i) direct intuitions about the 'meaning' of various prima facie referring expressions; (ii) intuitions about the potential cognitive significance (informativeness, etc.) of expressions of the form **a** = **b** (for co-referring instances of **a**, **b**); and (iii) intuitive failures of substitution (in particular, when singular terms are found under the scope of a belief operator). Call these *Fregean Intuitions*. Fregean Intuitions play into the hands of Neo-Fregeans.

There is a lot to say—and a lot has been said—about how the Direct Reference theorist can deal with various Fregean Intuitions. One idea that Soames (e.g. 2002, 2005) explores in some depth is that relevant descriptive information can be *asserted* by a speaker in uttering a sentence—for example as a result of a process of pragmatic enrichment—without it being part of the semantic content of the sentence. Moreover, even if one is suspicious of Soames' specific appeals, the general strategy of preserving an austere, Directly Referential semantics by appealing to some or another pragmatic resources appears to have a lot of potential.

The Threat of Collapse

So far so good for the Direct Reference theorist. Drawing a line between semantics and pragmatics promises to yield resources for explaining various hard cases in a way that is consistent with a referentialist semantics. Enter the Threat of Collapse. In a nutshell, advocates of (what I am calling) the Threat of Collapse claim that appeals to pragmatically imparted descriptive information are fated to conflict with Kripkean Intuitions. At this point, somewhat different conclusions are drawn by different advocates of the Threat. Braun & Sider (2006) and McKinsey (2005) criticize Soames' appeals to pragmatically conveyed descriptive information, but they do not claim that this undermines a basic commitment to a Millian semantics. In what follows, I am going to focus on Caplan's (2007) version of the Threat.[13] While Caplan doesn't put it in exactly these terms, the line of thought he pursues amounts to the following: First, everyone in the Neo-Fregean versus Millian debate needs to account for both Kripkean Intuitions and Fregean Intuitions. Second, each side has a semantic account that is tailor-made to handle one set of these intuitions but faces prima facie problems explaining the other. Thus, nothing in these two patterns of explanatory successes and failures privileges Millian Descriptivism over a (Neo-)Fregean alternative. Caplan says the following:

> *Even if the explanation that Millian Descriptivists offer of speakers' intuitions about some sentence appeals to the content of that sentence, Fregean Descriptivists can offer a parallel explanation, by saying that speakers*

_____

[13] Why am I focusing on Caplan? First, I find Soames' own response to Sider & Braun and McKinsey close to definitive and don't want to repeat what has been said elsewhere. Second, and more importantly, Caplan's discussion is much more in the spirit of unearthing a more fundamental tension in reference-based approaches, and I think unravelling the strands contributing to this tension is both important and illuminating.

> *communicate the proposition that Millian Descriptivists say is the content of that sentence.* As a result, Millian Descriptivism fares no better (or worse) than Fregean Descriptivism. (2007, p. 194, original italics)

Caplan's conclusion is cautious. Given Caplan's basic line of argument, one might go a step further and conclude that the situation facing the Direct Reference theorist is in fact *worse* than a stalemate. Arguably, the Neo-Fregean can claim that their semantics does a better job of according with pre-theoretic intuitions. I myself think the Direct Reference theorist ought to concede this last point. Fregean Intuitions are very easy to generate, whereas Kripkean Intuitions can take more work to probe—it isn't any surprise, I think, that descriptivism has seemed so attractive, and for so long. Nevertheless, for reasons that will become clear later in this paper, I also think that this difference in the immediacy of the respective intuitions is to be expected once Direct Reference has been embedded in a larger theoretical framework. Therefore, I do not think these considerations should bear any weight in the final analysis. This, however, is to jump ahead. For now the goal is to clarify the potential force of the Threat of Collapse. The basic worry, to repeat, is that the Neo-Fregean can hope to mimic moves made by the Direct Reference theorist, and they can do so against the backdrop of a semantics that claims better to accord with pre-theoretic intuitions.

A digression about descriptivism

Advocates of the Threat of Collapse, at least as I have described them, attack Millian *Descriptivism*. But one might wonder if the only way the Direct Reference theorist can hope to appeal to pragmatics is in terms of pragmatically conveyed *descriptive* information.

Speaking very generally, it is natural for the Direct Reference theorist to conceive of pragmatics as a potentially very messy business. And one might expect this messiness to leave open a broad range of options when it comes to information potentially conveyed by pragmatic mechanisms. A related point that I think is under-appreciated is that it isn't obvious that the Direct Reference theorist needs to come up with a fully general account of the source of Fregean Intuitions. Arguably, the Direct Reference theorist just needs to be able to invoke a relevant difference in the information conveyed in each *particular* problem case. Notice, for example, that different people can have different reasons for believing that "Hesperus" and "Phosphorus" fail to co-refer. One person might think this because they tend to apply the names at different times of the day; another because they were told as much by someone they believe to be reliable; yet another because of a general presumption that orthographically distinct names fail to co-refer. In itself, this suggestion about a lack of generality does not shed light on relevant pragmatic resources, but it does open the door for a potentially broad range of resources to be helpful. Soames himself registers doubts about the scope of the appeals to descriptive enrichment that he discusses in *Beyond Rigidity*. In his response to McKinsey, for example, Soames says the following:

> One might reasonably ask how often this sort of thing [an appeal to descriptive enrichment] occurs, and how much of a role it plays in understanding Frege's

Puzzle. These are legitimate questions about the scope and significance of the kind of descriptive enrichment I emphasize. My point is minimal; such enrichment does occur, and it plays some role in resolving Frege's Puzzle. (2005, p. 173)

In a related context, Soames gestures towards an appeal to different 'ways' of entertaining a proposition (also sometimes described as propositional 'guises'). This appeal is worth mentioning insofar as it is familiar among advocates of Direct Reference, and it isn't obvious that 'ways' or 'guises' need to be understood descriptively. (For related discussion, see Crimmins (1992), Thau (2002)).

In spite of Soames' acknowledgement that there might be alternatives to descriptive enrichment, it is striking, and to some extent suspicious, that Soames does little to explore such alternatives. More importantly, it is not obvious that the basic logic of the Threat of Collapse won't resurface in a different *guise* (sorry, I can't resist the pun). As I have described it, the core problem behind the Threat of Collapse is the appearance of a fundamental tension between the austerity of the semantic resources that best account for Kripkean Intuitions and the richness of the resources required to explain Fregean Intuitions. The Threat of Collapse appears to present an interesting and potentially robust problem for the referentialist.

The two-source solution

Simply put, I think that Kripkean Intuitions and Fregean Intuitions stem from two different sources, and that the appearance of tension between the two sets of intuitions is merely an appearance.[14] On the one hand, Kripkean Intuitions stem from the semantic content of the expressions in question; indeed, I'll sketch an account according to which these intuitions stem from the *literal meaning* of the expressions in question. On the other hand, I will suggest that Fregean Intuitions have their source in potentially much richer facts about what the speaker of an expression intends to express/communicate in uttering the expressions in question.

