

De Re Belief and Cumming's Puzzle*

James R. Shaw

Draft, please don't cite

Cumming (2008) uses a puzzle about belief ascription to argue against a Millian semantics for names, and in favor of a view on which names are assigned denotations relative to a shiftable variable assignment. I want to take the puzzle Cumming raises as an opportunity to showcase the virtues of a rival, broadly Stalnakerian, framework for attitude reports that safeguards Millianism.

I begin by describing the puzzle, and arguing that Cumming's proposed resolution fails. A lesson of this discussion is that the puzzle shouldn't be addressed by appeal to shiftable parameters of an index. As such, the puzzle seems to militate in favor of a Millian view, not against it. After arguing for this claim, I defend three theses needed to apply a Stalnakerian apparatus to the case, arguing that it not only accounts for the data, but harmoniously integrates the puzzling reports into a general framework for dealing with all *de re* intensional constructions.


1 Cumming's Puzzle

Cumming's puzzle is set in the context of a would-be Shakespearean prank.

Rosaline, Maria, and Katherine are going to a masked ball. Each lady anticipates (tepidly) the attentions of a different suitor, and each suitor has given his lady a favor to wear, by which he hopes to recognize her during the masque. The ladies, with malice aforethought, decide to swap favors and so cause the suitors to woo contrary. The exchange of favors in this (counter-fictional) version of *Love's Labour's Lost* is represented in [the figure below].¹

*For helpful comments and advice I'm grateful to Jacob Beck, Dilip Ninan, Anders Schoubye, audience members at Texas Tech and, especially, to Bryan Pickel.

¹Cumming (2008) p.529.



Provided Biron, one of the suitors, is duped, the following reports seem true on their most natural interpretations.

(KR) Biron thinks Katherine is Rosaline.

(RM) Biron thinks Rosaline is Maria.

(MK) Biron thinks Maria is Katherine.

Whereas the following reports, interpreted in parallel fashion, are false.

(RK) Biron thinks Rosaline is Katherine.

(MR) Biron thinks Maria is Rosaline.

(KM) Biron thinks Katherine is Maria.

Biron doesn't seem to have illogical or incoherent beliefs, so the puzzle is to explain why the identities believed appear to violate symmetry and transitivity.²

Cumming's resolution begins by noting a tempting parallel between these cases and belief ascriptions involving definite descriptions. Suppose, for example, that the ladies have entered the room in this order

Katherine, Rosaline, Maria

when, in keeping with Biron's confusion, he believes them to have entered in this order.

Rosaline, Maria, Katherine

Then (FL) can come out true while (LF) does not (again on natural, 'parallel' readings).

²It's worth noting, as Cumming does, that although a natural reading of (RK) is false, there is another reading of it that is true (*ibid.* p.530 n.12). This second reading is perhaps easiest to get when framed with the right kinds of question. I might say: "I remember that Biron thinks Rosaline is here somewhere. Who does Biron think Rosaline is?" A natural and seemingly truthful response can be "Biron thinks Rosaline is Katherine". Shifts in intonation can influence whether we can, or cannot, get these 'inverted' readings for each of (KR)–(KM). This is another datum we need to explain. I will also be following Cumming in taking the "is" here to be the "is" of identity (Cumming notes the uses fail diagnostics for the predicative copula (*ibid.* p.530 n.10)). But I suspect worries about coherence, analogous to those raised with identity, could be pressed for a predicative use.

- (FL) Biron thinks the person who entered first is the person who entered last.
- (LF) Biron thinks the person who entered last is the person who entered first.

A familiar treatment posits something like a scope ambiguity in the logical forms of (FL) and (LF).³ The initial descriptions tend to be read *de re* (that is, outside the scope of “believes”) and the second descriptions are read *de dicto* (inside the scope of “believes”). So the logical forms might look something like the following (say), if we treat definite descriptions as binary generalized quantifiers.

- (FL') $(\lambda x : x \text{ entered first})(\text{Bel}_{\text{Biron}}((\lambda y : y \text{ entered last})(x = y)))$
- (LF') $(\lambda y : y \text{ entered last})(\text{Bel}_{\text{Biron}}((\lambda x : x \text{ entered first})(x = y)))$

The importance of the difference is that the denotations of definite descriptions are sensitive to a shiftable parameter of the index: the world parameter. “Believes” is associated with an intensional operator which shifts the world parameter of the index, so definite descriptions in its scope can change their denotations. “The person who entered first” picks out Kathryn in the world of the story. But in the worlds compatible with what Biron believes, that description picks out not Kathryn, but Rosaline. Instead, it is “the person who entered last” which picks out Kathryn in the worlds compatible with Biron’s beliefs. This is why (FL') is true: it asserts an identity between Kathryn, picked out in one way in the actual world, and Kathryn, picked out in a different way in Biron’s belief-worlds. (LF') on the other hand falsely asserts an identity between Maria and Rosaline, each again picked out in different worlds.

The moral Cumming draws from (KR)–(KM) is that names should share the semantic feature of definite descriptions that allows for this ambiguity: we should allow the denotation of a name to be sensitive to some parameter which the belief operator can shift. Of course it needn’t, and perhaps shouldn’t, be the world parameter that plays this role for names. Cumming accepts, and helpfully supplements, the classic arguments of Kripke (1980) that the denotations of names are constant from world to world. Instead, he defends the view that we take the denotation of a name to be sensitive to a separate shiftable parameter of the index: a variable assignment parameter. If the denotations of names are shiftable in this way, we rule out a prevalent Millian position, on which the semantic contribution of a name is exhausted by its actual referent, leaving no room for sensitivity in the semantic value of a name to *any* parameter.

Let me, very briefly, sketch some important elements of Cumming’s particular proposal. First, according to Cumming, names behave syntactically like variables—a position he calls “variabilism”. So the logical form of (1) is roughly given by (1'), where v_k is the variable associated with the name “Katherine”.

³This is a slight oversimplification since accounting for the ‘scope paradoxes’ of Fodor (1970) may require more flexibility in readings than allowed by position in logical form—perhaps accomplished by introducing bindable world variables. But for the purposes of this paper, this technicality can safely be set aside.

(1) Katherine is tall.

(1') Tv_k

A context c supplies a variable assignment g_c mapping variables to objects. It may, for example, assign v_k to Katherine. Unbound names have their denotations fixed by this assignment. So the truth of (1) uttered in context c turns on whether $g_c(v_k)$ (i.e., Katherine) is tall.⁴

Next we introduce some complexity in attitude states. As already alluded to, attitude states like belief are often taken to encapsulate information about a set of metaphysically possible worlds—those worlds compatible with the belief state. By contrast, Cumming takes belief states to encapsulate information not only about what the world is like, but how variables are assigned to objects. Very roughly, these latter assignments give information about ‘the reference relation as Biron sees it’ at the relevant belief worlds.⁵ It can, for example, encapsulate information about who Biron takes uses of “Katherine” (v_k) to refer to at his belief worlds. Biron’s total belief state is accordingly encapsulated by a set of world/variable-assignment pairs (pairs of the form $\langle w, g \rangle$).

Recall that Cumming wants to resolve his puzzle by suggesting that there are two readings of names in attitude report contexts—what he terms “*de dicto*” and “*de re*” readings. The logical form of the first, *de dicto** reading of a report like (2) is given by (2'), where v_b is contribution of “Biron”.

(2) Biron thinks Katherine is tall.

(2') $Bel_{v_b}(Tv_k)$

In (2') the belief operator binds the variable v_k . Assuming the contextual variable assignment g_c assigns v_b to Biron, (2') is true at c just in case in for all world-assignment pairs $\langle w, g \rangle$ compatible with what Biron believes, $g(v_k)$ is tall in w . Roughly, but more colloquially, it's true if in every world compatible with what Biron believes, and for every reference relation Biron takes to be compatible with that world, the referent of uses of “Katherine” (v_k) at that world is tall. Note that since a given thinker can misconstrue the reference relation in all sorts of ways, there is no logical guarantee that the variable assignments compatible with a believer’s belief state always assign v_k to Katherine.

An important feature of this proposal is that coreferring names read *de dicto** in attitude report contexts are not generally substitutable *salva veritate*. “Hesperus” and “Phosphorus” may contribute different variables to logical form. If these names actually corefer, the contextual variable assignment will reflect this, so that these names are intersubstitutable in extensional contexts. But Biron may not recognize that they corefer, so that the variable assignments compatible with his beliefs assign different objects to the two variables at some of Biron’s belief-worlds. In this way Biron may coherently believe, for example, that Hesperus_{*de dicto**} is bright without believing Phosphorus_{*de dicto**} is bright.

⁴Truth is relativized not only to a variable assignment, but to a world and a model. I'll suppress both relativizations for simplicity in what follows.

⁵*ibid.* p.531.

In addition to the foregoing *de dicto** reading of names, Cumming introduces a second, *de re** reading. To obtain it, Cumming exploits a technique for quantifying into intensional constructions owing to Kaplan (1968). First, Cumming introduces ‘metavariables’ into the syntax of the language, which I’ll denote with greek letters (α, β, \dots) . Quantifiers binding these variables behave substitutionally, with metavariables taking ordinary variables as their values. So, for example, (3) is true just in case there is some (ordinary, non-meta-)variable, call it “ v_a ”, such that (3’) is true.

$$(3) (\exists\alpha)\Phi(\alpha)$$

$$(3') \Phi(v_a)$$

The *de re** reading of the name in (2) is then given by (2’’).

$$(2'') (\exists\alpha)(\alpha = v_k \wedge Bel_{v_b}(T\alpha))$$

Uttered in a context c this says, roughly, that there is some variable which actually (and according to the contextual variable assignment g_c) corefers with uses of “Katherine”, and which Biron takes to refer to someone who is tall.

