

Cet obscur objet du ‘désir’

Kyle Blumberg & Milo Phillips-Brown

February 16, 2020

Abstract

The Relational Analysis of attitude reports posits a transparent link between language and mind: the agent is ascribed an attitude whose object can be seen by looking at the grammatical object of the attitude verb. (For example, ‘Frank wants to have an orange’ is true iff Frank stands in the desire relation to the proposition that he has an orange—i.e. the proposition denoted by the grammatical object of ‘want’.) We argue that with desire reports and wanting, the link between language and mind is not transparent. Rather, it is *obscure*: as a previously unrecognized form of context sensitivity in desire ascriptions shows, the objects to which agents stand in the desire relation cannot usually be seen merely by looking at the object of the attitude verb. This is the *obscurity problem*, which falsifies both the Relational Analysis and all mainstream semantics for ‘want’. We offer a replacement for the Relational Analysis—the *Obscure Object Analysis*—that solves the problem and allows one to recover versions of the mainstream semantics. Along the way, we discuss how the obscurity problem and the Obscure Object Analysis relate to a question, brought to prominence by Fara (2013), of whether desire reports are in some sense “underspecified.”

1 Introduction

When you utter an attitude report, what do you ascribe to a person’s mind? For example, when you say, ‘Priya believes that a hard rain is falling’, or ‘Priya doubts that she’ll be promoted’, or ‘Priya wants to donate to charity’, what are you saying about how Priya’s mental states connect to hard rain falling, to being promoted, to donating to charity?

The orthodox view of attitude reports—the Relational Analysis, widely en-

dorsed by philosophers and linguists alike¹—answers straightforwardly. Taking belief reports and belief as an example, the view says that ‘ S believes p ’ is true just if the agent denoted by S bears the belief relation to the proposition denoted by p . Thus, ‘Priya believes that a hard rain is falling’ is true just if Priya bears the belief relation to the proposition that a hard rain is falling. That’s what your language ascribes to Priya’s mind.

More rigorously, the view is this:²

Relational Analysis

Where Φ is an attitude verb, and c is a context:

- (i) ‘ $S \Phi s p$ ’ is true in c iff S bears the relation expressed by Φ to p in c ;
- (ii) S bears the relation expressed by Φ to p in c iff S bears the relation that’s intuitively expressed by Φ to p in c .
(e.g. S bears the relation expressed by ‘believe’ to p in c iff S bears the belief relation to p in c , S bears the relation expressed by ‘want’ to p in c iff S bears the desire relation to p in c , etc.)

The Relational Analysis posits an intuitive and theoretically pleasing *transparent* link between language and mind: the agent is ascribed an attitude whose object can be seen simply by looking at the grammatical object of the attitude verb. For example, ‘Priya believes that a hard rain is falling’ is true just if Priya bears the belief relation to an object—the proposition that a hard rain is falling—that can be seen simply by looking at the grammatical object of the attitude verb, ‘that a hard rain is falling’. A person couldn’t be blamed for thinking that language and mind *must* be linked in this way!

Yet they are not. With desire reports and wanting, the link between language and mind is not transparent. Rather, it is *obscure*. That is the thesis of this paper, a thesis that breaks into two parts: one negative and one positive.

The negative part is a rejection of the Relational Analysis as applied to ‘want’:

¹In the philosophical literature, see e.g. (Stalnaker, 1988) and (Schiffer, 2003). In the linguistics literature, see e.g. von Stechow and Heim’s influential textbook on intensional semantics (von Stechow and Heim, 2011, ch. 2).

²We let ‘ S ’ range over the names of agents and let ‘ S ’ range over the corresponding agents denoted by ‘ S ’. Similarly, we let ‘ p ’ range over the logical forms of proposition-denoting strings and let ‘ p ’ range over the corresponding propositions denoted by ‘ p ’.

Relational Analysis for ‘want’ (RAW)

- (i) $\lceil S \text{ wants } p \rceil$ is true in c iff S bears the relation expressed by ‘want’ to p in c ;
- (ii) S bears the relation expressed by ‘want’ to p in c iff S bears the desire relation to p in c .³

A case will illustrate why RAW does not deserve your approval. Imagine that Frank—who loves juicy oranges—is at a lunch where each person will get a piece of fruit. He thinks that he might get an orange, so he’s happy about that. But then a friend delivers some bad news: all the oranges available are dessicated. Frank says, ‘That’s a pity, I really want to have an orange’. In this context, *Orange* is true:

(*Orange*) Frank wants to have an orange.

Imagine further that as the caterers are preparing, one of them comes by Frank’s empty seat carrying a basket of fruit. He asks, ‘Would Frank like to have an orange?’ You reply, ‘No, Frank doesn’t want to have an orange’.⁴ In *this* context, *Orange* is false.

What’s different between the two contexts? Intuitively, it’s that in the context where *Orange* is true, its truth owes to Frank looking favorably on having *juicy* oranges. In the context where *Orange* is false, its falsity owes to Frank not looking favorably on something *different*—namely, having a *dessicated* orange. Put another way, the truth value of *Orange* is not the same in the two contexts because the relevant object of the desire relation is not the same in the two contexts. In the context where *Orange* is true,

³RAW assumes that ‘want’ expresses a relation between subjects and propositions. However, a number of philosophers have argued that not all “object-oriented” attitude reports, e.g. ‘John loves Mary’, can be successfully analyzed in propositional terms (Ben-Yami, 1997; Crane, 2001; Montague, 2007; Grzankowski, 2012). Some want reports are also object-oriented, e.g. ‘Bill wants a house’, so one might worry that RAW rests on a false assumption. We want to remain neutral about this concern here. We only maintain that want reports such as *Orange* (introduced just below) are propositional, and can be analyzed as in (1) (also just below). Even proponents of the existence of non-propositional attitudes recognize that there are compelling reasons to think this is the case—see e.g. (Montague, 2007).

⁴‘want’ is a so-called “neg-raising” predicate (Collins and Postal, 2014). That is, some speakers hear $\lceil S \text{ doesn’t want } p \rceil$ as equivalent to $\lceil S \text{ wants } \neg p \rceil$, where negation takes narrow scope with respect to the attitude verb. The relevant interpretation of your assertion is *not* one on which it is neg-raised—i.e. the relevant interpretation is ‘it’s not the case that Frank wants to have an orange’.

the relevant object of the desire relation is Frank having juicy oranges; in the context where it's false, the relevant object is Frank having dessicated oranges.

This is not what RAW would lead one to expect. RAW says that in every context, the truth of *Orange* hinges on whether Frank bears the desire relation to the *same* object—having an orange *simpliciter*. To bring out the point a bit more, consider the context where *Orange* is true. As mentioned, the report's truth in this context owes to Frank bearing the desire relation to a certain object—having a juicy orange. This object cannot be seen simply by looking at the object of 'want'. *That* object is the complement of 'want' in the logical form of *Orange*, a logical form that we assume—following standard linguistic practice⁵—is the following:

- (1) Frank wants [PRO_{Frank} to have an orange].

The nonfinite complement construction in *Orange* contains a “silent” subject term PRO_{Frank} that refers to Frank. Thus, the object of 'want' expresses the proposition that Frank has an orange.⁶ And so, to repeat, the object of the desire relation, which is Frank having a juicy orange, cannot be seen simply by looking at the object of the 'want' report, which is that Frank has an orange *simpliciter*. That's not transparency; it's obscurity. Consequently, we call this the *obscurity problem* for RAW.

The positive part of the paper presents a replacement for RAW. Rejecting RAW presents a choice-point: give up (i), i.e. the claim that $\lceil S \text{ wants } p \rceil$ is true just in case *S* bears the relation expressed by 'want' to *p*, or give up (ii), i.e. the claim that *S* bears the relation expressed by 'want' to *p* iff *S* bears the desire relation to *p*. We consider both options and ultimately go for the second.

⁵See e.g. (den Dikken et al., 1996) and (Larson, 2002).