The basic structure of the two-source solution is obvious enough; what one wants is some reason to think that it can be implemented in a way that is both plausible and consistent with Direct Reference. In short, the Direct Reference theorist needs a larger framework that leaves obvious room for a relatively profligate view of pragmatics and yet carves out an important role for an austere, directly referential semantics. I'm going to rough out such a framework in the next two sections. In section 2, I will help myself to a broadly speaking *pragmatist* or *use-theoretic*

---

[14] Note that explaining away the appearance of conflict is not, on its own, enough to dissolve the force of the Threat of Collapse. The Direct Reference theorist needs not only to distinguish two sources for the intuitions in question, but to say something to privilege the source of the Kripkean Intuitions as deserving the title of 'semantics'. Although I won't explicitly focus on this second issue in what follows, an appropriate asymmetry between the two sources of intuitions will be an obvious feature of the diagnosis that I will offer.

conception of natural language. Much of this I am going to borrow from recent work by Stephen Neale (ms, 2004a, 2004b). Neale articulates a "collection" of theses that he groups together under the heading of *Linguistic Pragmatism*.[15] In what follows, I will focus on several ideas that are especially relevant for present purposes, and add a few details of my own. I am going to diverge from Neale on one crucial issue. The notion of literal meaning that I am going to introduce is different in emphasis, if not in substance, from Neale's own appeals to what he calls *blueprints*. This point of divergence is important enough to warrant its own section (§3).

2. Linguistic Pragmatism

The core idea behind Linguistic Pragmatism (henceforth "LP") is that natural language is best understood as a tool cooperatively exploited by speakers and hearers (mutatis mutandis for writers and readers) for the purposes of expressing/communicating meaningful thoughts. Thus, LP involves a strong claim about the semantic priority of thought over language. A point of central importance to LP is the idea that the primary explanatory notions at work in understanding language are not what a *sentence* means (or says, or whatever) but what a *speaker* uses a sentence to say/communicate/express/etc. I think Neale gets these issues almost exactly right, and I will quote him at length.

> Saying and implying are things *people* do. Following ordinary usage, *the speaker* is taken to be the understood subject, so to speak, of the verbs 'say' and 'imply', the verbs in talk about 'what is said' and 'what is implied.' (Similarly, with verbs such as 'communicate', 'convey', and 'get across'.) We should be initially suspicious of talk about what *uses of sentences* say (imply, communicate, etc.) and talk about what *sentences-relative-to-contexts* say (imply, communicate, etc.), unless such talk is taken to be straightforwardly translatable into talk about things that speakers are doing. […] At the same time, we should be open to the idea that new, *technical* uses of the verbs 'say' ('imply', 'communicate', etc.) may need to be defined, or at least developed, in the course of our inquiries, such uses earning their keep because of ineliminable theoretical work they do. (2004a, p. 183, original italics)

One minor point of possible contention is the extent to which Neale thinks that pre-theoretic intuition ("Following ordinary usage…") respects the priority of speaker meaning. I have my doubts about this. For present purposes, however, the issue is moot; even if the primacy of speaker meaning is not pre-theoretically intuitive or countenanced by 'ordinary usage', it *is* well

_____

[15] Citing Neale is convenient insofar as he brings together various theses under a common heading. However, Neale himself (2004a, p. 178) is quick to cite a long list of people who are "operating in a broadly pragmatist spirit", including (but not limited to): Austin, Bezuidenhout, Carston, Chomsky, Crimmins, Evans, Heal, Fodor, Grice, Quine, Perry, Récanati, Searle, Sellars, Sperber & Wilson, and (Peter) Strawson.

motivated by the extent to which it is a core feature of an imminently plausible account of natural language use.

Thus far I have been steering clear of any need to distinguish between what *is* communicated and what a speaker *intends* to communicate. As a matter of fact, I assume—and here I am still following Neale—that notions such as 'what is said' or 'what is communicated' are not primary in the explanatory order. Again, I'll quote Neale at length:

> In the first instance, we should separate (i) *what A intended to say by uttering X on a given occasion*, and (ii) *what a rational, reasonably well-informed interpreter in B's shoes would think A intended to say by uttering X on that occasion* (which is not to say there are not problems with the idea of *a rational, reasonably well-informed interpreter in B's shoes*). In cases where (i) = (ii), we can talk freely about *what the speaker said*. In cases where (i) ≠ (ii), certainly we *could* argue about which of (i) or (ii) or some third thing has the 'right' to be called *what is said*, but what would be the point? First, what third thing distinct from (i) and (ii) could be of any significance to a theory of interpretation? There is simply no rôle for a transcendent notion of what is said upon which (i) and (ii) converge when all goes well. It is the *coincidence itself* of (i) and (ii) that constitutes success, and it is the *potential* for such coincidence, independently of some third thing, that *gives sense* to the very idea of saying. (Contrary to linguistic appearances the concept of *intending* to say is, in fact, more basic than the seemingly simpler concept of saying.)   (2004a, p. 182, original italics)

To summarize: any notion in the vicinity of 'what is communicated' or 'what is said' is at best something to be derived from the dual notions of (i) what a speaker *intends* to say/communicate/ etc. and (ii) what a speaker reasonably can believe about their audience's ability to recover these intentions. Importantly, these notions come together in at least two ways: First, when communication is successful, a speaker's intentions line up with what their audience can and does recover. Second, speakers are constrained in their choices of utterances by their beliefs about how speakers will understand various sentences they *could* utter. It is here that linguistic conventions enter the LP picture: speakers and hearers must cooperatively take themselves to be operating with a shared set of conventions about how language can be used.[16]

---

[16] There is little reason to think that a complete account of linguistic interpretation will be anything other than extremely complex, if not theoretically intractable. Neale describes at least two pragmatists, Chomsky and Fodor, as deep pessimists who "have argued that asking for a theory of interpretation is tantamount to asking for a 'theory of everything', a complete cognitive psychology, because virtually anything can impinge upon the holistic process of interpretation." (2004a, p. 178 n. 25).

The next question to ask—given the present topic of the source of Kripkean and Fregean Intuitions—is how intuitive data fits into the LP framework. Neale (ms) acknowledges the following:

> As the pragmatist sees things, our intuitive judgments about what *A* meant, said and implied, and judgments about whether what *A* said was true or false in specified situations constitute the primary data for a theory of interpretation, the data it is the theory's business to explain.