Note that since the variable v_k in (2’’) lies outside the scope of the belief operator, it can now be freely substituted *salva veritate* with coreferring terms (e.g., other variables that actually, and relative to g_c , refer to the same object). If (2’’) is true, this is because there is some variable, say v_x , that actually corefers with “Katherine” (v_k) according to the contextual variable assignment, and such that Biron thinks v_x refers to someone tall. Suppose this is true, and let N be any name (actually, relative to g_c) referring to Katherine. Then there will again be some variable coreferring with N (namely, the earlier v_x) which Biron takes to refer to someone tall. So the *de re** reading of (2) with N substituted for “Katherine” will preserve truth.

With these readings in place, (KR) and (RK) are handled by reading the first occurring names *de re**, and the second *de dicto**. So the logical forms of those readings are (KR’) and (RK’).

$$(KR') (\exists\alpha)(\alpha = v_k \wedge Bel_{v_b}(\alpha = v_r))$$

$$(RK') (\exists\alpha)(\alpha = v_r \wedge Bel_{v_b}(\alpha = v_k))$$

There are two salient virtues of this proposal (among others). First, (KR’) doesn’t entail (RK’). (KR’) may be true because there is some variable referring to Katherine *other* than that contributed by “Katherine”, that Biron takes to corefer with “Rosaline” (Cumming suggests: an “internal deixis trained on Katherine”⁶). If this holds, there needn’t also be an expression actually referring to Rosaline (whether it be the variable contributed by “Rosaline”, or some other expression) that Biron takes to corefer with “Katherine”. Second, when (KR’) is true and (RK’) is false, this needn’t owe to any logical incoherence in Biron’s belief state. So we are able to take the truth-value intuitions about (KR) and (RK) at face-value, without attributing logical incoherence to Biron.

⁶*ibid.* p.550.

The foregoing has admittedly been a very quick sketch of Cumming’s view. But we have enough of the framework in place to raise my concern for the proposal. I’m not the first to express such worries. Pickel (2013), for example, argues that Cumming’s framework licenses odd entailments because of how belief operators bind variables within their scope. But my concerns are a little more basic: I worry that Cumming’s framework doesn’t adequately deal with the data from the very puzzle that is used to help motivate the variabilist treatment of attitude reports.

The concern begins with a simple observation: the readings of (KR)–(MK) which are true, and the readings of (RM)–(KM) which are false, are all readings in which coreferring names (and indeed any coreferring expressions, like descriptions, pure indexicals, and demonstratives) are intersubstitutable *salva veritate*.

Here is an example to bring this out. Suppose that Katherine and Rosaline have two lesser known names: “de Longueville” and “de Chevreuse”. You are at the masque with friends, and recall that Biron is confused, but forget whether Katherine gave her favor to Rosaline, or vice versa. A disagreement arises. One friend reports (KR). The other reports (RK). Two things seem true. First, you know one of these reports is true and the other false. And second, you know that whichever report is true, it would still be true if we substituted the lesser-known names. You know that if Biron thinks Katherine is Rosaline, then (LC) will have a ‘parallel’ true reading.

(LC) Biron thinks de Longueville is de Chevreuse.

And if instead Biron thinks that Rosaline is Katherine, (CL) will have a ‘parallel’ true reading.

(CL) Biron thinks de Chevreuse is de Longueville.

The reports would seem to have the relevant true readings even if Biron is not in any way familiar with these names, or even if he has independent, positive, erroneous views about who they refer to (e.g., to ladies far off, in no way associated with Katherine or Rosaline).⁷

Interestingly, Cumming seems to recognize and acknowledge this bit of data. In distinguishing his case from Frege’s puzzles involving failures of substitutivity, he says: “Note, too, that the puzzle does not involve the substitution of coreferring terms. On the contrary, if we suppose that Rosaline, Katherine, and Maria have alternative names—de Chevreuse, de Longueville, and de Motteville, respectively—it seems we can substitute these names in the sentences above [including (KR) and (RK)] *salva veritate*.”⁸

⁷Note that (LC) and (CL) may have several other readings, especially if Biron has positive, erroneous views about the names. This won’t detract from my point, which only concerns the readings of (LC) and (CL) that intuitively ‘parallel’ those of (KR) and (RK). A related point holds of descriptions. As I’ve already claimed, descriptions that corefer with these names may be substituted *salva veritate* as well. Of course these substitutions will also have other readings that may change the truth-values, owing to a *de re de dicto* ambiguity. The readings with shifted values aren’t relevant to the point I’m making.

⁸*ibid.* p.530.

Here, then, is the problem: Cumming’s proposed logical forms for (KR) and (RK) involve treating at least one name in each report *de dicto**. But it is integral to Cumming’s view that coreferring names are *not* intersubstitutable *salva veritate* in such contexts, as we’ve already discussed. It could be suggested that Cumming was hasty in acquiescing to the substitutivity data, and should retract his commitments on that score. But the truth-value intuitions about the parallel readings of (LC) and (CL) are just as robust as for (KR) and (RK). Maintaining the latter judgments are reliable semantic intuitions while the former are not is *ad hoc*. These are just two manifestations of the same linguistic phenomenon.

Given this, I can’t see any easy way for Cumming to adjust his proposal to account for the substitutivity data while preserving the other virtuous features unique to his variabilist account. To clarify the worry, let me run through three options that might seem promising.

First, one could suggest reading all names *de re**, instead of alternating *de re** and *de dicto** readings. The resulting logical forms would be (KR'') and (RK'').

$$(KR'') (\exists\alpha)(\exists\beta)(\alpha = v_k \wedge \beta = v_r \wedge Bel_{v_b}(\alpha = \beta))$$

$$(RK'') (\exists\alpha)(\exists\beta)(\alpha = v_r \wedge \beta = v_k \wedge Bel_{v_b}(\alpha = \beta))$$

Since *de re** readings of names permit substitution of coreferents *salva veritate*, this will capture the substitutivity data. The problem is that if Biron’s belief state is logically coherent, then if (KR'') is true, (RK'') will be true as well. Given Biron’s logical coherence, any variables that can be substituted truly for α and β in the doubly *de re** (KR''), can be substituted truly for β and α in the doubly *de re** (RK'').

Could we just bite the bullet and claim Biron *is* logically incoherent to block the inference from (KR'') to (RK'')? Certainly. However, this concession not only adds a counterintuitive consequence to the view, but undermines the motivation for the variabilist treatment of the puzzle. To account for the puzzle we can no longer model Biron’s belief state with a set of world-assignment pairs where the worlds are metaphysical possibilities. We will need new tools like ‘impossible worlds’ in which some objects a and b are such that $a = b$ but $b \neq a$, or structured propositions so that $\langle v_a = v_b \rangle$ but not $\langle v_b = v_a \rangle$ could be compatible with Biron’s beliefs for a single assignment g . But once we do this, none of the variabilist machinery is doing any essential work in resolving the puzzle. It is all done by the tools modeling incoherent belief. If we’re allowed the counterintuitive recourse of logical incoherence, we can get the same result, and much more simply, by just being Millians. Indeed, the resolution being considered is the most elementary strategy the Millian can pursue—one where Biron believes something like a simple contradiction—with variabilist tools overlaid superfluously. The variabilist should not be happy with this outcome. At the very least, she can’t pursue this route and continue to claim the puzzle distinctively motivates a variabilist treatment.

If *de re** readings alone are insufficient to resolve the puzzle, perhaps we could alter the semantics for *de dicto** readings of names to ensure substitutivity is guaranteed for them as well. Two problems for this avenue stand out. First, it's not easy to see how to adjust the framework to achieve this result. As I say, it was an integral part of Cumming's framework that coreferring *de dicto** names fail to be intersubstitutable *salva veritate*. Second, even if the change could be made it will come with high costs and uncertain benefits. If *de dicto** readings yield extensional contexts, then the new variabilist view predicts that coreferring names are intersubstitutable in *all contexts*, on *all readings*. This comes with a major concession: variabilism can no longer be used to treat Frege-style puzzles. But, worse, even with that concession it is unclear how the variabilist tools—the shiftability of a name's denotation—can be doing the work of resolving the puzzle. The importance of sensitivity to shifts in a parameter is precisely that it *forbids* substitution of coreferring terms (compare the sensitivity of descriptions to various parameters of an index). Once we say that the shifts don't lead to such substitution failures, there will be serious worries as to whether the shifts themselves could be playing any essential role in resolving the puzzle.

The last worry here can be clarified and sharpened by asking what would happen if, instead of securing substitutivity for *de dicto** readings logically with semantic changes, we secure it on Cumming's current semantics with added contingent assumptions about Biron's belief state. Provided Biron is logically coherent, general substitutivity can be achieved if he also exhibits a kind of 'infallibility' concerning the denotations of names read *de dicto**. (4) expresses such infallibility as regards "Katherine".

$$(4) (\forall\alpha)(\alpha = v_k \rightarrow Bel_{v_b}(\alpha = v_k))$$

This is the most problematic route for dealing with the problem. First, it is implausible that (4) holds on the interpretation Cumming gives of *de dicto** belief, given that (4) requires infallibility concerning names Biron has never heard of. Second, even if we grant Biron the infallibility in (4), he will need such infallibility for all three names (since all occur *de dicto** in one of (KR)—(MK)). But if Biron is logically coherent, this will just again ensure that if (KR') is true (RK') is true.⁹ So we'll once more have to use logical incoherence to account for the core of the puzzle, again undermining motivations for variabilism.