⁶It is worth remarking on two features of our logical form for *Orange*. First, ambiguities can arise when determiner phrases (such as indefinite descriptions) interact with attitude verbs. In particular, such constructions give rise to the *de dicto/de re* ambiguity. We assume that the indefinite in *Orange* is interpreted *de dicto* with respect to 'want'. This is because the *de re* reading is clearly irrelevant here. It can be paraphrased as, 'There exists a particular orange that Frank wants'. But we may simply suppose that Frank has no particular orange in mind, so this reading of *Orange* is false in context.

Second, some might think that the complement in *Orange* contains covert material that isn't represented in (1), e.g. the predicate 'juicy'. We consider this proposal in detail in §3, but ultimately reject it.

Our account—the *Obscure Object Analysis*—has two key aspects. First: *S* can stand in the relation expressed by ‘want’ to *p* without bearing the desire relation to *p*. Instead, *S* can stand in the relation expressed by ‘want’ to *p* in virtue of bearing the desire relation to some proposition *q* which entails *p*. For instance, in the context where *Orange* is true, the report is true because Frank bears the desire relation to the proposition that Frank gets a *juicy* orange, and Frank getting a juicy orange entails that he gets an orange. The second key aspect of the Obscure Object Analysis is that in order to explain the context sensitivity of *Orange*, i.e. the fact that it can also be false even though Frank’s beliefs and desires don’t change, we propose that the quantification over propositions is *contextually restricted*; context determines a domain of propositions over which the existential quantifier ranges. So, in the context where *Orange* is false, the proposition that Frank gets a juicy orange is excluded from the relevant domain.

It is worth being clear about how the obscurity problem impacts existing accounts of desire ascriptions. As we will show, this problem falsifies every mainstream semantics for desire reports—e.g. Heim’s (1992), von Stechow’s (1999), Levinson’s (2003), Villalta’s (2008), and Condoravdi and Lauer’s (2016). These semantics either predict that *Orange* should be true in every context, or that it should be false in every context, but none captures the pattern that we saw above: that *Orange* is true in some contexts and false in others. However, the problem with these accounts does not stem from their respective conceptions of desiring or the desire relation, but rather from the fact that each is implicitly committed to RAW. Thus, if one rejects RAW and replaces it with the Obscure Object Analysis, one can—without running afoul of the obscurity problem—recover a version of each semantics that preserves its basic spirit. So, although we will be posing a problem for existing accounts, we will also show how these semantics may be salvaged in the face of the obscurity problem.

The obscurity problem calls to mind the question—brought to prominence by Fara (2013)—of whether a sentence like *Orange* can be true even if Frank lacks a desire that is satisfied in exactly the worlds where he has an orange. More specifically, theorists wonder: can *Orange* be true even if Frank lacks a desire that is satisfied if he gets a dessicated orange? Some, such as

Fara, say yes, and thus that the proposition expressed by the complement in *Orange*—the proposition that Frank has an orange—“underspecifies” the satisfaction conditions of Frank’s desire. Others say no. Call the debate between these parties the *satisfaction underspecification debate*.⁷

We, like Fara, subscribe to the notion that desire reports are underspecified in a certain sense (although the sense we have in mind differs somewhat from hers—see our §4). And various elements in the satisfaction underspecification debate parallel elements in our dialectic, as we note throughout this paper. However, the obscurity problem is not fundamentally about desire satisfaction. Rather, the problem is about the (previously unnoticed) context sensitivity we introduced above. Consequently, while the satisfaction underspecification debate concerns a particular instance of RAW—what we call the “Content-specification Semantics” in §2.3—the obscurity problem has wider-reaching consequences, consequences for RAW itself and all of the mainstream semantics (the Content-specification Semantics among them).

The rest of this paper is structured as follows. In §2 we discuss the obscurity problem in greater depth and illustrate how it afflicts mainstream semantics for desire reports. Then in §3 we consider, and reject, a response to the problem that appeals to quantifier domain restriction. In §4 we present the Obscure Object Analysis and in §5 we conclude.

2 The obscurity problem

In the introduction, we claimed that the obscurity problem falsifies RAW. In this section, we present that claim in further detail.

RAW says that $\lceil S \text{ wants } p \rceil$ is true in c iff S bears the desire relation to p in c . The view takes different forms depending on how one understands the nature of S , i.e. of an agent; the nature of the desire relation; and the nature of p , i.e. the nature of propositions.

We believe that questions about the nature of agents and propositions are orthogonal to the obscurity problem. Why the nature of agents *would* bear on the obscurity problem we cannot imagine, and we make no special

⁷Those who answer no are Braun (2015) and Prinz (ms). In addition to Fara (2013), the following theorists answer yes: Fara (2003), Lycan (2012, ms), and Grant and Phillips-Brown (2019).

assumptions about them. What about propositions? We'll assume, for convenience and concreteness, that they are sets of possible worlds. Where p is a declarative sentence, the proposition expressed by p is the set of worlds in which p is true. For example, the proposition expressed by 'Petrar is a fan of the Beach Boys' is the set of worlds in which 'Petrar is fan of the Beach Boys' is true. Thus, we will assume that the proposition expressed by the complement of *Orange* ('Frank wants to have an orange') is the set of worlds where Frank possesses an orange.⁸ Of course, the possible worlds approach to propositional content is controversial (Soames, 1987), but again, we see no reason to believe that its truth or falsity bears on the obscurity problem.

The action is with the desire relation, which we'll address in this section. Our approach will be indirect. Specifically, we'll discuss three prominent mainstream semantics for 'want'—Heim's (1992) semantics, a decision theoretic semantics, and the Content-specification Semantics. Each semantics can be seen as instances of RAW—for, each is entailed by RAW in conjunction with different views about the nature of the desire relation—and each semantics is falsified by the obscurity problem. The falsity of these semantics tells us that the trouble for RAW runs deep. Precisify the view in any number of ways; it's still false.⁹

2.1 Heim's (1992) semantics

Heim patterns her semantics after Stalnaker's view that:

[W]anting something is preferring it to certain relevant alternatives, the relevant alternatives being those possibilities that the agent believes will be realized if he does not get what he wants.
(1984, p. 89)

Heim models an agent S 's beliefs with the set of worlds compatible with S 's beliefs— S 's belief set, or $\text{Dox}_{S,c}$. And Heim's notion of preference is between

⁸Our arguments will also go through if we instead assume that the proposition expressed by the complement of *Orange* is the set of worlds where Frank *eats* an orange. In that case, we would just stipulate that Frank hates wasting food, so that even if he was given a dessicated orange with his lunch, he would force himself to eat it.

⁹Indeed, the obscurity problem falsifies not just the semantics we target, but many others besides. Here is a sampling: Geurts (1998), von Fintel (1999), Villalta (2008), Crnić (2011), Anand and Hacquard (2013), and Phillips-Brown (2018).

worlds: roughly, S prefers a p -world w to its relevant alternative just if S prefers w to the most similar $\neg p$ -world to w . (This most similar $\neg p$ -world to w being w 's relevant alternative.)

The central idea behind Heim's semantics is this: given what S believes, does the truth of p in a world make that world an improvement over an otherwise similar world where p is false? If so, and only if so, is S truly said to want p . 'You want to get a model train set for your birthday' is true just if, given your beliefs, worlds where you get a model train set for your birthday improve upon otherwise similar worlds where you don't.