One needs to be careful in interpreting what Neale is asserting here. Two important points: First, remember Neale's insistence that the primary explanatory target is not what *sentences* mean, but what *speakers* mean in uttering these sentences. Second, in passages like the above, Neale does not intend to rule out notions along the lines of 'what is said' or, for that matter, the *literal meaning* of a sentence. Rather, I take Neale to be committed to the following two claims about the priority of speaker-meaning: (i) Any notion that goes beyond what speakers mean, for example by assigning meaning to sentences (or to their constituents), needs to earn its keep by making a substantial—if not practically ineliminable—explanatory contribution. (ii) Any notion that goes beyond what speakers mean needs to be explained in terms of—or as I prefer to put it, *extracted from*—the basic use-theoretic framework at the heart of LP.

So, intuitive judgments are a (if not *the*) primary source of data for the Linguistic Pragmatist; however, within the LP framework, there is the potential for (at least) two sources of intuitive judgments. The core data for a theory of interpretation is judgments properly treated as intuitions about what *speakers* mean (/say/imply/etc.); but the approach also leaves room for more constrained notions of, for example, the *literal meaning* of a sentence, as long as such a notion can both earn its keep and be derived from within a use-theoretic framework.

Before saying more on the topic of literal meaning, I want to discuss several potentially distracting details. Some of these details I mention only for the purpose of setting them aside; however, I also want to introduce a few substantive assumptions that I will take for granted in what follows.

Assertion, what is said, and literal meaning. Let's go back, for a moment, to Soames' notion of descriptive enrichment. Recall that Soames' appeals to descriptive enrichment, as well as his subsequent responses to McKinsey (2005) and Sider & Braun (2006) are couched in terms of information that is not part of the semantic content of a sentence but nevertheless is *asserted* by a speaker (in uttering the sentence in a particular context). Soames also makes similar claims about *what is said* by a speaker in uttering a sentence. These are potentially controversial claims. Indeed, I think it does much to cloud the discussion between Soames and McKinsey that they disagree about whether descriptive enrichment is a fact about what is asserted (mutatis mutandis for 'what is said') or if, as McKinsey (2005, p. 166) puts it at one point, "the descriptive assumptions that underlie the use of proper names typically remain in the background, and play at most a reference-fixing role". This specific debate is not one into which I wish to enter, nor are

various other points of contention surrounding different accounts of 'what is said' and related notions.[17] I am going to side-step these issues by jumping right to a relatively simple-minded contrast between the overall information that speakers intend to communicate in uttering a sentence and a potentially much more constrained conception of the *literal meaning* of a sentence. Sketching how a highly constrained conception of literal meaning can be extracted from the basic LP framework is the burden of the next section of the paper. Just to keep track of the point of all this, recall that the aforementioned distinction is intended to serve as an implementation of the two-source solution. In particular, I want to claim that Kripkean Intuitions plausibly are construed as intuitions about the *literal meaning* of a sentence, whereas Fregean Intuitions have their source in something potentially much richer.

Communication/expression. Language is used not only for communication but also in cases where there isn't an intended audience—consider a frustrated parent who blurts out an obscenity. The basic spirit of LP is focused on communication, and it will be easier to take this for granted in what follows (and thereby drop the annoying use of "communicate/express"). Simplifying the discussion is one motivation here, but I also think the case can be made that expressive uses of language are parasitic on communicative uses. For example, it isn't obvious how to make sense of the use of an obscenity as an *expression* of frustration unless the same expression could be used in a relevantly similar context to *communicate* the same content.[18]

Language and thought; words and concepts. A pragmatist or use-theoretic approach to natural language need not entail a pragmatist approach to the semantic content of *thought*. To the contrary, the basic outline of LP is consistent with an *anti*-pragmatist conception of the content of psychological states.[19] Although it will not play a central role in what follows, I myself am inclined to assume that thoughts have a roughly language-like structure, at least to the extent that whole proposition-sized thoughts break down into robust (repeatable and, for the most part, context-independent) constituents, or *concepts*. Of course, this is to commit to the heart of the so-called *Language of Thought hypothesis* (Fodor, 1981, 1975, 2008). I am also inclined to accept, at least as a working hypothesis, that words have semantic properties in virtue of their bearing a relation to corresponding concepts—more along these lines in a moment. Finally, I am inclined to follow the suggestion outlined above to the effect that conceptual content should receive an

---

[17] See, as a starting point, Bach (2001), Carston (2002), Recanati (2001) .

[18] For all I know, expression might be best understood as a kind of communication with oneself.

[19] One indication that LP is consistent with a *non*-pragmatist conception of thought is the fact that Neale describes Fodor as an ally when it comes to LP. Fodor is perhaps *the* arch *anti*-pragmatist when it comes to mental content: In his recent *LOT-2,* Fodor writes, "In my view, pragmatism is perhaps the worst idea philosophy ever had" (2008, p. 9). Fodor's target here is pragmatist conceptions of the nature of intentional states; in particular, the intentional properties of these states.

independent, non-pragmatist treatment.[20] I should add that while Neale flirts with these kinds of assumptions in various places, I can't with a clear conscience burden him with a firm commitment to either the Language of Thought or the plausibility of a non-pragmatist approach to conceptual content.

3. Literal meaning as what is communicated when push-comes-to-shove

Here is where I take us to be: According to LP, the primary explanatory notions underwriting claims about the semantics and pragmatics of natural language have to do with facts about how speakers use sentences to communicate thoughts. Of particular relevance is the nature of various conventions constraining the nature of the thoughts that a sentence can be used to communicate. In particular, then, anything in the vicinity of a claim about the *literal meaning* of a sentence needs to bottom out in such facts. I'm going to assume—and I don't think this much is particularly tendentious—that once such facts are spelled out, they can be treated as defining a relation that holds between each sentence (type) and the thought (or set of thoughts) that grounds claims about the literal meaning of the sentence. For the sake of introducing a relatively neutral term, I will call this relation the *backing* relation. As I indicated above, I am also inclined to accept, at least as a working assumption, that claims about the literal meanings of words are underwritten by the fact that they bear the *backing* relation to corresponding concepts.[21] One detail: I assume it is a mistake to describe the backing relation as holding between a word/ sentence and its *meaning*. The backing relation is more aptly described as the relation whereby natural language expressions *inherit* their meanings from their corresponding concepts.[22]

Thinking about the notion of literal meaning in terms of a relation is very natural within the LP framework; nevertheless, one might prefer to run the discussion in terms of some kind of *entity* that either defines or is defined by such a relation. This is more or less what I take Neale to be

---

[20] Admittedly, questions about how this story might go are fraught and I cannot hope to do justice to them here. For what it is worth, here is a quick summary of my own views on the subject: (1) In the context of a presumptive pragmatist reduction of linguistic content to mental content, questions about the potential naturalization of mental content become all the more pressing. (2) The basic geography of naturalistically-respectable accounts of mental content comes down to (i) one or another version of Conceptual Role Semantics (e.g. those articulated by Block (1986), Harman (1982), Loar (1981), Peacocke (1992)) and (ii) one or another theory of representation (e.g. those articulated by Dretske (1981), Fodor (1992), Millikan (1993)). (3) CRS faces a list of deep problems, including various difficulties catalogued by Fodor (1998, 2004) and Fodor and Lepore (2002), as well as analogues of the Kripkean Intuitions (Edwards, ms).