These options already countenance some radical changes in application of the variabilist machinery, but none are of any help. And a pattern is emerging from which we can draw a general lesson.

None of the foregoing considerations show that variabilism is false. I take no stand on that issue. But they do seem to show that variabilist tools aren't of help in dealing with the puzzle. In fact, we can say something more general, and much more important: looking to resolve the puzzle by treating the denotations of names as sensitive to *any* shiftable parameter of the index seems misguided. As recently noted, it is precisely such sensitivity that characteristically leads to failures of substitutivity. Definite descriptions, for example, are sensitive to

⁹See the appendix for a proof sketch.

parameters of place, time, and world. As a result, substituting descriptions that corefer here, now, and actually will not generally preserve truth-values when those descriptions are influenced by location-, time-, or world-shifting operators. It is no coincidence that Cumming’s introduction of sensitivity in the denotations of names to a shiftable assignment parameter led to failures of substitutivity when the denotations of the names were actually shifted in *de dicto** readings. That is precisely what such sensitivity is introduced to account for. But since uses of names in the relevant readings of (KR) and (RK) *don’t* exhibit failures of substitutivity, those readings seem *precisely not* to be exhibiting sensitivity characteristic of shifts of an index.¹⁰

Cumming bills his puzzle as posing a special challenge to the Millian. But the substitutivity data seems to show that this is getting things backwards. Positing a sensitivity in the denotations of names to a shiftable parameter is unhelpful in accounting for the puzzle because the relevant contexts are, surprisingly, just extensional contexts—contexts where the substitution-failures characteristic of intensional, shifty constructions aren’t manifested. Put more simply, what matters to understanding what is said with (KR), (RK), (LC), or (CL) merely involves understanding which individuals the used names *actually* refer to. This is why substituting other names (and indeed any expressions) that actually corefer with them doesn’t change their truth-values on the relevant readings.

If anything, then, this puzzle is a perfect candidate for treatment with Millian tools—tools designed to account for surprising truth-value judgments in spite of a name’s semantic contribution being only its actual referent. The pressing question before us is which such tools to use.

2 Inroads through Quine

There are many interesting Millian treatments of attitude reports, and no space to give them all their due here. I’ll be setting aside, without argument, views which claim that our intuitions about Cumming’s sentences are not semantic, wheeling in pragmatic tools to account for the truth-value judgements. I’ll likewise be setting aside views that attribute logical incoherence to Biron.

I’m setting these approaches aside in part because of suspicions about their viability, but also because I think there is much to be learned by trying to explain the puzzles as what they appear to be: manifestations of peculiarities in holding beliefs about particular objects. Doing this holds the promise of integrating a

¹⁰Note, the claim is not that it’s technically impossible to have expressions be sensitive to shifts while exhibiting substitutivity. Rather it’s that this maneuver would be surprising, unprecedented, and unhelpful. It would be surprising, and unprecedented, since no other intensional shifts work like this. It would have to be the case that although belief operators shift parameters to which the denotation of a name is sensitive, every pair of expressions (names, descriptions—any referring terms really, since substitution is permitted for any) that actually corefer continue to corefer under the relevant shifts. In other words, how the shift takes place ultimately depends only on the denotation of the referring expression relative to contextually initialized parameters. However, again, not only would this be surprising and unprecedented but, as noted in the second and third cases above, it would be unclear how this shiftability could actually do the real work in resolving the puzzle.

treatment of Cumming’s puzzle into a broader account of intensionality *de re*—for example, integrating it with other more familiar puzzles about *de re* belief ascription involving definite descriptions.

My goal in the next three sections is to develop the outlines of such an approach. I’ll do this by articulating three theses about *de re* belief ascription which form a ‘skeletal view’ that can be fleshed out in alternative ways. A key virtue of the position is that it relies almost exclusively on distinctions and concepts required to deal with any intensional construction. To help articulate the first of my three theses, I need to begin a little earlier, with Quine.

Though suspicious of intensional constructions, Quine felt compelled (at least early on) to give an account of systematic ambiguities that arise in belief attributions involving determiner phrases. This kind of ambiguity arises, for example, in (5).

(5) Ralph thinks someone is a spy.

(5) can communicate that Ralph has the humdrum belief that spies exist, or the more surprising belief that some particular individual is a spy. A tempting way to represent this distinction is with something like a scope ambiguity, with (6a) attributing the more common of the two beliefs to Ralph.

- (6) a. Ralph believes $[(\exists x)(x \text{ is a spy})]$
 b. $(\exists x)(\text{Ralph believes}[x \text{ is a spy}])$

Quine refused this option, claiming that (6b) made no sense because it quantified into an intensional context. The problem is that to make sense of (6b) we need to make sense of the objectual satisfaction of “Ralph believes x is a spy”. And to do that we need to make sense of how an object, independently of any way of specifying it, satisfies this open formula. But substituting different coreferring expressions for x —two different ways of picking out the same individual—results in a change of truth-value. Doesn’t this interfere with our understanding of how an individual herself satisfies the open formula? Won’t such an understanding arbitrarily privilege some ways Ralph thinks about her?

To help stress the problem, Quine raised the now familiar case of Ralph and Ortcutt.

There is a certain man in a brown hat whom Ralph has glimpsed several times under questionable circumstances on which we need not enter here; suffice it to say that Ralph suspects he is a spy. Also there is a gray-haired man, vaguely known to Ralph as rather a pillar of the community, whom Ralph is not aware of having seen expect once at the beach. Now Ralph does not know it, but the men are one and the same. Can we say of this *man* (Bernard J. Ortcutt, to give him a name) that Ralph believes him to be a spy?¹¹

¹¹Quine (1956) p.179.

Quine’s case is meant to dramatize obstacles to the intelligibility of objectual satisfaction in belief reports. Orcutt is both (say) the man in the brown hat at the bar and the gray-haired man at the beach. Ralph believes the man in the brown hat at the bar is a spy, but Ralph doesn’t believe the gray-haired man at the beach is a spy. Quine worried because of this that we couldn’t say of ‘the man himself’, Orcutt, that Ralph has beliefs about whether or not he is a spy.

Our reluctance to pronounce on this case is interesting and important. But Quine’s identification of intensionality as the culprit here is suspect. Parallel arguments for the unintelligibility of objectual substitution into other intensional constructions, like tense operators, are clearly fallacious. Suppose I report:

(7) A senator will greet you on the Capitol steps.

(7) exhibits a parallel ambiguity to (5). It can predict that a future (and perhaps not present) senator will greet you on the Capitol steps, or it can predict a greeting by a present senator (perhaps no longer a senator by the time of greeting). Again, a tempting way to represent the distinction is roughly as one of scope, with **F** marking a tense operator shifting the time index to future times.¹²

- (8) a. **F**[($\exists x$)(x a senator \wedge x greets you on the Capitol steps)]
 b. ($\exists x$)(x is a senator \wedge **F**[x greets you on the Capitol steps])

But we can raise analogous worries to those Quine raised for belief here. Can we allow for (8b) which quantifies into an intensional context? Suppose I wanted to communicate this reading because Menendez, presently the oldest senator from New Jersey, will greet you in six years on the steps of the Capitol. Couldn’t understanding my utterance as (8b) be problematic, because it involves quantification into an intensional context?

Consider, for example, that (9) is now true, because both descriptions pick out Menendez.

(9) The oldest senator from New Jersey is the winner of the 2006 New Jersey senatorial election.

Let’s suppose there is a future time at which the embedded claim in (10) is true, because Menendez will meet you, but there is no future time at which the corresponding embedded claim in (11) is true—there is no time at which you are greeted by the oldest senator from New Jersey at that time.

- (10) **F**[the winner of the 2006 New Jersey senatorial election greets you on the Capitol steps]
 (11) **F**[the oldest senator from New Jersey greets you on the Capitol steps]

¹²This may be an oversimplification of the semantics of tense, but the simplification won’t interfere with any the points I want to make here.

So substitution of coreferring terms in the open formula

(12) $\mathbf{F}[x \text{ greets you on the capitol steps}]$

results in a change of truth-value—(10) is true and (11) is false. So how can we claim that some *individual*, Menendez, satisfies this open formula when substituting two different ways of picking him out results in a change of truth-values? Won't making this claim arbitrarily privilege some way of singling Menendez out?

Of course no one would be persuaded that objectual satisfaction of (12) was problematic on these grounds. Obviously there are good reasons for 'privileging' the description (13) over (14) in this case.

(13) the winner of the 2006 New Jersey senatorial election

(14) the oldest senator from New Jersey

Though (13) and (14) *presently* corefer (to Menendez), which accounts for the truth of (9), they do not corefer *in the future*. (13) continues to refer to Menendez in the future, but (14) clearly need not. In the intensional context created by (12), it would be odd to expect substitution of presently coreferring terms to preserve truth-value. Rather, we should expect only expressions that pick out the same object or individual in the future to preserve truth-values when substituted for one another.

Quine's worries, when transposed to the temporal case, are easily resolved. Moreover, this resolution helps make sense of the *de re* readings of determiner phrases occurring in tensed claims. This reading of (15), for example, is the one which would be true if Menendez were to meet you on the steps of the Capitol long past his final term.

(15) The oldest senator from New Jersey will greet you on the Capitol steps.

