

A Relationist Theory of Intentional Identity

Dilip Ninan
dilip.ninan@tufts.edu

Abstract

This essay argues for a *relationist* treatment of intentional identity sentences like (1):

- (1) Hob believes that a witch blighted Bob's mare and Nob believes that she killed Cob's sow. (Geach, 1967)

According to relationism, facts of the form *a believes that ϕ and b believes that ψ* are not in general reducible to facts of the form *c believes that χ* . We first argue that extant, non-relationist treatments of intentional identity are unsatisfactory, and then go on to motivate and explore a relationist alternative in some detail. We show that the general thesis of relationism can be directly motivated via cases already discussed in the literature, and then develop a particular version of relationism couched in the possible worlds framework. The resulting theory avoids the problems facing its non-relationist rivals, and yields a natural account of the truth-conditions of (1), truth-conditions which can be generated in a compositional manner by a version of dynamic semantics. The theory also helps us to cleanly separate semantic questions about intentional identity from metasemantic ones.

1. Introduction

Suppose that I believe that ϕ and you believe that ψ , for some ϕ and some ψ . You and I thus stand in a certain two-place relation, which we may visualize as follows:

_____ believes that ϕ and _____ believes that ψ .

Let's call such a relation a *dyadic belief relation*. Suppose now that I also believe that χ , for some χ . Then I have a certain monadic property, which we may visualize as follows:

_____ believes that χ .

Let's call such a property a *monadic belief property*. It is natural to think that the facts about which individuals stand in which dyadic belief relations are determined by the facts about which individuals have which monadic belief

properties. Indeed, one might think that the logic of conjunction alone guarantees this. Nevertheless, I think that there is a case to be made that this natural view is wrong, and that some dyadic belief facts are in fact not reducible to the monadic belief facts. Let us call this broad view *relationism about belief*.¹

My argument for relationism involves the phenomenon of *intentional identity*, which Geach (1967) first introduced with examples like this:

- (1) Hob believes that a witch blighted Bob's mare and Nob believes that she killed Cob's sow.

While sentence (1) may have more than one reading, Geach was interested in a reading of it on which it does not entail the existence of witches, and on which the pronoun *she* occurring in the second conjunct is in some sense anaphoric on the indefinite description *a witch* occurring in the first conjunct.

What is the connection between relationism and intentional identity? The idea is this: as the extant literature reveals, it turns out to be surprisingly difficult to find a pair of monadic beliefs, b and b' , such that (1) is true iff Hob has b and Nob has b' . A tempting conclusion to draw is that the reason we can't find a pair of monadic beliefs meeting this description is that there is no such pair; in other words, the dyadic belief fact asserted by (1) is not reducible to any conjunction of monadic belief facts. But whether we ought to embrace this relationist conclusion depends on two things: (i) how difficult it really is to find the needed pair of monadic beliefs, and (ii) whether switching from non-relationism to relationism helps matters at all. While I shall present arguments that cast doubt on non-relationist approaches, my principal aim in this essay is to demonstrate that switching from non-relationism to relationism really does advance our understanding of intentional identity.

It will help to distinguish three questions we can ask about intentional identity sentences like (1):

TRUTH-CONDITIONS: Under what conditions are sentences like (1) true?

COMPOSITIONAL SEMANTICS: Assuming we know what the truth-conditions of (1) are, how are those truth-conditions compositionally determined by the meanings of the parts of the sentence?

METASEMANTICS: In virtue of what does a sentence like (1) have the truth-conditions that it does?

To get a sense of the three-way distinction being drawn here, consider a monadic belief ascription like (2):

- (2) Kripke believes that Feynman was a physicist.

¹I take the term *relationism* from Fine (2009), who uses *semantic relationism* for the view that the facts about which semantic relations a pair of linguistic expressions e and e' stand in are not in general determined by the facts about which intrinsic semantic properties e and e' each have separately. My form of relationism, which concerns beliefs rather than linguistic expressions, is not an instance of Fine's, but there is family resemblance between them.