The inner workings of the semantics are these. Where $>_{S,c}$ represents S 's preferences between worlds (in a context c), and $\text{Sim}_w(p)$ is the p -world most similar to w :

*Heim's semantics*¹⁰

$\lceil S$ wants $p \rceil$ is true in c iff

$$\forall w \in \text{Dox}_{S,c} : \text{Sim}_w(p \cap \text{Dox}_{S,c}) >_{S,c} \text{Sim}_w(\neg p \cap \text{Dox}_{S,c})$$

This semantics can be conceived of as an instance of RAW, since it (the semantics) is entailed by RAW and the following account of the desire relation:

Heimian Desire Relation

S bears the desire relation to p in c iff

$$\forall w \in \text{Dox}_{S,c} : \text{Sim}_w(p \cap \text{Dox}_{S,c}) >_{S,c} \text{Sim}_w(\neg p \cap \text{Dox}_{S,c})$$

Return to Heim's semantics and consider the obscurity problem. *Orange*, recall, is true in one context and false in another. Let us call the former context OT , and the latter OF . The issue is that Heim's semantics predicts that *Orange* is false in *both* contexts. Frank believes that only dessicated oranges are available: he gets a dessicated orange in every orange-world in both $\text{Dox}_{\text{Frank},OT}$ and $\text{Dox}_{\text{Frank},OF}$. Any such orange-world is *not* an improvement on an otherwise similar no-orange-world—that is, a world where Frank

¹⁰We've taken two liberties here. First, we've translated Heim's dynamic semantics into a static one. Second, Heim uses not $\text{Dox}_{S,c}$ but rather a certain superset of it that screens off what the agent believes about her own intentions. Our use of $\text{Dox}_{S,c}$ makes no difference to our arguments.

gets some other fruit, which (we may assume) he prefers to dessicated oranges—in which case *Orange* is false.¹¹

The intuitive diagnosis of *Orange* is that in *OT*, what’s at issue is juicy oranges, while in *OF*, what’s at issue is dessicated oranges. Heim’s semantics gets this wrong, saying instead that it’s dessicated oranges at issue in both *OT* and *OF*: in each context, the only orange-worlds that figure in the truth conditions are those in Frank’s belief set, which are dessicated-orange-worlds.

In response, one might try to somehow admit juicy-orange-worlds into the truth conditions in *OT*. Following Rubinstein (2017), one can do this by generalizing Heim’s semantics: replace $\text{Dox}_{S,c}$ with a contextually determined domain of worlds, D_c , and say that *only in certain contexts* does D_c equal $\text{Dox}_{S,c}$. (In such contexts the new truth conditions are exactly the same as Heim’s.) What is D_c in other contexts? Rubinstein suggests that it’s a superset of the belief set:

Rubinstein-inspired Semantics

‘S wants p’ is true in c iff

$\forall w \in D_c: \text{Sim}_w(p \cap D_c) >_{S,c} \text{Sim}_w(\neg p \cap D_c)$

Constraint on D_c : $\text{Dox}_{S,c} \subseteq D_c$

Rubinstein’s amendment is partly motivated by “counterfactual” desire reports such as (2):

(2) I want this weekend to last forever. (Heim, 1992, 199)

(2) can be true, but Heim’s account predicts that the report should always suffer from presupposition failure. This is because I know, and thus believe, that weekends don’t last forever, so the set of worlds where the weekend lasts forever intersected with $\text{Dox}_{\text{Me},c}$ is the empty set. But in order to deal with concerns about overgeneration, Heim stipulates that the similarity function is undefined for the empty set. So, on Heim’s semantics the report cannot be true. By contrast, if we assume that D_c includes some worlds where weekends last forever, Rubinstein’s account predicts that (2) will be true, as required.

Although Rubinstein’s amendment might improve on Heim’s account, it doesn’t solve the obscurity problem. Even if there are worlds in D_{OT} where

¹¹Our objection to Heim’s semantics applies just as well to more sophisticated versions of it, like Dandeleit’s (ms) view or Blumberg’s (2018) account of ‘wish’ adapted for ‘want’.

Frank gets a juicy orange, this semantics *still* predicts that *Orange* should be false in *OT*. This is because the similarity function is taken to obey Strong Centering, i.e. if w' is a p -world, then $\text{Sim}_{w'}(p) = \{w'\}$ (Lewis, 1973). Hence, where w' is one of Frank's belief worlds where he gets a (deseccated) orange, the most similar world to w' where he gets an orange will just be w' itself. And as we've seen, Frank does *not* prefer w' to the most similar world to w' in which he gets a different piece of fruit. So, expanding D_{OT} in the way that Rubinstein suggests doesn't help to solve the obscurity problem.¹²

2.2 Decision-theoretic Semantics

Turn now to another species of semantics for 'want': a decision-theoretic one. There are various subspecies—see van Rooij (1999), Levinson (2003), Lassiter (2011), Jerzak (2019), and Phillips-Brown (ms)—and the obscurity problem afflicts them all. To keep the discussion manageable, we spend our time on one endorsed by both van Rooij and Levinson. It goes as follows. Where $\text{ExpVal}_{S,c}$ is a function that takes a proposition and returns the expected value that S assigns to that proposition in c :

Decision-theoretic Semantics

$\lceil S \text{ wants } p \rceil$ is true in c iff $\text{ExpVal}_{S,c}(p) > \text{ExpVal}_{S,c}(\neg p)$

So, 'You want to get a model train set for your birthday' is true just if the expected value that you assign to getting a model train set for your birthday exceeds the expected value that you assign to not getting a model train set for your birthday.

¹²It might be thought that the obscurity problem could be solved by relaxing Rubinstein's constraint that D_c must be a superset of the subject's belief set. But this weakens the explanatory power of the account. Suppose I think that either 1 or 4 people are coming to my dinner party, and that my preference ordering for the number of guests is as follows: $2 > 1 > 4 > 3$. Now consider (3):

- (3) I want an even number of people to come to my dinner party.

Intuitively, (3) is false in this scenario. Suppose that similarity of worlds goes with closeness of guest number. Then Rubinstein's entry delivers the following: for each of my belief worlds w in which a single guest turns up, the closest worlds to w in which an even number of guests turn up *as constrained by my beliefs*, are worlds in which four guests turn up. But if what I want weren't necessarily constrained by my beliefs—i.e. if $\text{DOX}_{Me,c} \not\subseteq D_c$ —then it would be difficult to see why the report should be false.

We can see the Decision-theoretic Semantics as an instance of RAW, since the semantics is entailed by RAW and a decision-theoretic account of the desire relation:

Decision-theoretic Desire Relation

S bears the desire relation to p in c iff $ExpVal_{S,c}(p) > ExpVal_{S,c}(\neg p)$

Back to the obscurity problem. The Decision-theoretic Semantics predicts that *Orange* should be false in *OF*, but not that it should be true in *OT*. Frank believes—he *expects*—that if he gets an orange, it’ll be a dessicated one: the expected value he assigns to getting an orange is the same as the expected value he assigns to getting a dessicated orange. That’s a low expected value, certainly lower than the one you assign to not getting a fruit that’s not a dessicated orange. So, *Orange* is predicted to be false in *OT*.

What could an advocate of the Decision-theoretic Semantics do to make her account predict that *Orange* is true in *OT*? One thing she might try has already been done by Levinson (in response to a different problem): hypothesize that the value function used to calculate expected value varies by context. In one context, the value function represents one value of the agent, and so the expected value function represents one value of the agent’s, while in another context, it represents a different value.¹³ Call this the ‘Multiple-values Hypothesis’.

Decision-theoretic Semantics & Multiple-values Hypothesis

‘ S wants p ’ is true in c iff $ExpVal_{S,c}(p) > ExpVal_{S,c}(\neg p)$, where $ExpVal_{S,c}$ represents the value of S ’s that’s relevant in c .

The Multiple-values Hypothesis is compelling in certain cases. For example: does Otto want to buy that \$1,000 fountain pen? He does and he doesn’t; it would give him great pleasure to write with it, but he knows that spending \$1,000 would be financially irresponsible.

- (4) Otto wants to buy the \$1,000 fountain pen.

In one context, (4) is true, but in another it’s false, and the Multiple-values Hypothesis neatly explains why. Relative to Otto’s value of pleasurable writing, the expected value of buying the pen exceeds that of not buying; things

¹³Levinson is not explicit that it’s different values that the value functions are supposed to represent, but that is how we find it most natural to read him.

are the other way around relative to his value of being financially responsible. In a context where the value function represents Otto's pleasurable-writing value, (4) is true; in a context where the value function represents the financial-responsibility value, (4) is false.