What helps us make sense of the *de re* reading of the description in (15), despite requiring quantification across a temporal operator, is that we seem to have a firm grip on relations of identity over time—over the parameter the operator is effectively shifting. We have a good sense what it is for some future person on the Capitol steps to be identical with Menendez, the man we presently identify as the oldest senator from New Jersey. This is so even though, of course, articulating the criteria which underlie those judgments of personal identity across time is philosophically quite challenging.

Arguably the same sort of resolution should be applied to cope with *de re* readings of determiner phrases occurring in sentences with modal operators.

(16) The United States president elected in 2012 could have lost the 2012 presidential election.

There is a reading of (16) which is true because Barack Obama could have lost the 2012 presidential election. There is a trivially false reading of the same sentence, which claims it could have been the case that someone both won and lost the 2012 presidential election. What helps us make sense of the *de re* reading of (16), though it involves quantification across a modal operator, again seems to be a grasp of a notion of identity across shifts in the relevant parameter—the world of evaluation. We know, somehow, what it would be like for the actual president, Barack Obama, to be identical with someone bested in the 2012 election in alternative circumstances. This is so even though articulating criteria for that identity (and, in fact, settling whether we even need to articulate such criteria) is a philosophically challenging problem.

A caveat before I continue. I've claimed that we seem to have a grip on relations of identity across time and across metaphysically possible worlds. But need the relation we are tracking be *identity*? Not obviously. In the modal case, Lewis has defended a form of counterpart theory in which claims about what is possible, or necessary, for an actual object are settled on the basis of facts about possibilia that are counterparts of, but not identical with, the actual object.¹³ Sider among others has explored the application of counterpart theory to the temporal case, where stages of objects stand in counterpart relations that secure the intelligibility of *de re* temporal discourse.¹⁴ I don't want to presuppose that the relations facilitating modal and temporal discourse *de re* are simply the relation of identity, counterpart relations, or something else. The important point is that there are some kinds of relations—I'll call them temporal and modal *intensional relations*—that do the facilitating work. We have a good grip on when these relations hold in the temporal case, and perhaps the modal case as well. And objects which bear these intensional relations to each other, whatever those relations turn out to be, may sometimes fail to satisfy the same descriptive material despite being so-related.

What we learn from the temporal case is that although Quine's Ortcutt raises some interesting worries, those worries are not on their own enough to cast suspicion on quantification across intensional operators. Intensionality on its own cannot be to blame, or we would find the same problems pervading the temporal case, which we don't. Quine's claims here could conceivably be pressed if supplemented with additional arguments targeting the intelligibility of belief itself, as Quine famously did supply for modal notions. But Quine's arguments against the intelligibility of modality are themselves suspect, and anyway do not bear directly on belief.¹⁵

As such, despite the interest of Ortcutt's case, Quine supplied no reason against taking the ambiguity in (5) to be a product of something like the simple scope distinction that Quine rejected in (6).

(5) Ralph thinks someone is a spy.

(6) a. Ralph believes[($\exists x$)(x is a spy)]

¹³Lewis (1968), Lewis (1986) ch.4.

¹⁴Sider (1996), Sider (2001) ch.5, §8.

¹⁵See Fine (2005a,b) for discussion.

b. $(\exists x)(\text{Ralph believes}[x \text{ is a spy}])$

Moreover blocking Quine’s worries by examining the cases of tense and metaphysical modals invites a natural strategy for understanding the quantification in (6b). It does so by motivating a general principle undergirding the intelligibility of *de re* readings generally. This is the first thesis of my skeletal account of *de re* ascription.

Intensional Relations. For any parameter p of the index, to understand quantification across an intensional operator shifting p , we must be able to understand and apply the relevant intensional relation connected with p .

Quantification across tense operators requires an understanding of temporal intensional relations—relations which we bear to our future and past selves. Quantification across metaphysical modal operators requires an understanding of modal intensional relations—those which relate us to the people we could have been.

If we take this simple idea to heart, then we should strive to make sense of Quine’s original case of (5), and *de re* belief ascriptions generally, in the same spirit. The strategy should appear independently attractive, because it allows us to understand *de re* ascriptions in the simplest way possible: as a semantic phenomenon arising from something like a scope ambiguity. Moreover, the strategy promises to draw illuminating ties between *de re* belief and other forms of *de re* intensionality. If we adopt this strategy, we need a clearer understanding of two things: First, how are we to understand the parameter shifted by a belief operator? Second, what sorts of facts, if any, undergird intensional relations across shifts in that parameter? Our next task is to supply answers to these questions, and to see what they might teach us about *de re* ascriptions generally, and Orcutt’s case and Cumming’s Puzzle in particular.

3 Belief Worlds

In continuing to develop my skeletal proposal for interpreting *de re* belief ascription, I’ll be drawing selectively from an account proposed by Stalnaker.¹⁶ Some of the specifics of Stalnaker’s view—especially some of the more controversial elements of the framework in which he sets up his account—won’t be essential components of my proposal. Part of the work to be done in what follows is to separate out the elements in Stalnaker’s work that are essential to dealing with puzzles like Cumming’s, and to highlight how defensible and powerful those elements are when taken in isolation.

To begin, let’s return to the first of my two questions: what parameter is shifted by the belief operator, and what is the result of the shift? The most natural proposal here surfaced in our discussion of the *de re/de dicto* distinction.

¹⁶Stalnaker (1988, 2009).

The belief operator seems to shift the world of evaluation away from the way the world actually is, and to something like the way (or ways) the world might be according to one or more believers.

For this idea to make sense, there must be something like ‘a way the world might be (according to a thinker)’, something approximating a world state, or a set of world states. Stalnaker believes there are independent grounds to think that mental states are structured precisely to yield such states in the simplest way possible: coherent belief states are exhaustively characterized by a single set of metaphysically possible worlds.¹⁷ If this is right then we can think of a belief operator as working very much like a metaphysical modal operator, shifting the world of evaluation from one metaphysically possible world to others.

Not all theorists are comfortable with this assumption. Some think that mental states can be related to epistemically possible, but not metaphysically possible, worlds. The latter worlds are ways thinkers suppose to be genuine candidates for how the world might be that are not genuine candidates at all. For example, theorists might posit merely epistemically possible worlds in which Mark Twain is not identical with Samuel Clemens, though these individuals are identical at all metaphysically possible worlds. Other theorists might depart from Stalnaker in a different way, by taking mental states to have a special kind of structure that parallels that of a theory couched in a compositionally structured language. On these views, the structure of a belief state is something structurally like a set of natural language sentences. Or, as yet a further departure, mental states can also be supposed to have internal structure of a kind which is specially *non-linguistic*—irreducibly pictorial structure, for example.

Since the view I’ll articulate builds on Stalnaker’s account, it harmonizes well with his ‘unstructured’ representations of mental states. But this conception is not necessary to understand the core elements of my proposal. All that is required is the following very weak assumption, which is my second thesis.

Belief Worlds. Logically coherent belief states are structured in a way that makes it possible to identify qualitatively coherent possibilities (metaphysical or otherwise) compatible with them, which are populated by individuals bearing identifiable properties.

By ‘qualitatively coherent’, I mean that the possibilities exhibit no form of internal incoherence up to non-qualitative facts (e.g., that Mark Twain is not Samuel Clemens). As such we can coherently describe these possibilities on a qualitative level and say, for example, how many individuals are in them, where they are located relative to each other, and so on.

Belief Worlds is innocuous enough to be compatible with the views of very many theorists, and accordingly is a thesis I’ll be assuming without further explanation from now on. I’ll call the possibilities it alludes to *doxastic alternatives*.

Once we take up this assumption we have the materials to ask the second of the questions I raised at the end of §2: what sorts of facts, if any, undergird

¹⁷See Stalnaker (1984).

intensional relations bridging shifts between the actual world and various doxastic alternatives? Or, to give these relations a more convenient name, on what basis are *doxastic relations* settled?

Note that this question is distinct from the question of what doxastic relations are. Are they simply the relation of identity? Are they counterpart relations? I won't take a firm stand on these questions. Instead I want to ask: what, if anything, settles whether an object from the actual world is related in the appropriate way to an object in one of an agent's doxastic alternatives.¹⁸ (Compare the corresponding question for trans-temporal relations: what settles whether a person (or time-slice of a person) is the one on the basis of which *de re* claims about my future are settled? We can answer this question by claiming that psychological and causal facts (say) settle who is related to my present self in the relevant ways, without immediately taking a stand on whether the relations that the psychological and causal facts undergird is identity, merely a counterpart relation, or something else.)

Stalnaker gives a simple answer to my question that I find a helpful point of departure for explaining *de re* ascriptions about non-abstract objects.¹⁹ On this view, what settles whether an object in a doxastic alternative is suitably related to an object from the actual world turns on whether facts about the former objects causally covary (perhaps under 'normal' circumstances) with facts about the latter.²⁰

Here's a simple illustration of how causal relations allow an object to shape a thinker's doxastic alternatives in the relevant way. A single blue rock—Marty—lies at Jones's feet, in plain sight and broad daylight. As a result, in all of Jones's doxastic alternatives there is a blue rock at Jones' feet. If Marty were red then (other things equal) Jones's belief state would doubtless be different: it would be compatible only with doxastic alternatives where a single red rock lies at Jones's feet. This counterfactual truth is the product of simple causal relations tying facts about the shape, size, and color of the rock to the character of an

¹⁸If one takes doxastic relations to be identity relations, then one might think that this question could have no answer: maybe the identities are brute. It's true such a *brute identity* view is coherent, and yields a coherent way of rejecting my question. But the view faces a question roughly analogous to the one I'm asking here, namely: "in virtue of what does a believer's doxastic alternative contain such-and-such an actual person or thing?" This is no longer a question about what relates objects from the actual world to those in a doxastic alternative, but about what makes a doxastic alternative, some of whose entities are brutally identical to those in the actual world, compatible with the believer's belief state. This question still needs an answer. And the answer I give to the question of what settles doxastic relations below can, for the brute identity view, simply be transposed to answer this analogous question it faces. The result is an equally informative extension of the brute identity view, which enables it to address puzzles about belief reports in the way I'm suggesting.