[21] One reason to think this might be an oversimplification is the potential for sentences in natural language to be *backed* by thoughts containing constituents that do not correspond to anything 'articulated' in the surface structure of the sentence.

[22] See note 20 for more about conceptual content.

doing in his discussion of what he calls *blueprints*. I should note, however, that insofar as I want to offer an alternative to Neale's blueprints, the relation-vs.-entity issue is not what I take to be at issue.

Here is how the notion of a blueprint arises for Neale:

> Qua description of semantic competence, a semantic theory for a language will explain how the syntactic structure of a sentence (or sentence fragment) X and the meanings of the individual words in X conspire to constrain what speakers can *say* using X. Flushing out the modal: a semantic theory for a language L will provide, for each sentence X of L, a *blueprint* for (a *template*, a *schematic* or *skeletal* representation of) what someone will be taken to be saying when using X to say something. The blueprint associated with X is its *semantics*, and the set of such blueprints, one for every sentence of a language L, is the *semantics* for L. (2004a, p. 189)

It is important to note that Neale is careful about the theoretical burden that blueprints can and should be asked to bear in an overall account of linguistic interpretation. It is part and parcel of Neale's approach that blueprints leave a lot of slack to be picked up by other theoretical resources. I think Neale is right to limit the role of anything in the vicinity of a notion of literal meaning; however, I'm worried that he doesn't go far enough. In particular, there seems to be little to prevent blueprints themselves, at least as Neale construes them, from being infected by the obvious messiness of language use. To explain what I mean here, I need to take a step back and make some claims about plausible constraints on any workable notion of literal meaning.

I take it, and I take it that Neale would agree with this, that a substantial part of the explanatory bona fides of a notion of literal meaning is that it promises to ground linguistic communication in something relatively stable across agents and contexts. The need for a stable core is especially pressing in the context of a commitment to the basic LP framework insofar as it is part and parcel of LP that language use can be very plastic, context-dependent, creative—in short: *messy*. The tool metaphor is instructive here: Once a tool is available, it can be used in no end of creative ways, uses that can far outstrip its original raison d'être.

In spite of the obvious messiness of language use, there is substantial pressure in the opposite direction; there are some relatively uncontroversial ways in which language is a strikingly *robust* phenomenon. The most obvious point here is just that linguistic communication works, and that it works reliably and non-arbitrarily. To be sure, temporal, geographical and cultural discontinuities present obstacles to communicative success; but, one way or another, speakers/ hearers know how to overcome these obstacles. To introduce a slogan of which I am going to make heavy use in what follows: *when push comes to shove*, speakers and hearers know how to retreat to a common ground, and know that doing so will help to ensure communicative success. Retreating to a common ground presupposes that there is a common ground, and that speakers

and hearers recognize it (in some sense) as such.[23] Moreover, to repeat a now familiar point, the need for a common ground becomes even more pressing against the presumptively messy background of a use-theoretic conception of language.

I submit that underwriting the robust nature of communication should be at the heart of any notion of literal meaning worth wanting. Given this assumption, my worry about Neale's blueprints is that there is little reason to think they will be sufficiently robust. One might respond by pointing out that Neale's invocation of blueprints is merely a starting point, one that needs to be filled in by, among other things, substantial empirical research.[24] Moreover, Neale's motivation for introducing the notion of a blueprint is not so much to articulate a positive account of literal meaning as to make the negative point that *whatever* literal meaning is, it is apt to vastly underdetermine (i) what a speaker can intend to say/communicate and (ii) what a speaker reasonably can expect their audience to understand them as intending to say/ communicate.

These caveats noted, I want to introduce what I take to be an alternative way of thinking about how a notion of literal meaning might impose constraints on speakers. Rather than each actual usage being within a range of acceptable uses (as the notion of a blueprint seems to suggest) or being an elaboration or extension of acceptable use (as various discussions of 'what is said' seem to suggest), I want to suggest that the literal meaning of a sentence imposes constraints on speakers *indirectly,* by determining how they *would* use a sentence *were* they to be in a kind of ideal context (what I will call—in spite of obvious worries about circularity—a *literal* context). The basic idea builds on the aforementioned suggestion that language users share a common ground to which they retreat when push comes to shove; in particular, they retreat in the direction of a shared expectation that sentences will be used literally.

Let me add some flesh to this picture. I assume that in most conversational contexts, speakers start out with a wide range of information that they would be happy to communicate, some pieces of information more important than others. Speakers select from among various sentences that they could utter based on their estimates of what speakers would be able to extract from these sentences (as uttered in the particular context in question). To clarify, I am not claiming that speakers need to have an explicit conception of exactly what information they intend to

---

[23] Another way in which language is robust, one that is also in prima facie tension with the manifest flexibility, context-dependence, and creativity of language use, is the extent to which language is both *productive* and *systematic*. The natural explanation for these properties— perhaps the only plausible explanation—is that language and/or thought is both syntactically and semantically *compositional*, and arguably in a particularly strict way (Fodor and Lepore, 2002).

[24] There are suggestions in Neale (ms) that any workable notion of 'what is said' will need to be carefully circumscribed. The point, as I understand it, is that 'what is said' needs to be able to underwrite notions of commitment and contract across a potentially broad range of contexts (conversational or otherwise), and that the relevant constraints on 'what is said' bottom out in constraints imposed by the relevant blueprints.

communicate. Nor do speakers need to have an explicit conception of exactly what information needs to be communicated for the interaction to count as a success. Nor do they need to have a direct or infallible way to figure out what or how much information has been communicated. What speakers do have, according to the picture I want to advocate, is a potentially complex set of dispositions to react in certain ways given different kinds of feedback from their audience. Most important for present purposes is the suggestion that when a speaker perceives a threat to communicative success, they are disposed to retreat to uses of language that are increasingly literal. The reason they retreat in this direction is obvious given the view that I am advocating: literal meaning is a matter of conventions that are *robust*, inter alia that are shared across users of a language. In short, while it might not be efficient, literal meaning is relatively safe.

Without committing to details that outstrip the grain of the present analysis, more can be said about what it means to retreat to increasingly literal uses of language. In some cases, merely *restating* an initial sentence can stave off the threat of communicative failure. Consider a situation in which Ann and Ben are sitting at opposite ends of a table. The table is very long and thus Ann has a legitimate reason to wonder if Ben is able to pass various condiments located at his end of the table. Consider the following dialogue, the heart of which is a *restatement* of an initial sentence:

   Ann: "Can you pass the salt?"