According to a familiar view, (2) is true just in case Kripke stands in a certain binary relation (the monadic belief relation) to the singular proposition that Feynman was a physicist—this is a proposal concerning the *truth-conditions* of this sentence. Assuming that these *are* the truth-conditions of (2), we may then seek a *compositional semantics* that shows us how those truth-conditions can be determined as a function of the meanings of the parts of the sentence. Finally, a familiar *metasemantic* claim says that Kripke can only stand in the belief relation to the singular proposition that Feynman was a physicist if he (Kripke) is casually related to Feynman in the right way.

A full theory of intentional identity would answer all three of the above questions concerning sentences like (1). The present discussion focuses principally on the first two questions, questions concerning truth-conditions and compositional semantics, though we will have something to say about where the metasemantic issue fits in.

I begin in §2 by examining extant, non-relationist accounts of intentional identity. My aim here is not to refute non-relationism definitively, but only to impress upon the reader the difficulty of constructing a satisfactory non-relationist theory of these matters. §3 is then devoted to motivating and developing a relationist account in some detail. We begin in §3.1 by offering a more precise statement of relationism and then provide some initial motivation for that view. In §3.2, we present a particular version of relationism, one couched in possible worlds semantics. We show that this approach yields a natural account of the truth-conditions of intentional identity sentences, one that avoids the problems facing non-relationist accounts. §3.3 discusses the metasemantic question, arguing that the present proposal helps us to cleanly separate metasemantic issues from semantic ones. Finally, §3.4 sketches a version of dynamic semantics that assigns to intentional identity sentences our proposed relationist truth-conditions; the details of this semantic theory are presented in an appendix.

2. Non-relationism

2.1. Descriptivism

To see the sort of reading of (1) that Geach has in mind, imagine Hob and Nob having the following sort of exchange:

HOB: There's a witch going around town these days. I think she blighted Bob's mare last night.

NOB: I heard about that witch. I bet she also killed Cob's sow.

Focussing on such cases suggests that Geach's sentence is equivalent to something like the following:

- (3) Hob believes that a witch blighted Bob's mare and Nob believes that the witch that blighted Bob's mare killed Cob's sow.

The intended reading is one on which both belief ascriptions are read *de dicto*. On this approach, the pronoun *she* in (1) somehow goes proxy for the underlined definite description in (3). Alternatively, perhaps Nob is unsure as to whether the witch in question really did blight Bob's mare, but nevertheless realizes that *Hob believes* that she did. In that case, we might instead interpret Geach's sentence as follows:

- (4) Hob believes that a witch blighted Bob's mare and Nob believes that the witch that Hob believes blighted Bob's mare killed Cob's sow.

Here the pronoun in (1) is understood to go proxy for the more elaborate description *the witch that Hob believes blighted Bob's mare*. Approaches along this line fit naturally with what are known in the semantics literature as E-type approaches to anaphora (Evans, 1977; Cooper, 1979; Heim, 1990; Elbourne, 2005).

Notice that, on either of these views, (1) becomes straightforwardly equivalent to a conjunction of *de dicto* belief reports. On either view, there is a pair of beliefs *b* and *b'* such that (1) is true iff Hob has *b* and Nob has *b'*. For instance, on the first of these views, *b* is the belief that a witch blighted Bob's mare, while *b'* is the belief that the witch that blighted Bob's mare killed Cob's sow. So it seems that, on this view, there is a clear sense in which the dyadic belief fact reported by (1) is reducible to a pair of monadic belief facts.

Now while (3) and (4) might both be possible readings of (1), Geach observed in his original article that these are not the only available readings of (1). For it seems that there are situations in which (1) is true while both (3) and (4) are false. Here is an example, adapted from Edelberg (1986):

Newspaper Case

A number of animals in Gotham Village have recently died quite unexpectedly. Rumors have begun to circulate that these unfortunate events are due to the machinations of a witch. The local newspaper, the Gotham Star, has picked up on these rumors and dubbed the witch *Samantha*. The paper has reported that Samantha has been attacking animals and destroying crops. In reality, there is no such individual: the animals in question all died of natural causes, the crops withered from drought. But Hob and Nob both read the Gotham Star and both believe the articles about the witch. Hob thinks that the witch must have blighted Bob's mare, which fell ill recently, while Nob thinks that the witch killed Cob's sow. But Nob is unaware of Hob's and Bob's existence, and so has no beliefs about Hob or Bob at all.