¹⁹Is it a count against this view that it can't cope with abstract objects as well? We don't take an account of temporal intensional relations for human beings that appeals to facts about psychology to be inadequate simply because it can't be transposed directly to buildings or rocks, let alone numbers. In general, the facts undergirding intensional relations for objects of different kinds may differ, even though the intensional relations are connected with the same parameter of the index. Absent arguments to the contrary, we should be open to the idea that this might likewise be true of doxastic relations.

²⁰See, e.g., Stalnaker (1988) p.164.

object in Jones's doxastic alternatives. And this is what makes facts about the rock in Jones's doxastic alternatives establish what Jones believes about Marty. This is obviously just a very crude sketch. But for my purposes it will do. Even so briefly stated, the view provides enough of an answer to classify a substantial range of simple cases of *de re* ascription for illustrative purposes.

A warning: this particular conception of doxastic relations isn't an essential component of my proposal. I need some specific view of doxastic relations to show how the other elements of my view engage in particular cases. Any reasonable account of doxastic relations will suffice for me to bring out two key features of my view: the similarities it forges between all forms of *de re* intensionality, and the corresponding differences that would distinguish *de re* belief in particular. The former similarities will come out in §4. Let me say something about the latter differences now.

We can see differences between forms of intensionality *de re* arise in a less dramatic way for temporal and metaphysical modal shifts. Some philosophers, as I have already noted, take strong psychological connections to help settle temporal *de re* facts for human beings. But few take the same connections to settle modal intensional relations between humans in distinct metaphysically possible worlds. If this is right we may be able, simply on the basis of psychology, to rule out some future or past individual as a candidate to settle how Jane was or will be. But we may be unable to use the very same criteria to rule out some individual in counterfactual circumstances as settling how Jane *could* have been. In this way, the distinction between the facts grounding the intensional relations seems to allow for a kind of greater flexibility in *de re* modality than *de re* temporality: Jane could have been many more things than she was or will be.

If doxastic relations are settled in part by causal sensitivities, this would introduce a still greater permissiveness in belief *de re*—a special permissiveness in what we do, or could, believe that Jane is. And of course, familiarly, this is true to the data. Eli sits in his camping tent reading, and looks up to see a dark shadow cast by what appears to be a tree on the tent surface. Unexpectedly, the shadow begins changing in size, growing and appearing to come closer until the tent opens, revealing Jane as the source of the shadow. Later, recounting the event, Eli reports with apparent truth:

(17) I initially thought Jane was a tree.

An account of doxastic relations as the product of causal sensitivities can contribute to a very simple account of (17)'s truth, without necessarily imposing constraints on the semantics of names that conflict with Millianism. At the time of Eli's confusion, he is located in a tent in his doxastic alternatives next to a looming tree. But features of that tree in Eli's doxastic alternatives, such as its physical location, are sensitive to the corresponding features of Jane. If Jane were initially standing a few feet to the left, then the tree in Eli's doxastic alternatives at that time would have been further to the left, for example. On the view I'm provisionally working with, this is sufficient to relate Jane to the

tree in Eli's doxastic alternatives in ways that make facts about that tree settle what Eli believes about Jane.

I've aimed to remain agnostic about the exact nature of doxastic relations. But if we take the truth-values of sentences like (17) at face value, then we've seen how those relations can and should exhibit a great deal of permissiveness. Will this constrain which relations we opt for? Perhaps. It depends on other aspects of one's views of belief and modality. If one takes doxastic alternatives to be exhausted by metaphysical possibilities, this will apply pressure against construing doxastic relations as identity. But even taking doxastic relations to be identity relations across metaphysically possible worlds is possible with suitable 'give' in other parts of one's view. For example, we can accommodate the data by adopting a permissive ontology which takes individuals-at-a-world as modal 'slices' which can be combined permissively to generate entities that span worlds (though this requires a bit of indeterminacy in which particular modally extended individual "Jane" may refer to from use to use). An alternative is to maintain doxastic relations as identities, but proliferate worlds beyond metaphysical possibilities. And still another option is to construe doxastic relations as counterpart relations, a position which may, with relative ease, be combined with almost any view on belief and modality. I suspect, for this reason, my skeletal view is best framed as a form of counterpart theory, but I won't explore the details here.

Instead, since the core framework of an account of *de re* belief ascription is already in place. Just one more thesis is needed to treat Cumming's puzzle. To appreciate that thesis, though, we must reconsider Ortcutt.

4 Problematic Intensional Relations

On our view so far, as long as we can settle which objects in the doxastic alternatives of A are doxastically related to an object o from the actual world, ascriptions to A of *de re* beliefs about o should be readily intelligible. So what is going wrong in the case of Ortcutt? Why is it unobvious to some how to answer Quine's rhetorical question about objectual satisfaction in Ortcutt's case? We can address these worries by noting that Ortcutt's case is actually one instance of a much more general phenomenon that arises for all intensional relations: sometimes settling the relations can be challenging, or unprincipled, and when this is the case it will adversely affect the interpretability of *de re* ascriptions. As always, it helps to revisit other forms of intensionality *de re* to see the point.

Consider, for example, tensed claims. Let's grant for the moment a bit of science fiction required for a classic thought experiment of Wiggins: that all the information in our brain is stored twice—once in each hemisphere.²¹ Menendez, the present oldest senator from New Jersey, will soon be in a car accident. Doctors will harvest his brain, implant each of its two hemispheres in a different body, which will live on with the memories and personality of Menendez. One of these two persons dies in a few years. At that time the other

²¹See Parfit (1971) for an discussion of the case.

(who will then be no senator) meets you on the steps of the Capitol. Is the *de re* reading of (15) true *now*?

- (15) The oldest senator from New Jersey will greet you on the Capitol steps.

Obviously it depends. It depends on whether the person who meets you on the steps of the Capitol with half of Menendez's brain, and all his memories and personality, settles what will be true of Menendez. Does he? This is a controversial matter. It is so controversial that it is not implausible to maintain there is *no answer*—in bizarre enough cases our grip on persistence conditions for humans, and other objects more generally, can be strained to the point where it breaks down. If this is right then of course there is no fact of the matter about the truth or falsity of (15) either. Our grip on temporal intensional relations may falter, and the intelligibility of relevant quantifications across temporal operators will then fall with it.

We can raise the same troubles for metaphysical modal operators. Suppose that if Barack Obama's parents had conceived a child one minute before they actually did, the same sperm and ovum which formed Barack Obama would have combined, but a second sperm would have fertilized a separate ovum at the same time. In the early stages of development the two zygotes would fuse, resulting in a rare form of 'chimerism'—a fetus with two genetically distinct populations of cells. In the counterfactual circumstances we are considering, the chimera grows into a normal adult who runs for president in 2012 and loses to Mitt Romney. Is (18) true?

- (18) If the president of the United States's parents had conceived one minute before they actually did, then he would have lost the 2012 presidential election.

Again it depends. It depends on whether the person who would grow out of the fused zygotes can settle what is possible for Obama. It seems plausible that if the resulting individual had all the same genetic material as Obama, he could. It seems plausible that if the resulting individual had *none* of the same genetic material as Obama he could not. Would one cell from the other zygote make a difference? Two? Is there some fixed cut off? If not, then we might have scenarios in which it is not merely difficult, but unprincipled, to say whether some possibility is a possibility *for* Obama or not. Again, when our grip on the modal intensional relations is sufficiently problematic, (18) becomes equally problematic in turn.

These are manifestations of a restatement, with the help of contraposition, of the principle of *Intensional Relations*, from §2.

Intensional Relations': For any parameter p of the index, if we can't understand and apply the relevant intensional relation connected with p , we can't understand quantification across an intensional operator shifting p .

When we transpose this principle to the case of belief attribution, we find that when doxastic relations become problematic, so do the *de re* belief ascriptions whose intelligibility relies on them. I've followed Stalnaker in taking these relations to be secured by something like causal covariation. But I want to stress that this isn't essential to the point I'm making here. The point is that *whatever* criteria help settle the relevant doxastic relations, those criteria *could* be problematized in ways that would interfere with the interpretation of *de re* belief ascriptions.

It is this phenomenon, whose existence can be motivated purely by consideration of the class of intensional constructions generally, that Stalnaker suggests we should use to understand Ortcutt. We know more or less what Ralph's doxastic alternatives are like. In these worlds, there is one man in a brown hat at the bar at the time of Ralph's visit who is a spy. There is a man who Ralph sees on the beach who is not a spy. In Ralph's doxastic alternatives, these two men are not identical. Which of them settles what Ralph believes about Ortcutt? We can't answer because, it seems, they are equally good candidates. Exactly why this is so might be controversial. Taking up Stalnaker's recommendation, we can see that both of these individuals in Ralph's doxastic alternatives carry information about Ortcutt, as facts about them covary with facts about Ortcutt. For example if Ortcutt were actually taller, both men in Ralph's doxastic alternatives would correspondingly be taller. This seems like a plausible explanation, but we don't need it to be the right one. All we need is some criteria underlying the application of doxastic relations which relate Ortcutt and the man on the beach in Ralph's doxastic alternatives, and Ortcutt and the man at the bar in those same alternatives.