Whatever the merits of the Multiple-values Hypothesis, it doesn't solve the obscurity problem. With *Orange*, the hypothesis would be that relative to one of Frank's values, the expected value he assigns to getting an orange exceeds that of not getting an orange, and *vice versa* relative to another one of his values. This is wrong twice over.

First, what are the relevant values supposed to be? The Multiple-values Hypothesis makes sense with Otto and his pen because it's natural to trace the context-sensitivity of (4) to multiple values of Otto's (pleasurable writing, financial responsibility). Not so in Frank's case, where there looks to be only one value: having pleasurable gustatory experiences. Why is *Orange* true in *OT*? Because Frank enjoys juicy oranges. Why is it false in *OF*? Because Frank doesn't enjoy dessicated oranges. In other words, when *Orange* is true, that's because Frank's value of having pleasurable gustatory experiences is promoted by getting a juicy orange; when *Orange* is false, that's because the very same value isn't promoted by getting a dessicated orange. The Multiple-values Hypothesis hypothesizes, well, multiple values, but with *Orange* there's only one.

Second, the Multiple-values Hypothesis works by saying that in the relevant cases, there is *one* thing at issue, one thing that the agent values in some respect and disvalues in another. For instance, in the case of (4), there is one thing at issue—the proposition that Otto buys the pen—that Otto values in a certain respect and disvalues in another. But there *is no such one thing* in Frank's case; rather, there are two. Indeed, that there is no such one thing is what motivated the obscurity problem in the first place: *Orange* is true in *OT* because Frank looks favorably on getting a juicy orange, while *Orange* is false in *OF* because Frank doesn't look favorably on a different thing, getting a dessicated orange.

2.3 Content-specification Semantics

We've considered and rejected two instances of the Relational Analysis for 'want' (i.e. RAW). We'll now consider and reject a final instance: the Content-specification Semantics, which Fara (2013, p. 252) characterizes as "typically taken to be constitutive of [the Relational Analysis for 'want']" among philosophers.¹⁴ Here is the view:

Content-specification Semantics

⌈*S* wants *p*⌋ is true in *c* iff *S* has a desire in *c* that is satisfied in exactly the worlds where *p* is true.

On this account, 'You want to get a model train set for your birthday' is true just if you have a desire that is satisfied in exactly the worlds where you get a model train set for your birthday.

The Content-specification Semantics is an instance of RAW because it is entailed by the conjunction of RAW and:

Content-specification Desire Relation

S bears the desire relation to *p* in *c* iff *S* has a desire in *c* that is satisfied in exactly the worlds where *p* is true.

What does the Content-specification Semantics predict about *Orange* in *OT* and *OF*? It's rather unclear. On the one hand, one could argue that in neither context does Frank have a desire that is satisfied in *all* the worlds where he gets an orange, since in both contexts the thing that Frank looks favorably on is getting a *juicy* orange. If that's right, the Content-specification Semantics predicts that *Orange* is false in both contexts. On the other hand, one might argue that since Frank looks favorably on having a juicy orange in both contexts, he thereby looks favorably on having an orange *simpliciter* in both contexts. If *that's* right, then the semantics predicts that *Orange* is true in both contexts.

The question of which truth-value the Content-specification Semantics assigns to reports such as *Orange* in contexts such as *OT* (i.e. contexts where

¹⁴Bach (1997) and Shier (1996) also describe this precisification of the Relational Analysis in this way.

such reports are acceptable) is at the heart of what we'll call the "satisfaction underspecification debate". Some, like Fara (2013), maintain that such reports are predicted to be false in these contexts, while others, like Braun (2015), argue that such reports are predicted to be true.¹⁵

The obscurity problem raises an issue for the Content-specification Semantics regardless of where one stands in the above debate. Since *Orange* is acceptable in *OT*, and bears none of the hallmarks of insincere or non-literal speech, it is problematic if *Orange* is predicted to be false in this context. On the other hand, if *Orange* is predicted to be true in *OT*, then it will be true in *OF* as well. But then it is difficult to understand why the report is unacceptable in the latter context. (The only way we see that the Content-specification account could handle this kind of context sensitivity is by saying that Frank's desires *change* from context to context. But they don't.)

A proponent of the Content-specification Semantics might try to respond in the following way. The Content-specification Semantics closely resembles the semantics of Condoravdi and Lauer (2016), who adopt a relative of the Multiple-values Hypothesis (from §2.2): they say that in different contexts, different "sources of [an agent's] preferences" are relevant—where sources of preferences include things such as "inclinations, personal moral codes, and obligations" (p. 28). So suppose we're in a context where the relevant source is the agent's personal moral code. Then, 'You want to shovel snow from your neighbor's driveway' is true just if you want to shovel snow from your neighbor's driveway *and* you want to because of your personal moral code. Call this the 'Multiple-sources Hypothesis'.

Condoravdi and Lauer state their semantics for 'want' in terms of preference and they accordingly talk about preferences having sources. We may suppose that desires have sources, too, and if we do, we can integrate the Multiple-sources Hypothesis into the Content-specification Semantics—which, of course, is stated in terms of desires, not preferences.

¹⁵Braun also sketches a response where the report is strictly speaking false on the Content-specification semantics, but is nevertheless pragmatically acceptable. §3 discusses a close variant of Braun's argument.

Content-specification Semantics & Multiple-sources Hypothesis

‘ S wants p ’ is true in c iff S has a desire in c that is satisfied in exactly the worlds where p is true *and* the source of the desire is the one that’s relevant in c .¹⁶

The Multiple-sources Hypothesis isn’t going to work in Frank’s case; it’s not going to work to say that Frank has a desire with one source directed at having an orange and that he fails to have a desire with a different source directed at having an orange. The reasons are analogous to the two reasons that sunk the Multiple-values Hypothesis. With Frank, there aren’t multiple sources of desire at play, but rather only one—gustatory pleasure. And it’s wrong to think there’s a single thing, having an orange, that Frank at once has one kind of desire towards and fails to have another. There’s a certain thing that Frank looks favorably on, juicy oranges, and there’s a different thing that he looks unfavorably on, dessicated oranges: an adequate theory has to reflect that somehow.

2.4 Summing up

In §1, we said that the problem with mainstream semantics for desire reports traces to an implicit commitment to RAW. You’ve now seen what we mean by this: each of the mainstream semantics is entailed by a distinctive view of the desire relation (e.g. a decision-theoretic one) in conjunction with RAW. The obscurity problem falsifies each of various conjunctions of RAW and a distinctive view of the desire relation. What this reveals, we believe, is that the problem is not with the distinctive views of the desire relation, but with RAW itself.

Thus, we won’t try to develop some new view of the desiring relation and throw out all of the existing ones. Rather, we’ll replace RAW. An adequate replacement for RAW should be able to combine with any number of views of the desire relation and in each case be immune to the obscurity problem.

¹⁶Compare to Condoravdi and Lauer’s semantics, stated in terms that more clearly mirror those of the Content-specification Semantics: ‘ S wants p ’ is true in c iff S “has a preference that is satisfied if and only if p is true” *and* the source of the preference is the one that’s relevant in c (p. 31). (Condoravdi and Lauer formulate their view in two ways. The formulation we’ve just given doesn’t capture the full structure of their view—the full structure not being important to the point we’re making. Find the other formulation on their page 30.)

In §4, we propose such a replacement: the Obscure Object Analysis. Before getting to that, we'll consider another attempt to save RAW from the obscurity problem.