Since Nob has no beliefs about either Hob or Bob, Nob does not believe that the witch who blighted Bob's mare killed Cob's sow, nor does he believe that the witch that Hob believes blighted Bob's mare killed Cob's sow. Thus, neither (3) nor (4) is true in this scenario. But it is widely thought that (1) has a reading

on which *it* is true in this scenario.² If that is correct, then neither (3) nor (4) captures the intended reading of (1).

A number of authors take cases like the Newspaper Case to show that the truth of (1) requires that Hob's belief and Nob's belief have a common causal source (Glick, 2012; Cumming, 2014; Lanier, 2014).³ For note that, in the Newspaper Case, Hob's belief and Nob's belief are both partially caused by the articles in the Gotham Star, or by the rumors circulating in town. This observation can be used to motivate an alternative 'descriptivist' story. On the alternative approach, (1) is instead equivalent to something like:

- (5) Hob believes that a witch blighted Bob's mare and Nob believes that the witch described by the actual common causal source of Hob's belief and Nob's belief killed Cob's sow. (Lanier, 2014, 298)

The inclusion of the adjective *actual* is intended to make (5) (and thus (1)) equivalent to:

- (6) Hob believes that a witch blighted Bob's mare and the common causal source *S* of Hob's belief and Nob's belief is such that Nob believes that the witch described by *S* killed Cob's sow.

Note that if (6) is true, then *the common causal source of Hob's belief and Nob's belief* is a non-empty definite description. Since (1) is equivalent to (6) on this proposal, the account predicts that the truth of (1) entails that Hob's belief and Nob's belief have a common causal source. Furthermore, this proposal avoids the problem facing the previous descriptivist approaches, since its proposed truth-condition doesn't require that Nob is aware of Hob or Bob, only that he is aware of the causal source of his belief.

The principal difficulty with this approach is that while (1) might require that Hob's belief and Nob's belief *have* a common causal source, it would not appear to require that *Nob believe* anything about this source. For example, while it is natural to assume that, in the Newspaper Case, Nob has a *de dicto* belief to the effect that the witch described in the Gotham Star article killed Cob's sow, a nearby variant of the case lacks this feature (Azzouni, 2013, 341). Imagine, for example, that Nob reads the article in the Gotham Star, forms the belief that the witch killed Cob's sow, and then proceeds to forget how he formed this belief. Maybe he later comes to believe that he learned about the witch from his friend Janice, or maybe he simply forms no new beliefs about the source of his witch-beliefs. After all, we often forget how we formed certain beliefs, but retain those beliefs nevertheless. The article in the Gotham Star is the causal source *S* of Hob's belief and of Nob's belief, but since Nob has forgotten all about that article, he has no beliefs about *S*. So (6) is false in this version of the Newspaper Case. Nevertheless, it seems that (1) is still true, which suggests that (1) is not equivalent to (6) after all.⁴

²Though see King (1993) and Braun (2012) for some doubts about this.

³Here and in what follows, by *Hob's belief* I mean Hob's belief that a witch blighted Bob's mare, and by *Nob's belief* I mean Nob's belief that a witch killed Cob's sow.

⁴The considerations in this paragraph also cast doubt on the counterpart-theoretic account

2.2. Mythical objects

The foregoing considerations provide modest evidence in favor of relationism. For we've been struggling to find a pair of monadic beliefs b and b' such that (1) is true iff Hob has b and Nob has b' , and one explanation of this fact is that there is no such pair, just as relationism would predict. But one might, of course, draw an alternative conclusion from our inability to find the needed pair of monadic beliefs. For what we've seen so far is that, if we restrict our search to beliefs concerning ordinary objects and their properties, it is difficult to find the needed pair of monadic beliefs. But perhaps this just shows that the class of monadic beliefs is larger than we initially thought: in addition to including beliefs whose content may be characterized by ordinary objects and their properties, it includes beliefs concerning certain kinds of *extraordinary* objects and their properties.