As long as this holds, both men will be equally good candidates to settle Ralph's beliefs about Ortcutt. Perhaps this means that neither actually should settle what Ralph believes about Ortcutt. Perhaps it means that either, or both, should. Either way the case is problematic enough to explain our mixed intuitions about the case. As I argued before, the culprit here is *not intensionality* which, as we've seen, doesn't generally create insuperable difficulties in the temporal or modal cases. Problematized intensional relations are to blame. Though Quine misdiagnosed his case, the mistake is forgivable. In cases of belief ascription, unlike with metaphysical modals and tense operators, problematic intensional relations and failure of substitution of coreferring terms tend to come together, obfuscating the real source of our confusions.

Interestingly, obstacles to settling doxastic relations may easily disappear, as Stalnaker notes, if we simply eradicate the basis for relating him to one of the men in Ralph's doxastic alternatives.²² Suppose just as in Quine's original case that Ralph has heard of Ortcutt as a pillar of the community and has seen him at the beach, but that he never sees Ortcutt in questionable circumstances. Does Ralph think of Ortcutt (the man himself) that he is a spy? It seems safe to say in this scenario that Ralph thinks Ortcutt is not a spy. Why can we say

²²Stalnaker (1988) discusses these cases under the heading of "divide and conquer arguments", citing Kripke (1979) as an early appreciator of their importance.

this? Because now it is once again easy to find the person in Ralph's doxastic alternatives who settles what Ralph believes of Ortcutt, so the relevant *de re* ascriptions again regain their utility. This phenomenon again has an interesting potential parallel in the case of temporals (and perhaps metaphysical modals as well): Menendez gets into his earlier accident, which immediately causes him to lose one hemisphere of his brain before any chances of transplanting it. The other hemisphere, along with Menendez's personality and memories, remains intact and is used in a successful transplant operation. When given this case, some report more confidently that Menendez will survive the accident. Some problems for temporality *de re* might only arise when there is, as it is sometimes put, a suitable 'competitor' to settle the facts about Menendez. What Ortcutt's case reveals is that a perfectly parallel phenomenon arises for belief.

5 Contextual Relocation

We're now in a position to articulate the third and final thesis about *de re* belief ascription required to account for Cumming's puzzle: a sensitivity in those ascriptions to context. This thesis has the potential to be the most controversial element of my skeletal account. Some frameworks for separate intensional notions, like Lewis' counterpart theoretic treatment of modality *de re*, build a form of context sensitivity into them. Others do not. But even if no other form of intensionality *de re* exhibits context sensitivity, I submit that we should accommodate such sensitivity in ascriptions of belief *de re*. Given our first two theses, the empirical data alone gives us ample motivations.

Stalnaker arrives at this conclusion after noting that in Quine's original case, even if Ralph's mental state doesn't seem to change, we may nonetheless be able to felicitously make *de re* belief ascriptions about Ortcutt if the context of attribution is right.²³ This seems right, but I think we can see the motivations more clearly in a different case, where the operation of context is more transparent. Suppose there is a famous horror writer from Maine named "Carl Filler" who, as it happens, has briefly taken up residence in a neighborhood of St. Louis and frequents a local bar there. The bartender knows Carl as a regular named "Carl Filler" but does not know he is the author of the same name. If asked, the Bartender would report that he suspects the horror author lives in a log cabin in Maine. One day Carl comes into the bar at which point the bartender proceeds to pour him a complementary amount of a fine whiskey in the hopes of garnering better tips.

Over in a conversation among die-hard Filler fans, who know the bartender is ignorant of Filler's identity as author, we can overhear the following.

A: "Pah. Look at the bartender! I can't believe it. He still doesn't recognize Filler. He thinks Filler is off in some log cabin in New Hampshire. Carl's presence is wasted here."

²³Stalnaker (1988) p.166.

B: “You’re right that the bartender doesn’t realize Filler is here, but you’re wrong about where he thinks he is. He thinks Filler is off in Maine—that’s the local misconception.”

We can simultaneously hear a separate conversation, at which patrons, who are again privy to the bartender’s ignorance, discuss the virtues of different kinds of whiskey.

C: “Look. Carl Filler, that dime-store hack, is back again and currying the bartender’s favor. Amazingly the bartender thinks that Filler is the type of person to enjoy good quality whiskey. I’d rather suspect that Filler likes his drinks like he likes his fiction—cheap and trashy.”

D: “You’re mistaken there. The bartender thinks that though Filler will drink what he pours enthusiastically, he won’t really enjoy the whiskey all—he just thinks Filler will tip better if he gives him some of the expensive stuff.”

What we have here are two conversations where speakers attribute to the bartender apparently incompatible beliefs about the man, Carl Filler: that he is presently in Maine, or that he is presently at the bar and about to drink the whiskey being poured out, for example. What’s interesting about these conversations is that some of these attributions which collectively attribute incompatible beliefs seem, *prima facie*, to be true—those made by B and D for example. Moreover, some of the other attributions which collectively attribute incompatible beliefs about Filler seem *untrue*—the utterances of A and C, in their respective conversations for example.

If we take the corresponding truth-value judgments at face value, what governs the correctness of these ascriptions? As always, our first step is to get a clear grip on the bartender’s doxastic alternatives, qualitatively described. It is relatively transparent what these are like. In them, there is a famous horror author, writing in a log cabin in Maine, named “Carl Filler”. There is also a distinct individual likewise named “Carl Filler” who is no author, and a regular at the bar where the bartender works in St. Louis. To understand which beliefs the bartender has about Carl Filler, the man himself, we must locate the individual in these doxastic alternatives who settles such matters. Who is he: the man at the bar, or the man in Maine? We have here a case very much like that of Orcutt, in which an actual individual’s influence is ‘split’ between two individuals in various doxastic alternatives. So why is conversation proceeding so smoothly?

What seems to occur is that each conversation has its own standards for who settles the bartender’s beliefs about Filler. In the conversation in which A and B participate, speakers are presently taking the relevant individual to be the author in the bartender’s doxastic alternatives. This is the person who lives in Maine, and not in New Hampshire, in those worlds. Further ascriptions of belief about Filler in the same vein will be judged for truth or falsity according to how things stand with that individual in the bartender’s doxastic alternatives. The

conversation in which C and D are engaged has different standards. Clearly, they are taking the man at the bar in the bartender’s doxastic alternatives to settle what the bartender believes of Filler. This is the person in those doxastic alternatives who will merely pretend to appreciate whiskey about to be poured, and will leave a big tip. Again, further ascriptions in this same vein will be judged for truth or falsity according to how things stand with *that* individual in the bartender’s doxastic alternatives.

These ways of resolving the indeterminacy in who settles the bartender’s beliefs about Filler are incompatible with each other. Consider what would happen if the Bartender thought that the horror author who lived in Maine by the name of “Carl Filler” did enjoy fine whiskey. Even if this were supposed, it seems quite safe to label C’s utterance of “. . . the bartender thinks that Filler is the type of person to enjoy good quality whiskey” as false (note that typically the speaker would, upon learning more, be willing to retract this assertion). A corresponding utterance of this sentence in the course of the conversation of A and B, however, could be interpreted as a truth with relative ease.

Stalnaker suggests that we should explain such shifting standards as the product of a process of Lewisian accommodation.²⁴ So conceived, the problem of locating appropriate individuals via doxastic relations must be construed, at least to some modest extent, as a matter of convention among speakers who are using that individual to describe the character of a thinker’s belief state. When, say, facts about causal sensitivities don’t unambiguously fix who is doxastically related to whom, we sometimes use conversational cues to conventionally settle, at least for conversational purposes, which individuals do the work of settling the truth of *de re* ascriptions. This seems true to the facts. Aspects of context such as speaker intentions, salience, and relevance do seem to help settle what kinds of *de re* ascriptions are correct, or incorrect, in cases like those I’ve just given.

These empirical considerations, paired with our two earlier theses, motivate a final principle governing *de re* belief ascription.

Contextual Resolution. When there are multiple candidates at a belief world to be doxastically related with an object from the actual world (as per *Intensional Relations'*), features of the context (e.g. salience, relevance) may resolve the ambiguity in favor of one of these candidates.

The details of how the context of utterance affects interpretation here are complex, but no more so than any other kind of contextual accommodation. Cooperative speakers say things which are topical, interesting, and true, and sometimes it will be obvious to all participants in a conversation just which among multiple candidates should be doxastically related to an object in the actual world to make pertinent *de re* ascriptions satisfy these conditions—the Filler case is a paradigmatic example.

²⁴Lewis (1979).

Contextual Resolution requires context to select among or resolve doxastic relations. How context mediates this process may differ depending on how we chose to accommodate the flexibility which characteristically sets *de re* belief apart from other forms of *de re* intensionality. If, for example, we take doxastic relations to be identity and choose to accommodate that flexibility by allowing a flexible conception of individuals ‘extended’ across worlds, context will effectively settle which such individuals are salient to stand as the referent of a single name. If doxastic relations are counterpart relations, context will help select among, or disambiguate, the counterpart relations themselves. And there are alternatives besides. Importantly, on no view I’ve considered so far do doxastic relations *need* to be part of logical form (though of course, the context which does the selecting, or disambiguating, may be). Doxastic relations form part of an account of what it takes to have beliefs about objects—what constitutes having such beliefs.