3 Domain restriction

Indefinite descriptions, i.e. expressions such as $\lceil a/an F \rceil$, are standardly treated as a type of quantified noun-phrase on a par with expressions such as $\lceil every F \rceil$, $\lceil no F \rceil$, $\lceil most F \rceil$, etc. (Neale, 1990). Now, it is well known that natural language quantifiers exhibit *implicit domain restriction*. That is, the domain of a quantifier can be restricted in context even if this restriction isn't explicitly expressed (Neale, 1990; Bach, 1994; Stanley, 2000; Recanati, 2002; Elbourne, 2008). Here are some examples:

- (5) [We have just attended a party, and are discussing the night's events.]
- a. i. Every man sang karaoke.
ii. Every man *at the party* sang karaoke.
 - b. i. No woman danced the waltz.
ii. No woman *at the party* danced the waltz.
 - c. i. Most of the couples had fun.
ii. Most of the couples *at the party* had fun.

It's natural to understand utterances of the (i)-sentences as expressing what the (ii)-sentences express. For instance, an utterance of (5a-i) can be heard as saying that every man *at the relevant party* sang, even though the domain of quantification is left implicit. Quantifiers that occur in the complements of attitude reports also exhibit this effect. Assuming the same context as above, for example, it's natural to understand the following (a)-sentences as equivalent to the (b)-sentences:

- (6) a. Amelie thinks every man sang karaoke.
b. Amelie thinks every man *at the party* sang karaoke.
- (7) a. Joey hopes that no woman danced the waltz.
b. Joey hopes that no woman *at the party* danced waltz.
- (8) a. Ollie suspects that most of the couples had fun.

b. Ollie suspects that most of the couples *at the party* had fun.

A friend of the Relational Analysis could point to these observations about domain restriction in responding to the obscurity problem. When we presented the problem, we assumed that the proposition expressed by the complement of ‘want’ in *Orange* (‘Frank wants to have an orange’) denotes the proposition *that Frank gets an orange*. It is this that the response denies; it alleges that our presentation of the problem ignored implicit restriction. More specifically, the response is that what *Orange really* expresses in, for example, the context *OT* is that given in (9):

(9) Frank wants to have a *juicy* orange.

The thought is that the indefinite’s domain of quantification is further restricted in context by what amounts to the material *juicy*.¹⁷ But now there is no problem for the Relational Analysis, since Frank *does* stand in the desire relation to the proposition that he possesses a juicy orange. For instance, on the Rubinstein-inspired Semantics, (9) is *true*: for each world w' in D_{OT} , Frank prefers the world most similar to w' in D_{OT} where he has a juicy orange to the world most similar to w' in D_{OT} where he doesn’t. (We assume that D_{OT} contains juicy orange-worlds.)

Similarly, in the context *OF*, what *Orange* really expresses is that given in (10):

(10) Frank wants to have a *dessicated* orange.

It is straightforward to check that (10) is predicted to be false in *OF* on, e.g. the Rubinstein-inspired Semantics, as required. So, the obscurity problem can be solved by appealing to implicit domain restriction; it has nothing especially to do with desire reports.

In reply, we don’t deny that the domains of quantifiers can be implicitly restricted, but we doubt that this explains what’s going on with *Orange*. The central observation is that the relevant effect is *local* to desire reports: indefinites and other quantifiers occurring outside desire reports cannot be interpreted the same way as those occurring inside of them. For instance, consider (11a)-(13a) in our original context (which we repeat below):

¹⁷This argument parallels one of Braun (2015) defenses of the Content-specification Semantics. (See footnote 15.)

Frank is attending a lunch where each person is given a piece of fruit. Frank thinks he might get an orange. He loves juicy oranges, so he's happy about that. But then one of his friends reports that all the oranges available are dessicated. Frank says, "That's a pity, I really want to have an orange."

- (11) a. # Frank is sure that he will not be given an orange with his lunch.
b. Frank is sure that he will not be given a *juicy* orange with his lunch.
- (12) a. # Frank's friend convinced him that no oranges will be served.
b. Frank's friend convinced him that no *juicy* oranges will be served.
- (13) a. # If Frank gets an orange with his lunch, he'll be happy.
b. If Frank gets a *juicy* orange with his lunch, he'll be happy.

The (a)-sentences are all false in context (as indicated by the '#' preceding each example). However, if the domains of the relevant quantifiers could be restricted as in (9), then we would expect the (a)-sentences to be equivalent to the (b)-sentences, which are *true*.

To bear down on this point some more: to us and our informants, it is almost impossible to interpret the (a)-sentences as equivalent to the (b)-sentences. But if domain restriction was responsible for generating the relevant true reading of *Orange*, then this is surprising; we would expect the same restrictions to be available for the quantifiers in (11a)-(13a). As far as we can see, the only way that the proponent of the domain restriction response could account for this is to maintain that there is something special about desire reports, so that quantifiers occurring inside their scope can be restricted in ways that quantifiers occurring outside of them cannot. But this is completely *ad hoc*.¹⁸

¹⁸In arguing that the Content-specification semantics predicts that reports such as *Orange* are false in contexts such as *OT*, Fara (2013) considers, and rejects, a response that appeals to so-called "nominal restriction", i.e. the phenomenon whereby the interpretation of noun-phrases is implicitly restricted in context (it has been argued that quantifier domain restriction is an instance of nominal restriction (Stanley, 2002)). Fara objects to this response by considering examples such as (14):

- (14) Heidi wants to wear high heels (Fara, 2013, 266).

Fara claims that the desire of Heidi's that makes (14) true is one that is satisfied only if Heidi wears high heels on her feet (and not, e.g. on her hands). So, the Content-specification

To be clear, we think that there is something right about the appeal to implicit restriction. An essential element in this response is that *Orange* is true not because Frank bears the desire relation to the proposition that he has an orange, but rather because he bears the desire relation to the proposition that he has a *juicy* orange. We agree. However, we don't think that this goes by way of implicit restriction, or any other general-purpose mechanism for "enriching" logical forms. Instead, we think that this effect is specific to desire reports, and impacts the standing of the Relational Analysis.

4 Desire's obscure object

The Relational Analysis posits a *transparent* link between language and mind: the agent is ascribed an attitude whose object can be seen simply by looking at the grammatical object of the attitude verb. However, we have argued that the Relational Analysis is false. More specifically, we have argued that the Relational Analysis as applied to 'want' is false:

Relational Analysis for 'want' (RAW) (repeated from page 3)

- (i) $\lceil S \text{ wants } p \rceil$ is true in c iff S bears the relation expressed by 'want' to p in c ;
- (ii) S bears the relation expressed by 'want' to p in c iff S bears the desire relation to p in c .

Semantics predicts that (14) will be false, since Heidi has no desire that is satisfied in exactly the worlds where Heidi wears high heels (since some of these worlds will be ones where Heidi wears high heels on her hands). This is supposed to be a problem for the nominal restriction response since there is no suitable noun-phrase in (14) that can be restricted to yield the desired interpretation: 'Heidi wants to wear high heels *on her feet*' doesn't involve nominal restriction; instead it involves restricting the verb 'wear'. Thus, Fara claims that the nominal restriction response isn't general enough.

We agree with Fara that (14) can be true when all Heidi desires is to wear high heels on her feet. But we don't think that this constitutes a compelling argument against the appeal to nominal restriction. That's because the sort of restriction on verbs exhibited by 'wear' in (14) is also exhibited by verbs occurring outside of desire reports. For instance, if exactly four of the five women who attended the gala wore high heels on their feet, I can truthfully utter 'Not all of the women wore high heels', even if the fifth woman used a pair of high heels as mittens. Or if none of the women wore high heels on their feet, but one wore high heel-mittens, I can truthfully utter 'None of the women wore high heels'. So, it's no mark against the appeal to nominal restriction that it cannot account for the verb restriction in (14), since this phenomenon appears to be quite general. A more compelling argument against the restriction strategy comes from cases where the purported restriction effects appear to be local to desire reports. This is exactly the type of argument that we have presented in the main text.

If RAW is false, then what exactly is the relationship between language and mind—what do desire reports tell us about the subject’s desires? In this section, we propose an answer to this question.