One family of approaches to intentional identity draws precisely this conclusion. Salmon (2002), for example, holds that (1) is true iff there is a 'mythical witch' x such that Hob believes that x blighted Bob's mare and such that Nob believes that x killed Cob's sow. On Salmon's view, a *myth* is any false theory that has been held true; and a *mythical object* is a hypothetical entity erroneously postulated by a myth (Salmon, 1998, 304). Mythical objects are abstract objects; they are neither physical objects nor mental entities. Salmon's view is that whenever someone a believes that there is an F that is such-and-such when there is no F that is such-and-such, then there is a mythical F thereby believed by a to be such-and-such.

Since Hob believes there is a witch who blighted Bob's mare even though there is no such witch, it follows that there is a mythical witch that Hob believes blighted Bob's mare. Furthermore, Salmon holds that if two believers believe there is an F that is such-and-such when there is no F that is such-and-such, 'they may or may not believe in the same mythical F , depending on their interconnections' (Salmon, 2002, 105, n. 25). Thus, we may assume that Hob's and Nob's interconnections in the Newspaper Case are such that they believe in the same mythical witch. In that case, if we assume that (1), on the relevant reading, has the same truth-conditions as (7), then we predict that (1) is true in the Newspaper Case.

- (7) There is a mythical witch such that Hob believes that she blighted Bob's mare, and such that Nob believes that she killed Cob's sow.

And since Hob's and Nob's interconnections in the Newspaper Case are preserved even in the variant in which we stipulate that Nob forgets how he acquired his belief, this approach likewise predicts a true reading of (1) in that variant of the case. So the 'mythical objects' approach avoids the problems facing the descriptivist views discussed above.

As Salmon (2002, 107, n. 28) observes, this theory is not itself committed to the existence of mythical objects. What the theory implies is that *if* (1) is true,

of intentional identity found in Glick (2012). See, in particular, Glick (2012, 391) where Glick's characterization of the relevant counterpart relation relies on there being a newspaper article S such that Hob and Nob both believe that they have read S .

then there is at least one mythical witch. So if one accepts this analysis, but rejects the existence of mythical objects, then one will have to reject the truth of (1). But if this theory is combined with the rejection of mythical objects, the result is a radically revisionary view concerning a raft of ordinary attitude ascriptions. If Jones thinks that a U.S. senator embezzled funds, then if there is no such senator, there is a mythical senator that Jones believed embezzled funds. If Smith thinks that a mouse ruined her stamped collection, then if there is no such mouse, there is a mythical mouse that Smith believes ruined her stamp collection. If there are no mythical objects, then Jones does not think that a U.S. senator embezzled funds, and Smith does not think that a mouse ruined her stamp collection, appearances notwithstanding. Thus, the most plausible version of this view accepts both the truth of these ordinary ascriptions *and* the existence of mythical objects. But it seems that we should only agree to accept the existence of these mythical objects if doing so is absolutely necessary, i.e. only if there is no other plausible account of the truth of sentences like (1). This follows not from a general metaphysical objection to abstract entities, but simply from the methodological principle that we should not multiply entities beyond necessity. Even Salmon seems to concede the point (Salmon, 2002, 107, n. 28); he just doesn't think an adequate alternative analysis is available. But what I shall argue in what follows is that an adequate alternative analysis *is* available. The relationist analysis developed below allows us to accept the truth of (1) without having to accept the existence of mythical objects.⁵

3. Relationism

3.1. *The general thesis and some initial motivation*

The foregoing discussion of non-relationist approaches to intentional identity has not been exhaustive, but it does suggest that finding a satisfactory non-relationist account is no simple matter.⁶ That provides at least some motivation for considering relationist alternatives, and it is to this task that we now turn.