Though there are many ways of implementing *Contextual Resolution*, it is the more controversial of the three theses I’ve articulated. So it’s important to stress just how powerful and explanatory an assumption it is. Let’s witness this power now, by returning to Cumming’s puzzle.

6 Applications

We’ve now seen that admitting the three principles of *Intensional Relations*, *Belief Worlds*, and *Contextual Resolution* can allow for systematic ambiguities in the *de re* readings of (19), and accordingly in (20) as well even if we adopt a Millian semantics for names.

(19) Ralph thinks the *F* is *G*.

(20) Ralph thinks *n* is *G*.

Importantly, on the skeletal view I’ve been developing, (19) and (20) should be able to exhibit not just a two-way ambiguity, but as many readings as one likes provided the contexts of attribution afford them. This does seem to be the case. Consider the following example of a ‘twice-over’ confusion of identities.

Ralph is at the Miskatonic University graduate conference with the famous political philosopher George Smith as keynote. Ralph knows of George’s work, but is unaware that he is in attendance. When George introduces himself by name to Ralph before the first speaker, Ralph mistakenly takes him to be a technician and asks him to help set up the lecture slides. As George does this, a graduate student, Tom, comes up to assist. At first George is working with some cords on the left, and Tom is working with some cords on the right. Ralph turns his back for a moment to grab some cables for George, during which time George and Tom change places. Ralph

turns around, clearly unaware of the switch. He heads to hand the cables to Tom, on the left, while glaring at what he takes to be inappropriate dress of George, to the right.

The following three dialogs are taking place around the room.

Dialog 1:

- A: “Why does Ralph have that awful expression on his face?”
B: “He clearly thinks George Smith is not dressed appropriately for the conference.”

Dialog 2:

- C: “George Smith was the one who asked for cables. Why is Ralph heading to give cables to the guy on the left of the podium instead of the guy on the right?”
D: “He got confused when his back was turned. He now thinks George Smith is the person standing on his left.”

Dialog 3:

- E: “Why has Ralph been treating the keynote so disrespectfully?”
F: “He doesn’t realize who he is talking to. He thinks George Smith is in Princeton, and that he’s been talking with a technician all along.”

Our skeletal account gives the resources to see more clearly what’s happening here. Ralph’s doxastic alternatives are easy to characterize. In them, there are three individuals we need to keep track of. One is named “George Smith” and is a well-known philosopher working away in Princeton. Another individual with the same name is a technician come to fix the lectures slides. In Ralph’s doxastic alternatives the second individual is now working with cords to the left of the podium, while a third individual—a graduate student—works away to its right.

Which of these individuals settles what Ralph believes of George Smith? Now we have three candidates. Changes in facts about George Smith in the actual world can systematically lead to changes in *three* individuals in Ralph’s doxastic alternatives. If George Smith were to have specialized in mereology, then the philosopher in Princeton in Ralph’s current doxastic alternatives would have as well. If George Smith had worn different clothes, so would the graduate student in Ralph’s current doxastic alternatives. And if George Smith were to have sneered at Ralph when asking for cords, so would the technician presently working to the left of the podium in Ralph’s doxastic alternatives.

Dialogs 1–3 proceed as smoothly as they do because context allows speakers to easily resolve the ambiguities in whom to doxastically relate with George

Smith in Ralph’s doxastic alternatives. In the first dialog, it’s with the graduate student; in the second, it’s with the technician; and in the third, it’s with the philosopher. Moreover the divergences in these relations have a very natural explanation: they are the only ways to relate George Smith to individuals in the doxastic alternatives so as to make the replies suitably pertinent to the questions being asked in the respective conversations.

I bring this case up because the same kinds of ambiguities are present in Cumming’s original case of mistaken identities. Maria, masquerading as Kathryn, continues to twirl her carnation now visible to Biron who misidentifies it as a rose. Rosaline, masquerading as Maria, holds no flower, again in Biron’s sight. It should be easy to see how to get *prima facie* true readings of each of the following sentences, set in a suitable context.

(21) Biron thinks Maria is holding a rose, not a carnation.

(22) Biron thinks Maria is not holding any flowers.

We can also get readings on which these can be false. Note, for example, that assertions of the conjunction of (21) and (22) tend to have a contradictory feel because, it seems, no single context easily accommodates them both.

Cumming has the tools to try to account for the true readings of (21) and (22) without giving Biron contradictory beliefs by appealing to the alternating scope of “Maria”. But this strategy will not deal with the phenomenon in its full generality, as the George Smith case reveals: in general there can be more possible ambiguities than a simple scope distinction will allow.

(21) and (22) have both true and false readings because, in a now familiar manner, Maria is responsible for the features of two distinct women in Biron’s doxastic alternatives. Maria’s influence is accordingly split in those alternatives between two candidates to settle what Biron believes of her. The source of the ambiguities in (21) and (22) are clearly closely tied to the asymmetries in Cumming’s original statements (KR)–(MK). As such, our skeletal account gives us all the resources needed to account for the puzzle. To take one example, let’s revisit (KR).

(KR) Biron thinks Katherine is Rosaline.

There are two women in Biron’s doxastic alternatives who are responsive to facts about Katherine. There is a woman in Biron’s doxastic alternatives whose features are *presently* responsive to the actual Katherine (a woman named “Rosaline”, presently wearing Rosaline’s favor), and a distinct woman in Biron’s doxastic alternatives whose features owe to *past* causal relations with the actual Katherine (a woman named “Katherine”, presently wearing Katherine’s favor). If the actual Katherine were to refuse to dance at the ball, so would the woman in Biron’s doxastic alternatives wearing Rosaline’s favor. If the actual Katherine were to have rolled her eyes when receiving her favor, so would the distinct woman named “Katherine” in Biron’s doxastic alternatives. Similarly, two women in Biron’s doxastic alternatives are responsive to facts about

Rosaline: the woman named “Maria” wearing Maria’s favor (owing to present causal relations), and the woman named “Rosaline” wearing Rosaline’s favor (owing to past causal relations).

Accordingly, on my view, we need context to disambiguate both the doxastic relations associated with Katherine and Rosaline to get any intelligible reading of (KR). How should this be done? Suppose we know (as it seems we can) that Biron has no straightforwardly illogical beliefs, and he knows all three ladies are distinct from each other. Then if we interpret an utterance of (KR) charitably in this context, we must disambiguate the doxastic relations associated with Katherine and Rosaline so as to make it possible for the individuals to whom they are related in Biron’s doxastic alternatives to be identical. To do this we need to ‘alternate’ methods of interpretation—for example by resolving doxastic relations associated with a first mentioned name in accordance with present causal relations, and resolving doxastic relations associated with a second mentioned name in accordance with past causal relations. Note, this would require context to intervene intra-sententially, applying one method for resolving intensional relations (through present causal relations) to the first used name, and another (through past causal relations) to the second. But there is no obstacle to this happening, on the current view, since shifts in resolutions of intensional relations owe precisely to contextual cues, including linguistic context.

If we do this for (KR), we will take Katherine to be related to the woman in Biron’s doxastic alternatives whose features she presently shapes, and Rosaline to be related to the woman whose features she shaped in the past. These two women are, in fact, one and the same: the woman in Biron’s doxastic alternatives named “Rosaline” and wearing Rosaline’s favor. This explains how (KR) can be true. Note that a ‘parallel’ disambiguation applied to (RK) will yield a falsehood: it will incorrectly say that the woman in Biron’s doxastic alternatives whose features Rosaline presently shapes (the woman wearing Maria’s favor) is identical with the woman in Biron’s doxastic alternatives whose features Katherine once shaped (the woman wearing Katherine’s favor). The ‘alternating’ strategies for contextually locating Katherine and Rosaline in Biron’s doxastic alternatives is part of what accounts for the informativeness of (KR), even provided our assumption of Biron’s logical and conceptual prowess. Similar remarks apply to the rest of our asymmetric identities.²⁵

²⁵Note that of course there is another fruitful way of disambiguating doxastic relations: by resolving the doxastic relations to the first lady mentioned according to past causal relations, and resolving the doxastic relations to the second lady mentioned according to present causal relations. This will give us the ‘inverted’ readings mentioned in n.2. For example, on this reading (KR) will come out false and (RK) true, for exactly the same reasons (RK) was false and (KR) was true on the previous contextual disambiguation. Since context is responsible for these disambiguations, we should expect the set of readings we get to be responsive to changes in context and subject matter as desired. Why is one of these two sets of readings easier or more natural than the other? We have many options to explain this. I suspect there is probably a default practice, in ambiguities owing to multiple causal sources, to taking present causal relations as primary. This is why the first name in (KR) and (RK) tends to be read this way, so that the second name needs ‘accommodation’ to a distinct standard. The grounds for such a practice are obvious: in general these default relations will generally be most relevant to present psychological theorizing about believers. This is just one explanation,

In effect, we’ve recovered the structure Cumming appealed to in dealing with his puzzle, but via a different route—by appealing to indeterminacies in doxastic relations, and contextually driven shifts in the resolutions of those indeterminacies (in on our working example, according to temporally privileged sets of causal relations). But, crucially, we don’t run into the problem that Cumming did: we can account perfectly for the substitutivity data. Note that a ‘reading’ of a name on the proposed view, as it occurs in a belief attribution to someone with identity confusions, is a strategy for resolving indeterminate doxastic relations (e.g., by favoring present, or past, causal relations). Parallel readings of (KR) and (LC), then, are readings in which the parallel contextual resolution strategies are applied to the first names, and second names, respectively.