Our proposal is guided by two ideas. The first is that the objects to which agents stand in the desire relation are *underspecified* by desire reports. When a report ‘ S wants p ’ is true, this needn’t be because S bears the desire relation to p . Instead, it can be true because S bears the desire relation to a proposition q that entails p .

To illustrate, this explains how *Orange* can be true in *OT*: Frank bears the desire relation to the proposition that he gets a *juicy* orange—a proposition that’s more “specific” than the proposition that he gets an orange. The proposition expressed by the complement in *Orange* doesn’t fully specify the object to which Frank bears the desire relation. Nevertheless, the report is true, since this proposition is entailed by something to which Frank *does* stand in the desire relation.

This first idea can be implemented by either rejecting (i) in RAW and keeping (ii), or by rejecting (ii) and keeping (i) (our preferred option). We’ll consider these options in turn.

4.1 Rejecting (i) in RAW

The first account we’ll consider is the following:

Analysis 1

- (i) ‘ S wants p ’ is true in c iff there is a q such that:
 - (a) S bears the relation expressed by ‘want’ to q in c ;
 - (b) q entails p .
- (ii) S bears the relation expressed by ‘want’ to q in c iff S bears the desire relation to q in c .

Analysis 1 captures the intuitive reason why *Orange* is true in *OT*: there is a proposition, the proposition that Frank has a juicy orange, to which Frank bears the desire relation, and which entails the proposition expressed by the prejacent of *Orange*. In other words, Analysis 1 allows *Orange* to be true by virtue of the agent (Frank) standing in the desiring relation to a proposition

(having a juicy orange) that's not fully specified by the prejacent of *Orange* (having an orange).

Although Analysis 1 captures our first guiding idea, we are reluctant to reject clause (i) in RAW, as Analysis 1 does. An important consideration stems from intuitively valid natural language inferences such as the following (Schiffer, 2003, ch.2):

Schiffer Inference

(P1) Patricia wants the Seahawks to win the Super Bowl.

(P2) Petra wants the Seahawks to win the Super Bowl.

(C) So, there is something that both Patricia and Petra want.

The Schiffer Inference seems unimpeachable. But Analysis 1 doesn't validate it: (P1) could be true because Patricia bears the desire relation to a proposition q that entails that the Seahawks win the Super Bowl; and (P2) could be true because Petra bears the desire relation to a *different* proposition r that entails that the Seahawks win the Super Bowl. In this case, there needn't be any proposition to which both Patricia and Petra stand in the desiring relation, and so—by clause (ii)—no proposition to which they both stand in the relation expressed by 'want'.

4.2 Rejecting (ii) in RAW

To validate the Schiffer Inference, we propose to reject (ii) in RAW:

Analysis 2 (to be revised)

- (i) $\ulcorner S$ wants $p \urcorner$ is true in c iff S bears the relation expressed by 'want' to p in c ;
- (ii) S bears the relation expressed by 'want' to p in c iff there is a q such that:
 - (a) S bears the desire relation to q in c ;
 - (b) q entails p .

As with Analysis 1, Analysis 2 captures our first guiding idea: *Orange* is true in *OT* because Frank stands in the desire relation to the proposition that he gets a juicy orange.

As promised, Analysis 2 validates the Schiffer Inference: if (P1) is true, then, by (i), Patricia stands in the relation expressed by ‘want’ to the proposition that the Seahawks win the Super Bowl; if (P2) is true, then by (i) Petra also stands in relation expressed by ‘want’ to that same proposition; so, there is something that both Patricia and Petra want, which is to say that (C) is true.

However, Analysis 2 doesn’t solve the obscurity problem. That’s because it predicts that *Orange* should be true in *both* *OT* and *OF*. On Analysis 2, *Orange* is true in *c* just if there’s a proposition *q* to which Frank bears the desire relation in *c*, and *q* entails that Frank gets an orange. There is, of course, such a proposition in both *OT* and *OF*—it’s the proposition that Frank gets a juicy orange. But of course the truth value of *Orange* varies by context.

This brings us to the second guiding idea behind our proposal. The first idea, recall, is that a desire report may be true (or false) because the agent stands (or doesn’t stand) in the desire relation to a proposition that’s more specific than the prejacent of the report. The second idea is that there is contextual variation in *which* more specific proposition or propositions are relevant to evaluating a desire report. In the context where *Orange* is true, it’s true because Frank bears the desire relation to the proposition that he gets a juicy orange, while in the context where *Orange* is false, it’s false because Frank doesn’t bear the desire relation to the proposition that he gets a dessicated orange.

To capture this context-sensitivity, we propose a view just like Analysis 2—thereby (correctly) validating the Schiffer Inference—except that now the existential quantification over propositions from clause (ii) of Analysis 2 is *contextually restricted*. Specifically, context determines a domain of propositions Γ_c over which the existential quantifier ranges.

Obscure Object Analysis

- (i) ‘ S wants p ’ is true in c iff S bears the relation expressed by ‘want’ to p in c ;
- (ii) S bears the relation expressed by ‘want’ to p in c iff there is a q in Γ_c such that:
 - (a) S bears the desire relation to q in c ;
 - (b) q entails p .

The view works like this. Γ_{OT} contains the proposition that Frank gets a juicy orange. And so since the Obscure Object Analysis embodies the first guiding idea, it predicts that *Orange* is true in *OT* because Frank bears the desire relation to the proposition that he gets a juicy orange.

By contrast, Γ_{OF} does *not* contain the proposition that Frank gets a juicy orange. Indeed, Γ_{OF} contains no proposition q to which Frank bears the desire relation and that entails that Frank gets an orange. Γ_{OF} could, for instance, contain only the proposition that Frank gets a dessicated orange. So, the Obscure Object Analysis predicts that *Orange* is false in *OF* because Frank doesn't stand in the desire relation to the proposition that he gets a dessicated orange.

We won't try to resolve exactly how Γ_c gets determined in a given context, with one exception: in evaluating $\lceil S \text{ wants } p \rceil$ in c , Γ_c always contains p .¹⁹ This ensures that whenever S bears the desire relation to p in c , they can be truly reported as such, i.e. $\lceil S \text{ wants } p \rceil$ is true in c .

Clearly, in order for our account to be considered a complete theory, much more needs to be said about how exactly the value of Γ_c is fixed in context. This is an important issue, and as long as it goes unaddressed, there will be a major lacuna in the Obscure Object Analysis. However, we won't try to fill this gap here since our aim in this paper is relatively modest. We have tried to outline the general shape that a solution to the obscurity problem must take, rather than provide a complete account. And we are confident that virtually any adequate solution will appeal to something very much like the parameter Γ_c , even if specifying how exactly this parameter is determined is a non-trivial task.

4.3 Mainstream semantics, again

In §2, we showed that the obscurity problem falsifies various mainstream semantics, each of which implicitly couples RAW with a distinctive view of the desire relation. And we said that any adequate replacement for RAW should be able to combine with various different views of the desire relation and in

¹⁹This “prejacent sensitivity” that we propose for Γ_c parallels constraints endorsed by others, e.g. Lassiter (2011) and Mandelkern et al. (2017). For instance, in the course of providing a semantics for desire reports, Lassiter claims that $\lceil S \text{ wants } p \rceil$ is evaluated against a contextually determined set of alternative propositions, which always contains p .

each case evade the obscurity problem.

The Obscure Object Analysis fits this bill. Each of the semantics from §2 is entailed by RAW and a distinctive view of the want relation. Similarly, each of these views of the wanting relation, in conjunction with the Obscure Object Analysis, entails a variant of each of the semantics from §2. The spirit of the original versions of the mainstream semantics is preserved—at least to a good extent—by these variants that, unlike the mainstream semantics themselves, don’t suffer from the obscurity problem.

We’ll illustrate this by way of example: a decision-theoretic view of the desire relation combines with the Obscure Object Analysis to yield a decision-theoretic semantics that gets obscurity cases right.