We begin by giving more precise characterizations of relationism and non-relationism, respectively. Let's say that an agent *a* *has precisely the same monadic beliefs in world w as they have in world w'* iff: for all ϕ , *a* believes that ϕ in w iff *a* believes that ϕ in w' .⁷ And let's say that agents *a* and *b* *stand*

⁵Braun (2012) argues that even if there are mythical witches, it is doubtful that (1) entails this, a point I am sympathetic to. See Lanier (2014, 293-294) and Sandgren (2018) for additional worries about the mythical objects approach.

⁶One notable omission in our discussion of extant accounts of intentional identity is the approach due to Edelberg (1986, 1992, 1995) and further developed by Cumming (2014). Because this style of approach invokes a non-standard semantic apparatus (e.g. indefinite descriptions are referential, rather than quantificational, and they denote 'thought-objects' rather than ordinary objects), it would take us too far afield to examine it in any detail. Furthermore, the extant literature has already turned up some problems for this approach: Cumming (2014) points out flaws in the proposals of Edelberg (1986) and Edelberg (1992), while Lanier (2013, Ch. 3) raises problems for Cumming's own proposal.

⁷Here and in what follows, we employ a metalanguage that permits quantification into sentence position.

in precisely the same dyadic belief relations in w as they do in w' iff: for any ϕ and ψ , (a believes that ϕ and b believes that ψ in w) iff (a believes that ϕ and b believes that ψ in w'). Then *non-relationism about dyadic belief* is the following view:

NON-RELATIONISM: For any worlds w and w' and agents a and b , if a has precisely the same monadic beliefs in w as they have in w' , and b has precisely the same monadic beliefs in w as they have in w' , then a and b stand in precisely the same dyadic belief relations in w as they do in w' .

And *relationism about dyadic belief* is simply the rejection of non-relationism:

RELATIONISM: The negation of non-relationism.

The bulk of the rest of the essay develops in detail a particular version of relationism, one couched in possible worlds semantics. But it is worth separating that particular theory from the general thesis just stated, and worth observing that the general thesis can be motivated independently of the arguments for our favored version of relationism.

To see this, note that non-relationism is a particular way of saying that the dyadic belief facts supervene on the monadic belief facts. Thus, we can attempt to construct a counterexample to non-relationism by providing a pair of cases which differ with respect to the dyadic belief facts but do not differ with respect to the relevant monadic belief facts. Fortunately for us, we do not need to construct such a pair *ex nihilo*, since the extant literature already suggests a pair of cases which has precisely this feature. Consider, for example, the following pair of cases, lightly adapted from Lanier (2014, 292):

The Connected Case

Al and Bud both suspect that a witch has come to town and is poisoning livestock and destroying crops. Al and Bud get together, discuss their respective theories, and decide to warn the town, each from a separate location. Each man goes to his designated location and begins to warn passersby: ‘There’s a witch in town! She’s poisoning our livestock and destroying our crops! Be on your guard!’ Hob hears Al, and concludes that the witch in question must have blighted Bob’s mare, which fell ill recently. Nob hears Bud, and concludes that the witch in question must have killed Cob’s sow, which died unexpectedly last night. Of course, no witch (or any other person) caused any of the mishaps in question, all of which were due to natural causes. We may also assume that, as in the Newspaper Case, Nob knows nothing of Hob or Bob.

The Unconnected Case

This is exactly like the Connected Case, except that Al and Bud have never met and have no coordinated plan to warn the townspeople

about the witch. Al is delusional. Bud is bored and decides to start a witch-hunt by spreading a rumor to the effect that there is a witch in town causing trouble. Al goes to the same location that he goes to in the Connected Case, and similarly for Bud, and each man makes the same speech that he made in the Connected Case. And, again, Hob overhears Al and comes to believe that the witch Al is talking about must have blighted Bob's mare, while Nob overhears Bud and comes to believe that the witch Bud is talking about killed Cob's sow.