(LC) Biron thinks de Longueville is de Chevreuse.

As long as we choose the same strategies, (KR) and (LC) will have the same truth-values. Indeed, truth-values will be preserved on the substitution of *any* coreferring expressions, including definite descriptions (read *de re*), demonstratives, and indexicals. This is because intensional relations are, by definition, relations that hold between two *objects*. If a resolution strategy is held fixed, we have the same relation. And if we substitute coreferring terms, we have the same object. This means that the second object in the relation—the individual in Biron’s belief worlds who is used to settle the truth-values of the attitude report—*must* be the same. Accordingly, no change in truth-values arises.

This account of the case maintains several very desirable features. It treats Cumming’s puzzle as a semantic phenomenon by taking stable truth-value intuitions at face value. In particular, it gives a semantic treatment not only of the ‘inversion’ data, but also the substitutivity data. The account captures that data while keeping Biron’s belief state logically and conceptually coherent. Throughout, it makes use only of resources compatible with Millianism.²⁶ And

though. My view, like Cumming’s, is compatible with a number of possible accounts.

²⁶I should reiterate that though I do not think Cumming’s puzzle distinctively motivates a variabilist treatment, my view is perfectly *compatible* with variabilism as well. Indeed, a variabilist position that integrates the machinery of counterpart theory in ways that could possibly also integrate the claims of this paper can be found in Santorio (2012) (see especially pp.389ff.). Santorio, importantly, motivates a significant variabilist interaction with attitude reports in a much different way than Cumming. Still, Santorio’s extension of his framework to Cumming’s puzzles faces, I think, some special challenges precisely because of the enmeshing of variabilist and counterpart-theoretic machinery. For example, for Santorio, counterpart relations are tagged with numerical indices that link them to variables, not individuals (*ibid.* p.401)—precisely allowing the counterpart relations used in an attitude ascription to map coreferring expressions to different objects. This is reasonable given the puzzles he takes as his starting point. But the question of how to recapture the substitutivity data (with ‘parallel’ readings) for Cumming’s cases then arises again. Second, the account requires non-trivial extensions to deal with definite descriptions (and perhaps other referring expressions). Santorio’s framework requires treating all names in these puzzles getting bound ‘*de dicto*’ readings. Again, this makes sense given his starting puzzles. But then the general substitutivity data will seemingly commit him to using *de dicto* readings of descriptions in variants of Cumming’s puzzle with descriptions. This will sometimes require special readings of descriptions on which they have denotations relative to belief worlds where their descriptive material is not satisfied. I’m not sure whether there are such readings. But the account just

perhaps most importantly, it manages to view Cumming’s puzzle as an instance of a much more general phenomenon dealt with by a single overarching account of *de re* belief ascription, and in turn subsumes this account of *de re* belief under a completely unified account of all intensionality *de re*.

Though the account draws heavily on Stalnaker’s work, it requires a sparse set of resources which can be developed in ways incompatible with many of Stalnaker’s more controversial views if desired. The account consists of three theses. Two of these—*Intensional Relations* and *Belief Worlds*—are motivated by consideration of the resources needed to cope with, and understand, intensional operators generally, and belief ascriptions as one among many intensional constructions. Once these two theses are accepted, the final thesis—*Contextual Resolution*—can be motivated almost exclusively by empirical data. Taken together, the theses make no commitments on controversial questions about the fineness of grain of mental or linguistic content. A Russellian, or even a Fregean, about propositional content can adopt these theses, and the account if desired.²⁷ The account makes no assumptions about the structure of a mental state beyond the claim that logically coherent mental states determine qualitatively coherent doxastic alternatives of *some* kind populated by objects with identifiable properties. Finally, it makes no assumptions about the nature of intentionality or accounts of relations bridging the actual world and doxastic alternatives beyond those required, in conjunction with the three theses, to take intuitive truth-value judgments at face value.

Appendix

Suppose we treat (KR) as true given Cumming’s proposed logical form in some model M , world w and relative to some variable assignment g —that is, suppose that (a) holds.

(KR) Biron believes Katherine is Rosaline.

$$(a) \models_{M,w,g} (\exists \alpha)(\alpha = v_k \wedge Bel_{v_b}(\alpha = v_r))$$

One way of ensuring that substitution of coreferring names preserves truth-values for the expressions read *de dicto** in (KR)–(MK) is to assume that Biron exhibits the following kind of infallibility (along with logical coherence):

$$(b) \models_{M,w,g} (\forall \alpha)(\alpha = v_k \rightarrow Bel_{v_b}(\alpha = v_k))$$

$$(c) \models_{M,w,g} (\forall \alpha)(\alpha = v_r \rightarrow Bel_{v_b}(\alpha = v_r))$$

given has the virtue of not needing them: all substituted descriptions on my account will be read *de re* and so have their descriptive material satisfied in the actual world, where there is guaranteed to be an object satisfying it. As I say, these are only *prima facie* worries, and I would be happy to have the account given here be integrated with a variabilist proposal were it given independent motivations, as Santorio has done, and shown empirically adequate for treating the puzzles.

²⁷Though of course the account may threaten to make some aspects of a Fregean view appear superfluous.

$$(d) \models_{M,w,g} (\forall\alpha)(\alpha = v_m \rightarrow Bel_{v_b}(\alpha = v_m))$$

If Biron is logically coherent, we get substitutivity. (c) and (a), for example, would then ensure

$$(e) \models_{M,w,g} (\forall\beta)(\beta = v_r \rightarrow (\exists\alpha)(\alpha = v_k \wedge Bel_{v_b}(\alpha = \beta)))$$

Why? By (a) there is some variable, call it v_α , such that v_α corefers with v_k (at the ‘actual’ world w and assignment g) but such that Biron believes v_α corefers with v_r . Now, if we pick any variable—say v_β —that actually corefers with v_r , then by (c), Biron knows it corefers with v_r . So Biron thinks v_α corefers with v_r and that v_β corefers with v_r . He’s logically coherent, so he must think v_α corefers with v_β . So we have a variable, v_α that actually corefers with v_k , but that Biron thinks corefers with v_β . But v_β was an arbitrary variable coreferring with v_r . So (e) holds.

The problem is that we can then show that (a) (that is, (KR)), (b), and Biron’s logical coherence entail (RK). If (KR)/(a) is true, for some variable v_α we have

$$(f) \models_{M,w,g} v_\alpha = v_k \wedge Bel_{v_b}(v_\alpha = v_r)$$

Given the first conjunct of (f) and (b) we have

$$(g) \models_{M,w,g} Bel_{v_b}(v_\alpha = v_k)$$

By (g), the second conjunct of (f), and Biron’s logical coherence we have

$$(h) \models_{M,w,g} Bel_{v_b}(v_r = v_k)$$

So, trivially

$$(i) \models_{M,w,g} v_r = v_r \wedge Bel_{v_b}(v_r = v_k)$$

Hence

$$(j) \models_{M,w,g} (\exists\alpha)(\alpha = v_r \wedge Bel_{v_b}(\alpha = v_k))$$

That is, (RK) is true on its proposed logical form.

References

- S. Cumming (2008). ‘Variabilism’. *The Philosophical Review* **117**(4):525–554.
- K. Fine (2005a). ‘Quine on Quantifying In’. In *Tense and Modality: Philosophical Papers*, pp. 105–132. Clarendon Press.
- K. Fine (2005b). ‘The Problem of *De Re* Modality’. In *Tense and Modality: Philosophical Papers*, pp. 40–104. Clarendon Press.
- J. D. Fodor (1970). *The Linguistic Description of Opaque Contexts*. Ph.D. thesis, Massachusetts Institute of Technology.

- D. Kaplan (1968). ‘Quantifying in’. *Synthese* **19**(1):178–214.
- S. Kripke (1979). ‘A Puzzle about Belief’. In *Meaning and Use*, vol. 3 of *Studies in Linguistics and Philosophy*, pp. 239–283. Springer.
- S. Kripke (1980). *Naming and Necessity*. Harvard University Press.
- D. Lewis (1968). ‘Counterpart Theory and Quantified Modal Logic’. *Journal of Philosophy* **65**(5):113–126.
- D. Lewis (1979). ‘Scorekeeping in a Language Game’. *Journal of Philosophical Logic* **8**(1):339–359.
- D. Lewis (1986). *On the Plurality of Worlds*. Blackwell.
- D. Parfit (1971). ‘Personal Identity’. *The Philosophical Review* **80**(1):3–27.
- B. Pickel (2013). ‘Variables and Attitudes’. *Noûs* **48**(1).
- W. Quine (1956). ‘Quantifiers and propositional attitudes’. *The Journal of Philosophy* **53**(5):177–187.
- P. Santorio (2012). ‘Reference and Monstrosity’. *Philosophical Review* **121**(3):359–406.
- T. Sider (1996). ‘All the World’s a Stage’. *Australasian Journal of Philosophy* **74**:433–53.
- T. Sider (2001). *Four-Dimensionalism*. Oxford: Clarendon.
- R. Stalnaker (1984). *Inquiry*. MIT Press.
- R. Stalnaker (1988). ‘Belief Attribution and Context’. In *Context and Content*, pp. 150–167. Oxford.
- R. Stalnaker (2009). ‘What is De Re Belief?’. In J. Almog & P. Leonardi (eds.), *Essays in Honor of David Kaplan*. Oxford University Press.