Recall the Decision-theoretic Desire Relation:

Decision-theoretic Desire Relation (repeated from page 11)

S bears the desire relation to p in c iff $ExpVal_{S,c}(p) > ExpVal_{S,c}(\neg p)$

This view and the Obscure Object Analysis entail:

Obscure Object Decision-theoretic Semantics

$\ulcorner S$ wants $p \urcorner$ is true in c iff for some proposition q in Γ_c :

(a) $ExpVal_{S,c}(q) > ExpVal_{S,c}(\neg q)$;

(b) q entails p .

To illustrate: *Orange* is true in *OT* because Frank stands in the Decision-theoretic Desire Relation to the proposition that he has a juicy orange—which, recall, is a member of Γ_{OT} . The expected value Frank assigns to having a juicy orange exceeds the expected value he assigns to not having a juicy orange.²⁰ Similarly, *Orange* is false in *OF* because Frank does not stand in

²⁰There is a slight wrinkle here. We said when introducing the case that Frank is certain that he won’t get a juicy orange, which means, by common definitions of expected value, that the expected value that Frank assigns to getting a juicy orange is *undefined* (so it doesn’t exceed the expected value of not getting a juicy orange!). The problem, stated generally, is that one can stand in the desire relation to a proposition to which one is certain is false, as Frank does to the proposition that he gets a juicy orange. (Another example of this is was (2) (‘I want this weekend to last forever’) discussed in §2.1.)

Happily, there’s a fix. Wrenn (2010) gives a decision-theoretic account of wanting that uses not expected value but rather what he calls *extended* value to deal with cases where an agent wants something she’s certain won’t happen. While expected value is built out of a probability function (and value function), extended value is built out of a Stalnake-

the Decision-theoretic Desire Relation to the proposition that he has a dessicated orange—which, we may suppose, is the only member of Γ_{OF} . The expected value Frank assigns to having a dessicated orange doesn't exceed the expected value he assigns to not having a dessicated orange.

There are two things to note here. First: not only does the Obscure Object Decision-theoretic Semantics make the right predictions in *OT* and *OF*, the explanations for those predictions look (unsurprisingly) just like the ones we gave for Obscure Object Analysis predictions. That is, the truth value of *Orange* varies by context: when it's true, that's because of how Frank feels about juicy oranges; when it's false, that's because of how Frank feels about dessicated oranges.

Second: the Obscure Object Decision-theoretic Semantics preserves the spirit of the original Decision-theoretic Semantics—at least to a good extent—since both semantics have at their core the Decision-theoretic Desire Relation. So, although the obscurity problem poses a problem for the Decision-theoretic Semantics from §2, a close relative of that view can be salvaged once the Obscure Object Analysis is taken on board. We have the 'at least to a good extent' and 'close relative' riders in the above because the Decision-theoretic Semantics and the Obscure Object Decision-theoretic Semantics do differ in certain logical properties, as we bring out in the next section.

4.4 A problem for another day: monotonicity

We'll end this section by discussing the following feature of The Obscure Object Analysis: it confers upon desire reports a controversial logical property: (upward) *monotonicity*. That is, if (i) $\lceil S \text{ wants } q \rceil$ is true in c and (ii) q entails p , then (iii) $\lceil S \text{ wants } p \rceil$ is true in c . Note that our proposal validates monotonicity no matter how it is parametrized. In particular, no matter which conception of the desiring relation one endorses, a semantics based on The Obscure Object Analysis will validate monotonicity. This is worth remarking on since none of the mainstream semantics from §2 validate monotonicity. So,

rian (1970) *extended* probability function (and value function). The extended probability to which you assign a proposition may be defined if you're certain that the proposition is false, and in turn the extended value that you assign to a proposition may be defined even if you're certain that the proposition is false—which is exactly what we need in Frank's case. If we swap in extended value in place for expected value in the Obscure Object Decision-theoretic Semantics, then the problem we've discussed here can be solved.

moving to The Obscure Object Analysis has non-trivial logical consequences.

Why is monotonicity controversial? On the one hand, there are cases that support it. For instance, consider an example from Crnič (2011, 172) (building on (von Fintel, 1999)):

Dice: John is playing a game of dice. He bets his entire fortune that the sum of the two dice that are thrown will be either less than four, in which case he gets his fortune back, or more than eight, in which case he doubles his fortune. In case the sum of the dices is between four and eight he loses everything.

- (15) # John wants the dice to sum up to more than eight, but he doesn't want the dice to sum up to more than three.

(15) is non-nonsensical, and if 'want' is monotonic, we have an explanation ready to hand: the first conjunct of (15) entails the negation of the second.

However, as many have noted, there are also cases that suggest desire reports are not monotonic.²¹ For instance, the inference in (16) is invalid:²²

- (16) a. Camila wants to die peacefully
b. \nRightarrow Camila wants to die.

Theorists are split on monotonicity. Some, like Heim (1992), explicitly disavow it. These theorists then need to explain why examples such as (15) are unacceptable.²³ As far as we are aware, no such explanation has been proposed. Others, like von Fintel (1999), attempt to preserve monotonicity in a restricted form. That is, they attempt to invalidate the inferences in (16) while "preserving" monotonicity in the context where we evaluate (15).

One such attempt has been to take a view on which desire reports are monotonic, and add the following, widely endorsed constraint: $\lceil S \text{ wants } p \rceil$ cannot be true if S believes p .²⁴ The inference in (16) rightly doesn't go through: Camila believes that she'll die, in which case (16b) is not true, even if (16a) is true. And (15) is rightly predicted non-nonsensical: in John's case,

²¹For example, (Asher, 1987; Stalnaker, 1984; Heim, 1992).

²²The sentence is adapted from (Anand and Hacquard, 2013).

²³Granted plausible assumptions, it can be checked that (15) is predicted to be true on both Heim's account and the Decision-theoretic Semantics.

²⁴The constraint is endorsed, for example, by (Heim, 1992; von Fintel, 1999; Crnič, 2011; Rubinstein, 2012).

the constraint is inert—John doesn’t believe that the dice will sum up to more than three—and so the constraint doesn’t interfere with the monotonic inference from the first conjunct of (15) to the negation of the second. This “belief constraint” could easily be added to our (monotonic) account, just as it’s been added to von Fintel’s (monotonic) account.

This isn’t a satisfying response, though, because the constraint is wrong, as various authors have noted.²⁵ ‘ $\lceil S \text{ wants } p \rceil$ can be true even if S believes p . In some contexts, even a relative of (16b) can be true, if sad:

(17) Arlo believes he’s going to die, and indeed that’s what he wants.

Or consider the following from (Grano and Phillips-Brown, ms):

(18) a. Francisca: I’m hungry.

b. Bert: Good! I want you to be hungry; I’m about to cook dinner.

To sum up: some can explain why (16) is bad, while others, including us, can explain why (15) is bad. However, no one has a satisfying explanation for why *both* are bad. It’s well beyond the scope of this paper for us to try to resolve this issue. We have discussed it mainly so that our proposal may be more easily situated in the existing literature. Our analysis validates monotonicity, so we need an account of why monotonicity inferences are sometimes unacceptable. We must leave this for future work.

5 Conclusion

The Relational Analysis as applied to ‘want’, RAW, posits an intuitively pleasing, transparent connection between language and mind: according to RAW, the object of the desire relation can be seen simply by looking at the grammatical object of the desire report. But things are not transparent; rather, they are obscure. Desire reports say rather less about our desires than RAW dictates,²⁶ a fact we’ve brought out by noting a previously unappreciated kind

²⁵For example, Iatridou (2000).

²⁶Blumberg and Lederman (2020) suggest that true attitude reports sometimes say *more* about a subject’s mental life than is actually the case. They base their arguments on reports such as (19) in contexts such as *Tennis* (Blumberg and Holguín, 2018):

Tennis: Ann is a six-year-old girl whom Petra, an expert in tennis pedagogy, has never met and whose existence he is unaware of. Petra believes that every six-year-old can learn to play tennis in ten lessons. Jane, Ann’s aunt, is aware of Petra’s feelings on the matter. Jane wants to encourage Ann’s father, Jim,

of context-sensitivity in desire reports.