The Unconnected Case is essentially exactly like the Connected Case in all relevant respects, except that Hob's belief and Nob's belief do not derive from a common causal source. Lanier observes that (1) appears to be true in the Connected Case, for that case is essentially like the Newspaper Case. But Lanier also claims that (1) is false in the Unconnected Case, and Glick (2012, 392) reports a similar judgment.

Lanier's purpose in discussing pairs of cases like this is to argue for the common cause requirement. We shall return to that issue below (§3.3), but here we observe that such pairs of cases can also be used to argue for relationism. Since (1) is true in the Connected Case, but not in the Unconnected Case, there is a dyadic belief fact that holds in the Connected Case, but not in the Unconnected Case. But given how the cases are described, it is natural to suppose that Hob has precisely the same monadic beliefs in the Connected Case as in the Unconnected Case, and similarly for Nob. There is, at any rate, nothing in the description of these cases that stands in the way of our simply stipulating that Hob has precisely the same monadic beliefs in both cases, and similarly for Nob. For the only relevant difference between the two cases concerns whether or not Al and Bud are colluding, a difference that need not make any difference to either Hob's monadic beliefs or to Nob's monadic beliefs. If we accept that these two cases may be filled out in this way, then Hob and Nob have precisely the same monadic beliefs in them, despite the fact they do not stand in precisely the same dyadic belief relations in them. In that case, relationism will be true, non-relationism false.

3.2. Relationism vs. non-relationism

The foregoing argument provides some initial motivation for relationism, but we can extend the case by developing relationism in more detail, and then comparing the resulting theory to the non-relationist theories discussed in §2. The particular version of relationism I want to propose is couched in possible worlds semantics, and so it will be useful to consider briefly how the problems we discussed in §2 arise within that framework.

On standard possible worlds theories of attitudes (Hintikka, 1962), for any agent a and any ϕ , a believes that ϕ iff in every world w compatible with what a believes, ϕ . Framed in that way, the initial problem posed by (1) becomes the problem of saying what it is for a world w to be compatible with what Hob believes and what it is for a world w' to be compatible with what Nob believes,

given that (1) is true. The first of these questions has a natural answer: w should contain a witch x who blighted Bob's mare. The trouble comes with saying what w' , an arbitrary world compatible with what Nob believes, should be like. World w' should contain someone y who killed Cob's sow, but who in w' is y ? We can think of the various proposals considered in §2 as offering different answers to this question. For example, the first version of descriptivism we considered says that y is the witch who blighted Bob's mare in w' . The advocate of the mythical objects view says instead that y is identical to x , and x/y is a mythical object. But, as we argued above, none of these answers is wholly satisfactory.

Now the fact that this question has proved so difficult to answer suggests that there might be something wrong with the question itself. While it is hard to be sure of this, I think we have said enough at this point to motivate considering an approach that focusses on a different question altogether. I propose that instead of asking what it is for a world w to be compatible with what Hob believes and what it is for a world w' to be compatible with what Nob believes, we instead ask what it is for a *pair* of worlds (w, w') to be compatible with what the pair (Hob, Nob) believe. And *this* question turns out to have a comparatively natural answer: (w, w') should be compatible with what (Hob, Nob) believe only if *there is an x such that x is a witch in w , x blighted Bob's mare in w , and x killed Cob's sow in w'* . (Here, the first element of (w, w') is indexed to Hob, the second to Nob.) The view taken here is (to a first approximation) that (1) is true iff every pair (w, w') compatible with what (Hob, Nob) believe meets the italicized condition above. This basic idea will be further developed in the remainder of the essay.