   [Ben picks up the salt and pauses, evidently reflecting on whether he can get the salt down the table without getting up from his chair]

   Ann: "No, don't bother; I was really just asking if you *can* pass the salt."

In other cases, a mutually assumed shift towards an increasingly literal context will render a previous sentence misleading, but the speaker can get away with *refining* it in various ways. Cases of refinement include the familiar notion of a pragmatic implication being *cancelled*:

   Casey: "Did you hear that Jackson and Jillian had children?"

   David: "Yes, but I've heard two names for the baby; did they have twins?"

   Casey: "No, Jackson and Jillian had children, *but not together*."

In yet other cases, the threat of communication breakdown is so strong, or the shift to an increasingly literal context does so much to undermine the felicity of the previous utterance, that the speaker needs to *retract and replace* their previous sentences, as in the dialogue below.

   Evan: "I went to have a look at the aircraft carrier last night and, wow, that thing is a mile long!"

15

Frida: "Really? that seems unlikely."

Evan: "No, not *really*, but it is a lot longer than I expected"

A clarification: the fact that speakers *restate*, *refine*, and *retract and replace* does not entail that any particular use of a sentence is *purely* literal. What matters to the approach is that speakers/ hearers are disposed to retreat in the direction of literal usage, not that they ever are pushed all the way to a context where only strictly literal usage is felicitous. Note that a retreat to increasingly literal usage is driven by practical considerations, but that practical considerations cut in two directions: On the one hand, there is pressure to retreat to literal usage when communication threatens to break down. On the other hand, setting up a literal context—a context where hearers assume that the speaker is speaking literally, and speakers can rightly assume that their hearer so assumes, and so on—can be laborious and inefficient. Thus there is little reason to think that typical cases where an utterance is restated, refined, or retracted and replaced are anything other than steps in the direction of literal usage.

The potential unattainability of a purely literal context is a virtue of the account. It helps to make sense of how difficult it is to isolate exactly what a speaker has in mind in uttering a particular sentence, even in highly rarified contexts where a speaker is under substantial pressure to ensure that communication is successful. This is an example of what I take to be the most obvious, overarching virtue of the push-comes-to-shove approach to literal meaning, namely that it promises to capture the best of two prima facie competing worlds. On the one hand, a relatively austere notion of literal meaning promises theoretical tractability and robustness across agents and contexts.[25] On the other hand, the fact that actual conversational contexts are rarely, if ever, strictly literal leaves room for the obviously messy features of language use (context-dependence, plasticity, creativity, etc.)

Exactly what distinguishes contexts of increasingly literal use? I can't hope to offer a fully worked out account here (nor do I have one); but I can offer two broad suggestions regarding how the use of an expression can *fail* to be literal. First, a speaker can fail to use an expression literally for *secondary practical* reasons. Speaking very generally, in a context where speaker and hearer share a sufficient number of background beliefs, a speaker can assume (and assume that her audience will assume) that a sentence used non-literally will succeed in communicating more information than a different sentence used (relatively) literally. Of course, this is little more than a standard gesture to the need to invoke pragmatic mechanisms (e.g. Soames' appeals to descriptive enrichment). Nevertheless, the underlying point is that various practical considerations, while contributing to the efficiency with which language can be used to communicate information, are *secondary* in that they tend to be overridden in contexts where communicative success is under threat.

_____

[25] Austerity also makes it easier for a semantic theory to live up to a compositionality requirement (see note 23).

16

A second reason speakers can fail to use an expression literally is (again very roughly) because they are suffering from relevant *epistemic deficiencies*. Things get tricky here insofar as attempts to distinguish issues properly belonging to semantics from those better classified as having to do with the epistemic situation of speakers/hearers is bound to induce controversy. The basic idea, at any rate, is that various epistemic deficiencies can lead a speaker to use an expression in a way that outstrips or even violates its literal meaning.

The possibility of epistemic constraints on what counts as literal usage introduces the need to shift to a *theorist's* perspective. What I mean by this is that epistemic deficiencies can prevent a speaker (mutatis mutandis their audience) from recognizing that a use of an expression is, in a relevant respect, non-literal. At risk of courting controversy, let me offer an example. The example has the benefit of gesturing in the direction of how the push-comes-to-shove approach plays into the hands of Direct Reference theorists. The following is an example of a typical problem case for Direct Reference. Gary is aptly described as failing to realize *that Cilantro is Coriander* (more details about Gary's epistemic situation in a moment). In a suitable context involving a conversation with his friend Hattie, Gary utters (1):

(1)   Cilantro is not Coriander

According to Direct Reference, (1) is semantically equivalent to (2).

(2)   Cilantro is not Cilantro

There is little reason to think that raising the practical stakes will always be a way to generate contexts that expose the presumptive synonymy of expressions like "Cilantro" and "Coriander". Perhaps no amount of pragmatic pressure would give Gary a reason to let go of his commitment to the truth of (1); whereas, if queried, he would refuse to assent to (2). Nevertheless, my suggestion is that the Direct Reference theorist can explain Gary's inability to retreat to an increasingly literal context in terms of his suffering from various epistemic deficiencies. For example, Gary doesn't realize *that "Cilantro" and "Coriander" are names for the same plant*, and he doesn't realize *that the plant on his window sill produces the same kind of seeds that are in his cupboard in the jar marked "Coriander"*, and so on. These, among others, are true statements about what Gary (and perhaps also Hattie) doesn't know.[26] Importantly, there does not need to be any particular epistemic deficiency to which the Direct Reference theorist appeals in all situations in which a speaker utters (1).[27] What matters is that the Direct Reference theorist can gesture towards a relevant epistemic deficiency in each particular case.

---

[26] Depending on how details of the view are worked out, it can be important to the Direct Reference theorist that these statements about Gary's relevant epistemic deficiencies are describable in terms that do not presuppose distinctions that go beyond those sanctioned by a commitment to Direct Reference.

[27] See the related discussion in §1 under the heading, "A digression about descriptivism".

Before going any further, I want to clarify that I am not tying myself to the claim that what a speaker communicates when (practical and epistemic) push-comes-to-shove *defines* or provides an *analysis* of literal meaning. Defending this claim would require a much richer story about purely literal contexts than I have to offer. Nevertheless, I hope I have made it plausible (1) that the notion of how speakers (would) react when practical and epistemic push-comes-to-shove is one with intuitive force and explanatory potential, (2) that it is relatively easy to see how to elaborate on this basic picture in potentially productive ways (for example by filling in details about the pragmatic mechanisms that speakers/hearers are under pressure to abandoned when push-comes-to-shove).[28] I now want to add some flesh to the suggestion that the Direct Reference theorist can hope to exploit the push-comes-to-shove framework. And I want to do this in a way that brings Kripkean considerations about an expression's modal and epistemic profile back into the discussion.