This obscurity and context-sensitivity, we've shown, pose a problem for both RAW and what can be seen as its instances: mainstream semantics for desire ascriptions. We've proposed to solve the obscurity problem by rejecting RAW and putting in its place what we call the *Obscure Object Analysis*. We've also shown that versions of mainstream semantics can be recovered on the Obscure Object Analysis.

As we have noted, there is more work to be done. For one thing, we need a fuller metasemantic theory of how exactly the values of the context sensitive elements that we've posited are fixed. For another, we need to get clearer on the various issues surrounding monotonicity. However, we believe that this work is worth taking on, since the Obscure Object Analysis provides us with a better picture of how desire reports work.

to sign Ann up for tennis lessons, so in conversation with Jim she asserts the following:

(19) Petra believes Ann can learn to play tennis in ten lessons.

Blumberg and Lederman claim that (19) is true in context, and take this to show that 'believe' doesn't express the belief relation. Instead, the relation expressed by 'believe' is obtained by "revising" the belief relation with additional information, namely the proposition that Ann is a six-year-old. Blumberg and Lederman's conclusion is very much in keeping with our own, i.e. that attitude reports aren't transparent in the relevant sense. However, it should be clear that this phenomenon of "revision" is distinct from the obscurity problem. So, even if we have come to the same ultimate conclusion as Blumberg and Lederman, our motivations and arguments are quite different.

References

- Anand, P. and Hacquard, V. (2013). Epistemics and attitudes. *Semantics and Pragmatics*, 6(8):1–59.
- Asher, N. (1987). A typology for attitude verbs and their anaphoric properties. *Linguistics and Philosophy*, 10(2):125–197.
- Bach, K. (1994). Conversational implicature. *Mind and Language*, 9(2):124–162.
- Bach, K. (1997). Do Belief Reports Report Beliefs? *Pacific Philosophical Quarterly*, 78(3):215–241.
- Ben-Yami, H. (1997). Against characterizing mental states as propositional attitudes. *Philosophical Quarterly*, 47(186):84–89.
- Blumberg, K. (2018). Counterfactual attitudes and the relational analysis. *Mind*, 127(506):521–546.
- Blumberg, K. and Holguín, B. (2018). Ultra-liberal attitude reports. *Philosophical Studies*, 175(8):2043–2062.
- Blumberg, K. and Lederman, H. (2020). Revisionist reporting.
- Braun, D. (2015). Desiring, desires, and desire ascriptions. *Philosophical Studies*, 172(1):141–162.
- Collins, C. and Postal, P. M. (2014). *Classical NEG Raising: An Essay on the Syntax of Negation*. MIT press.
- Condoravdi, C. and Lauer, S. (2016). Anankastic conditionals are just conditionals. *Semantics and Pragmatics*, 9(8):1–69.
- Crane, T. (2001). *Elements of Mind: An Introduction to the Philosophy of Mind*. Oxford: Oxford University Press.
- Crnič, L. (2011). *Getting even*. PhD thesis, MIT.
- Dandele, S. (ms). Partial desires, blinkered beliefs. Unpublished manuscript.
- den Dikken, M., Larson, R., and Ludlow, P. (1996). Intentional “transitive” verbs and concealed complement clauses. *Revista De Linguistica*, 8.
- Elbourne, P. (2008). The argument from binding. *Philosophical Perspectives*, 22(1):89–110.
- Fara, D. G. (2003). Desires, scope, and tense. *Philosophical Perspectives*, 17(1):141–163. Originally published under the name ‘Delia Graff’.
- Fara, D. G. (2013). Specifying desires. *Noûs*, 47(2):250–272.
- von Fintel, K. (1999). Npi licensing, strawson entailment, and context dependency. *Journal of Semantics*, 16(2):97–148.
- von Fintel, K. and Heim, I. (2011). *Intensional Semantics*. MIT unpublished class notes.
- Geurts, B. (1998). Presuppositions and anaphors in attitude contexts. *Linguistics and Philosophy*, 21(6):545–601.
- Grano, T. and Phillips-Brown, M. (ms). (counter)factual want ascriptions.
- Grant, L. and Phillips-Brown, M. (2019). Getting what you want. *Philosophical Studies*.
- Grzankowski, A. (2012). Not all attitudes are propositional. *European Journal of Philosophy*, (3):374–391.
- Heim, I. (1992). Presupposition projection and the semantics of attitude verbs. *Journal of Semantics*, 9(3):183–221.
- Iatridou, S. (2000). The grammatical ingredients of counterfactuality. *Linguistic*

- Inquiry*, 31(2):231–270.
- Jerzak, E. (2019). Two ways to want? *Journal of Philosophy*, 116(2):65–98.
- Larson, R. (2002). The grammar of intentionality. In Peter, G. and Preyer, G., editors, *Logical Form and Language*, pages 228–62. Oxford University Press.
- Lassiter, D. (2011). *Measurement and Modality: the Scalar Basis of Modal Semantics*. PhD thesis, New York University.
- Levinson, D. (2003). Probabilistic model-theoretic semantics for 'want'. *Semantics and Linguistic Theory*, 13(0).
- Lewis, D. K. (1973). *Counterfactuals*. Blackwell.
- Lycan, W. G. (2012). Desire considered as a propositional attitude. *Philosophical Perspectives*, 26(1):201–215.
- Lycan, W. G. (ms). In what sense is desire a propositional attitude? Unpublished manuscript.
- Mandelkern, M., Schultheis, G., and Boylan, D. (2017). Agentive modals. *Philosophical Review*, 126(3):301–343.
- Montague, M. (2007). Against propositionalism. *Noûs*, 41(3):503–518.
- Neale, S. (1990). *Descriptions*. MIT Press.
- Phillips-Brown, M. (2018). I want to, but.... In Truswell, R., Cummins, C., Heycock, C., Rabern, B., and Rohde, H., editors, *Sinn und Bedeutung 21*, pages 951–968.
- Phillips-Brown, M. (ms). What does decision theory have to do with wanting?
- Prinz, J. (ms). No satisfaction? the mundane truth about desires. Unpublished manuscript.
- Recanati, F. (2002). *Literal Meaning*. Cambridge University Press.
- Rubinstein, A. (2012). *Roots of Modality*. PhD thesis, University of Massachusetts Amherst.
- Rubinstein, A. (2017). Straddling the line between attitude verbs and necessity modals. In *Modality Across Syntactic Categories*. Oxford University Press.
- Schiffer, S. R. (2003). *The Things We Mean*. Oxford University Press.
- Shier, D. (1996). Direct Reference for the Narrow Minded. *Pacific Philosophical Quarterly*, 77(3):225–248.
- Soames, S. (1987). Direct reference, propositional attitudes, and semantic content. *Philosophical Topics*, 15(1):47–87.
- Stalnaker, R. (1970). Probability and conditionals. *Philosophy of Science*, 37(1):64–80.
- Stalnaker, R. (1984). *Inquiry*. Cambridge University Press.
- Stalnaker, R. (1988). Belief attribution and context. In Grimm, R. H. and Merrill, D. D., editors, *Contents of Thought*, pages 140–156. University of Arizona Press.
- Stanley, J. (2000). Context and logical form. *Linguistics and Philosophy*, 23(4):391–434.
- Stanley, J. (2002). Nominal restriction. In Peter, G. and Preyer, G., editors, *Logical Form and Language*, pages 365–390. Oxford University Press.
- van Rooij, R. (1999). Some analyses of pro-attitudes. In *Logic, Game Theory and Social Choice*. Tilburg University Press.
- Villalta, E. (2008). Mood and gradability: An investigation of the subjunctive mood in spanish. *Linguistics and Philosophy*, 31(4):467–522.
- Wrenn, C. (2010). A puzzle about desire. *Erkenntnis*, 73(2):185–209.