Kripkean considerations enter the picture with the suggestion that being under pressure to respect an expression's epistemic and modal profile is partly constitutive of the notion of literal usage. As an illustration, consider a discussion between Ian and Jackie about how mathematicians tend to do their best work early in life. Given the right conversational background, one can imagine Ian using a proper name to communicate descriptive information, which I will indicate in square parentheses:[29]

> Ian: "Wiles [the discoverer of the proof of Fermat's last theorem] did great work when he was over forty."

Now imagine the conversation continuing as follows:

> Jackie: "Wiles is an exception; what if Wiles had *not* proved Fermat's last theorem?

---

[28] There are interesting questions about what kinds of considerations motivate shifts toward increasingly *non*-literal usage. For example, one can imagine conversations where a metaphor is discharged in terms of a more literal explanation but then speakers/hearers shift back to using the metaphor because it is more efficient. More interesting is the possibility of situations in which speakers are happy to be interpreted as intending to communicate a range of propositions but efficiency demands that they not be held hostage to spelling out each and every one. Perhaps typical utterances of "I'm ready" [to go to the party / to leave the house / to go out on the town / ?] might be like this (to pick a favorite example of Stephen Neale's).

[29] After writing the present paper, I came across the fact that Kripke (1977) draws on similar conversational considerations as partial support for a 'unified' Russellian account of descriptions. The two appeals are nicely complementary. Conversational considerations suggest that *both* descriptive uses of referring expressions and so-called referential uses of descriptions can be explained in terms of pragmatic features of language use, rather than in terms of the literal meaning of the expressions in question. (The *referential* versus *attributive* distinction is due to Donnellan (1970)).

Ian: "Wiles would still have done great work when he was past forty, even if he didn't prove Fermat's last theorem."

Here is what the aspiring Direct Reference theorist can say about this conversation. First, there is an obvious shift in context between Ian's two utterances. The shift is such that Ian can no longer use "Wiles" to express the relevant descriptive information. Ian is been pushed to consider counterfactual circumstances in which Wiles did not complete the proof. Because of this, he is under pressure to make sure that his use of "Wiles" does not presuppose otherwise. Second, the Direct Reference theorist can claim to have a principled reason for treating this shift as a step in the direction of increasingly literal usage. Recall that preserving something *robust* (sharable, context independent, etc.) at the core of language use is a central motivation for wanting to extract a notion of literal meaning from the LP framework. Robustness across counterfactual circumstances is an important form of robustness. Thus, it is natural for the Direct Reference theorist to claim that increased sensitivity to features of an expression's modal profile is part and parcel of a shift towards increasingly literal usage.

The same points can be made about pressure on speakers to manifest sensitivity to an expression's epistemic profile (or, more precisely, to the epistemic profile of sentences of which the expression is a constituent). As an illustration, imagine a conversation between two historically (and philosophically) ignorant students, who are discussing whether anyone in ancient Greece would have been smart enough to pass the exam they have just suffered through:

Karl: "Plato [the ancient Greek philosopher] was really smart; he would have aced it."

Laura: "I'm sure Plato was really smart, but you don't know that he was an ancient Greek; you just heard that from your roommate, who is a terrible source of information."

Let's assume that Laura's utterance is enough to make Karl worried about the source of his assumption. He can no longer use "Plato" as a means of communicating the relevant descriptive information. Again, it is natural for the Direct Reference theorist to claim that this shift in context is a shift towards an increasingly literal context. And, again, this is consistent with the general phenomenon of increasingly literal contexts involving uses of language that are increasingly robust, this time across differences in an agent's epistemic situation.

A quick recapitulation is in order. I have suggested a way to extract a notion of literal meaning from the basic LP framework. The heart of the approach is the idea that expressions in natural language are *backed* by a corresponding concept. The backing relation, I have suggested, is best approached by thinking about how speakers are disposed to respond when pragmatic and epistemic push-comes-to-shove. Kripkean considerations entered the picture with the suggestion that increased sensitivity to an expression's modal and epistemic profile can be understood as instances of the general phenomenon that speakers retreat to increasingly robust uses of language when push-comes-to-shove. The obvious suggestion is that the aspiring Direct Reference theorist can hope to claim that the notion of literal meaning that emerges from this picture is one that

vindicates Kripkean Intuitions. More specifically, they can hope to claim that the notion of literal meaning that emerges is Directly Referential.

4. Being pragmatic about the Threat of Collapse

I can now be explicit about how all of this bears on the two-source solution. Working within the basic framework of LP, it is relatively easy to account for Fregean Intuitions. At the core of LP is the idea that speakers and hearers exploit their linguistic resources in potentially open-ended ways. Given this framework, rich and fine-grained intuitions about 'meaning', cognitive significance, and so on, are surely to be expected. The substantive question facing the aspiring Direct Reference theorist is whether something can be extracted from the LP framework that is plausibly the source of the Kripkean Intuitions. My claim in the previous section was that the push-comes-to-shove approach supplies a notion of literal usage that has this feature.

Actually, I think the question of whether the relevant notion of literal usage can produce the Kripkean Intuitions is dialectically misguided. The Direct Reference theorist who buys into my suggestion that Direct Reference ought to be embedded in an LP framework does not need to treat standard Kripkean Intuitions as anything other than a starting point. The important question is whether the referentialist-qua-pragmatist can point to sufficiently powerful considerations—intuitive or theory-driven—to motivate a referentialist commitment within the framework of LP. My suggestion has been that taking seriously the push-comes-to-shove approach promises to do this. Of course, insofar as a Directly Referential notion of literal meaning can be extracted from the LP framework, it isn't surprising to find something aptly described as versions or analogues of the original Kripkean Intuitions; for example, considerations about how speakers react to various kinds of conversational pressure.[30]

Let's look at a specific example, in particular a case discussed by Caplan (2007). I will follow Caplan in taking for granted the following as background: First, in the relevant context, "Cary" and "Archie" are names for the same individual, "Cary Grant" being the screen name of Archibald Alexander Leach. Second, I'll assume that the descriptive information, *the actor who played C.K. Dexter in The Philadelphia Story*, is a plausible candidate for what the Millian Descriptivist might claim is conveyed by "Cary" but not by "Archie". Caplan's worry, then, centers on a sentence such as (3).[31]

---

[30] There is at least one way in which the referentialist-qua-pragmatist needs to reconfigure standard intuitions about the semantic, modal, or epistemic profile of prima facie referring expressions: According to the pragmatist, talk of what sentences (inter alia their constituent expressions) mean is at best short-hand for talk of conventional constraints on what speakers can hope to communicate in uttering a sentence.

[31] I've simplified Caplan's sentences by glossing over a point about which Caplan is rightly careful; namely, the possibility of counterfactual circumstances in which Cary/Archie fails to exist, and hence where (3) is trivially true.

(3)   Cary is an actor.

The problem starts with the fact that if "Cary" conveys the descriptive information *the actor who played C.K. Dexter in The Philadelphia Story*, then (3) should be equivalent to (4).

(4)   The actor who played C.K. Dexter in *The Philadelphia Story* is an actor.

(4) is both necessary and knowable a priori; therefore, the Millian Descriptivist ought to be committed to the claim that (3) is necessary and knowable a priori. But this runs contrary to Kripkean Intuitions. Hence the conclusion that I have credited to Caplan: *the Millian Descriptivist's motivations for a Direct Referential semantics are fated to collapse given their account of the relevant Fregean Intuitions.*

Here is what the resources that I have outlined allow the Direct Reference theorist to say about this case: On the one hand, Fregean Intuitions are to be expected given the background framework of LP. Utterances of sentences are mutually understood by speakers and hearers to communicate a potentially wide range of information. Insofar as there is any intuitive pull towards including *the actor who played C.K. Dexter in The Philadelphia Story* in the information conveyed by a particular utterance of "Cary", the referentialist-qua-pragmatist will have an easy time accommodating this. But now let's be more careful about the kind of context that might generate such an intuition. For example, imagine that Mary and Ned have just watched *The Philadelphia Story*. Imagine further that Mary is enamored with the leading star of the film and has been subjecting Ned to one after another of his films over the past few weeks. In an underhanded attempt to strike back, Ned has tracked down an article detailing Archie Leach's marital exploits; indeed, he has conveniently left the article open on a page detailing accusations to the effect that Archie Leach spanked at least one of his five wives. Unfortunately for Ned, the article is written without mentioning the name "Cary Grant". When Mary reads the article, she comments that "Archie Leach must have been a chauvinist pig", but she never connects the person described in the article with the star of the films she has been watching. Given these details of the case, how can the referentialist-qua-pragmatist hold onto the claim that even though Ned uses (3) to communicate the relevant descriptive information, this information is not part of the expression's semantics?

Here is how: First, it is clear that the conversational context in question is not a purely literal context, since Mary is suffering from an obvious cluster of epistemic deficiencies. For example, Mary doesn't realize that "Cary" and "Archie" (as used in this context) refer to the same person, she doesn't realize that the actor in the movie she just watched is none other than the person about which the article was written, and so on. In this particular case, Ned himself is in a better epistemic position, and he can take steps to shift the conversational context such that it is increasingly literal. For example, he can utter (5):

(5)   You know, the person described in the article *just is* the actor that you are in love with. If you don't believe me, look it up online.

The resulting context is more literal than the previous context, and hence a better insight into what constitutes a literal usage of the names in question. And note that in the resulting context, the relevant descriptive information can no longer be used to draw a wedge between "Cary" and "Archie", and hence cannot be part of the content of one but not the other. So, even if a use of "Cary" in a relevant utterance of (3) conveys the descriptive information "the actor who played…", the referentialist-qua-pragmatist should not countenance it as part of the name's literal meaning. Again, the general moral is that even if typical uses of an expression communicate descriptive information, there can be reasons to cleave to the claim that this information is not part of the expression's literal meaning. Thereby, I claim, the relevant Threat of Collapse can be disarmed.

What would Caplan say about this? While Caplan does not use the term "two source solution", it is clear that he has considered the basic strategy. Here is what he says:

> It is implausible that speakers' intuitions about modal profile have a different source than their intuitions about truth-value. After all, the modal profile of a sentence includes, not only whether it's necessary or contingent, but also which worlds it's true in and hence whether it's actually true; and it's implausible that speakers' intuitions about whether a sentence is actually true should have a different source than their intuitions about whether it's true simpliciter. Similarly, it's implausible that speakers' intuitions about epistemic profile have a different source than their intuitions about cognitive value. After all, the cognitive value of a sentence includes whether it's a priori or a posteriori, as does its epistemic profile. (2007, p. 189)

I want to underline two points in response, neither of them new. I'll focus on what Caplan says about the relationship between intuitions about modal profile and intuitions about truth-value. First, as clarified in the previous section of the paper, the Direct Reference theorist *qua pragmatist* does not need to hold themselves hostage to the project of rescuing original Kripkean Intuitions about the cases in question; they are free to appeal to whatever considerations motivate their commitment to a referentialist semantics. I have suggested that within the general framework of LP, there is good reason to embrace the existence of what amounts to a referentialist core of literal usage.

Second, and more importantly, I disagree with Caplan's claim about the implausibility of intuitions about modal/epistemic profile having a different source than intuitions about truth value / cognitive value. For starters, standard intuitions about truth value, informativeness, and more generally about what an expression 'means' typically are largely pre-theoretical, whereas generating robust Kripkean Intuitions often involves unpacking, or at least clarifying, theoretical issues. Moreover, there are some important differences between the kinds of questions that are

useful in prompting each category of intuition. For instance, when asked about the meaning of a belief ascription, it is easy to interpret the question relatively loosely, as a question about what the sentence typically is used to say about the believer. Without some reason to be careful about drawing a distinction between what is semantically expressed versus pragmatically conveyed by such an ascription, or between what the sentence itself means and what attributers mean in uttering the sentence, it is no wonder that intuitive responses tend to be informationally rich. Contrast this with questions intended to probe intuitions about the modal and epistemic profile of a sentence. In answering these questions, an informant is forced to look for features that are shared across modal and epistemic circumstances. This has the potential to *constrain* rather than *broaden* the scope of potentially relevant information. As such, it is natural for the informant to settle on something much narrower than what guides intuitions about informativeness, cognitive significance, and even truth value. *Pace* Caplan, it is Kripkean Intuitions that are more restricted in the relevant respects.

5. Conclusion

The discussion in this paper suggests conclusions at several levels of abstraction. Beginning with the most concrete, I hope to have made it plausible that if the Direct Reference theorist is willing to embed their core semantic commitments within the larger framework of a pragmatist conception of natural language, they can hope to lay claim to more than sufficient resources to avoid the Threat of Collapse. To be fair, I haven't argued for Linguistic Pragmatism so much as take it for granted; nevertheless, the view is certainly not implausible. Moreover I think there are various signs that Linguistic Pragmatism and Direct Reference are a powerful combination.

I also suggested what I take to be an alternative way of approaching the notion of literal meaning or, more precisely, literal usage. This is, no doubt, the most controversial territory explored in the paper. Nevertheless, the push-comes-to-shove approach has notable virtues; in particular the extent to which it draws on facts about how speakers are disposed to react in various conversational settings. This, in turn, focusses attention on the extent to which the referentialist-qua-pragmatist can hope to lay claim to what I have described as the best of both worlds. The Direct Reference theorist can hope to embrace a full-blooded, anything-goes-as-long-as-it-works conception of language use while holding onto a very austere conception of literal meaning. This overall package promises to explain the manifest plasticity, context-dependence, and creativity of language use without this posing a threat to the robust and non-arbitrary nature of linguistic communication. In a nutshell, the approach that I have outlined goes right to the heart of how language can be so disastrously messy and yet work so well.

Bibliography

Bach, K. (2001) 'You Don't Say?', *Synthese,* 128(1/2), 15-44.

Block, N. (1986) 'Advertisement for a Semantics of Psychology', In French, P. A., Uehling, T. E. J. & Wettstein, H. K. (eds.) *Midwest Studies in Philosophy* (Minneapolis: University of Minnesota Press), pp. 615-678.

Braun, D. (2005) 'Empty Names, Fictional Names, Mythical Names', *Nous,* 39(4), 596-631.

Braun, D. (2006) 'Illogical but Rational', *Nous,* 40(2), 376-379.

Braun, D. (2008) 'Complex Demonstratives and Their Singular Contents', *Linguistics and Philosophy,* 31(1), 57-99.

Braun, D. & Sider, T. (2006) 'Kripke's Revenge', *Philosophical Studies,* 128, 669-682.

Caplan, B. (2007) 'Millian Descriptivism', *Philosophical Studies,* 133(2), 181-198.

Carston, R. (2002) 'Linguistic Meaning, Communicated Meaning and Cognitive Pragmatics', *Mind & Language,* 17(1-2), 127-148.

Crimmins, M. (1992) *Talk About Beliefs* (Cambridge, MA: MIT Press).

Davidson, M. (2007) *On Sense and Direct Reference: Readings in the Philosophy of Language* (Boston: McGraw-Hill).

Donnellan, K. (1970) 'Proper Names and Identifying Descriptions', *Synthese,* 21, 335-358.

Dretske, F. I. (1981) *Knowledge and the Flow of Information* (Oxford: Basil Blackwell).

Edwards, K. (ms) 'Keeping Reference in Mind', Paper delivered at the Pacific meeting of the American Philosophical Association, April 2009.

Fodor, J. A. (1975) *The Language of Thought* (Cambridge, MA: Harvard University Press).

Fodor, J. A. (1981) *Representations: Philosophical Essays on the Foundations of Cognitive Science*, 1st MIT Press edn (Cambridge, MA: MIT Press).

Fodor, J. A. (1992) *A Theory of Content and Other Essays* (Cambridge, MA: MIT Press).

Fodor, J. A. (1998) *Concepts: Where Cognitive Science Went Wrong* (Oxford: Oxford University Press).

Fodor, J. A. (2004) 'Having Concepts: A Brief Refutation of the Twentieth Century', *Mind & Language,* 19(1), 29-47.

Fodor, J. A. (2008) *LOT 2: The Language of Thought Revisited* (New York: Oxford University Press).

Fodor, J. A. & Lepore, E. (2002) *The Compositionality Papers* (Oxford: Oxford University Press).

Frege, G. (1952) 'On Sense and Reference', *Translations from the Philosophical Writings of Gottlob Frege* (Oxford: Basic Blackwell), pp. 56-78.

Harman, G. (1982) 'Conceptual Role Semantics', *Notre Dame Journal of Formal Logic,* 23(2), 242-256.

Kaplan, D. (1989) 'Demonstratives', In Almog, J., Perry, J. & Wettstein, H. (eds.) *Themes from Kaplan* (New York: Oxford University Press), pp. 481-564.

Kripke, S. (1977) 'Speaker's Reference and Semantic Reference', *Midwest Studies in Philosophy,* II, 255-277.

Kripke, S. (1979) 'A Puzzle About Belief', In Margalit, A. (ed.) *Meaning and Use* (Dordrecht: D. Reidel), pp. 239-283.

Kripke, S. (1980) *Naming and Necessity* (Cambridge, MA: Harvard University Press).

Loar, B. (1981) *Mind and Meaning* (Cambridge; New York: Cambridge University Press).

Loar, B. (1987) 'Names in Thought', *Philosophical Studies,* 51, 169-185.

McKay, T. (1981) 'On Proper Names in Belief Ascriptions', *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition,* 39(3), 287-303.

McKinsey, M. (2005) 'Critical Notice of Scott Soames, *Beyond Rigidity*', *Canadian Journal of Philosophy,* 35(1), 149-168.

Mill, J. S. (1858) *A System of Logic, Ratiocinative and Inductive*, 5th edn (New York: Harper & Brothers).

Millikan, R. G. (1993) *White Queen Psychology and Other Essays for Alice* (Cambridge, MA: MIT Press).

Neale, S. (2004a) 'Pragmatism and Binding', In Szabo, Z. (ed.) *Semantics and Pragmatics* (Oxford: Oxford University Press).

Neale, S. (2004b) 'This, That, and the Other', In Bezuidenhout, A. (ed.) *Descriptions and Beyond* (New York: Oxford University Press).

Neale, S. (ms) 'Linguistic Pragmatism', Unpublished manuscript.

Peacocke, C. (1992) *A Study of Concepts* (Cambridge, MA: MIT Press).

Recanati, F. (2001) 'What Is Said', *Synthese,* 128(1-2), 75-91.

Russell, B. (1905) 'On Denoting', *Mind,* 14(56), 479-93.

Salmon, N. (1986) *Frege's Puzzle* (Cambridge, MA: MIT Press).

Salmon, N. (1989a) 'How to Become a Millian Heir', *Nous,* (23), 211-220.

Salmon, N. (1989b) 'Illogical Belief', *Philosophical Perspectives* (Atascadero: Ridgeview).

Salmon, N. (2005) *Reference and Essence*, 2nd edn (Amherst, NY: Prometheus Books).

Sider, T. & Braun, D. (2006) 'Kripke's Revenge', *Philosophical Studies,* 128(3), 669-682.

Soames, S. (1997) 'Direct Reference, Propositional Attitudes, and Semantic Content', In Ludlow, P. (ed.) *Readings in the Philosophy of Language* (Cambridge, MA: MIT Press), pp. 921-962.

Soames, S. (2002) *Beyond Rigidity: The Unfinished Semantic Agenda of Naming and Necessity* (New York: Oxford University Press).

Soames, S. (2005) ''Beyond Rigidity': Reply to McKinsey', *Canadian Journal of Philosophy,* 35(1), 169-178.

Soames, S. (ms) 'What Are Natural Kinds?', Forthcoming in *Philosophical Topics*.

Thau, M. (2002) *Consciousness and Cognition* (Oxford; New York: Oxford University Press).