Let us start with a simple question about what we've said far: what is it for a pair of worlds (w, w') to be compatible with what a pair of individuals (a, b) believe? We can approach this question by examining the parallel question that arises for the standard possible worlds semantics for attitude reports. As we said above, the standard view holds that a believes that ϕ iff in every world w compatible with what a believes, ϕ . How do we understand the notion figuring on the right-hand side of this biconditional, the notion of a world's being compatible with what an agent believes? The basic idea is that w is compatible with what an agent a believes iff: for all ϕ , if a believes that ϕ , then ϕ in w . So if, for example, Sam believes that it is raining in Tokyo, then w will be compatible with what Sam believes only if it is raining in Tokyo in w .⁸

The relationist can say something similar about the notion of a pair of worlds being compatible with what a pair of agents believe. The basic idea is that (w, w') will be compatible with what a pair of agents (a, b) believe iff: for all ϕ and ψ , if a believes that ϕ and b believes that ψ , then ϕ in w and ψ in w' . So if, for example, Sam believes that it is snowing in Chicago and Tomoko believes that it is raining in Seattle, then (w, w') will be compatible with what (Sam, Tomoko) believe only if it is snowing in Chicago in w and it is raining in Seattle in w' . More interestingly, if Sam believes that a senator from New England

⁸It should be clear that nothing we've said amounts to a non-circular analysis of the notion of belief, nor is such intended. Such an analysis, if there is one, is left to the philosopher of mind or to the cognitive scientist.

embezzled funds and Tomoko believes that she lied to the FBI, then (w, w') will be compatible with what (Sam, Tomoko) believe only if there is an x such that x is a senator from New England in w , x embezzled funds in w , and x lied to the FBI in w' .

Now these remarks, I believe, suffice to show that the relationist's key notion—that of a pair of worlds being compatible with what a pair of agents believe—is intelligible, or at least as intelligible as the parallel notion typically taken for granted in standard possible worlds theories of attitudes. So we shall take this notion for granted in what follows, and see where doing so leads.

Our proposal, recall, is that sentence (1) is true iff for all pairs of worlds (w, w') compatible with what (Hob, Nob) believe, there is an x such that x is a witch in w , x blighted Bob's mare in w , and x killed Cob's sow in w' . The first thing to observe about this approach is that it avoids the various problems facing the non-relationist theories discussed earlier. Note, for example, that, on this account, (1) does not entail (3):

- (3) Hob believes that a witch blighted Bob's mare and Nob believes that the witch that blighted Bob's mare killed Cob's sow.

On the present approach, (3) would be true iff for all pairs of worlds (w, w') compatible with what (Hob, Nob) believe, there is an x such that x is a witch in w , x blighted Bob's mare in w , x is the unique witch that blighted Bob's mare in w' , and x killed Cob's sow in w' . But given natural assumptions about what the space of worlds is like, this is clearly a stronger condition than the truth-condition we proposed for (1). For any pair of worlds (w, w') and individual x that witnesses the truth-condition for (1), x must blight Bob's mare in w and kill Cob's sow in w' . But that appears to be consistent with x not blighting Bob's mare in w' , since w' may well be a distinct world from w .

Similarly, the present approach avoids the problem facing the 'causal descriptivist' discussed earlier. For it is clear our proposed truth-condition for (1) does not require that Nob believe anything about the source of his belief. Note also that the present approach avoids the problems facing the mythical objects view, for it does not imply that (1) entails that there are mythical objects. So in these respects, our relationist approach appears to be a genuine improvement over its non-relationist rivals.

I have been calling the present account *relationist*, but I have to justify my doing so. We need to see why this account counts as relationist in the sense of §3.1. We argue as follows. Let a and b be fixed but arbitrary agents. We assume that we can extract the monadic belief facts from the dyadic ones in the following way. Let $Dox_{a,b}^w$ be the set of pairs of worlds compatible with what (a, b) believes in w , let Dox_a^w be the set of worlds compatible with what a believes in w , and let Dox_b^w be the set of worlds compatible with what b believes in w . Then we assume that:

$$Dox_a^w = \{v : (v, v') \in Dox_{a,b}^w \text{ for some world } v'\}, \text{ and}$$

$$Dox_b^w = \{v' : (v, v') \in Dox_{a,b}^w \text{ for some world } v\}.$$

And we assume that for any $c \in \{a, b\}$, c has precisely the same monadic beliefs in w as they have in w' iff $Dox_c^w = Dox_c^{w'}$. We also assume that a and b stand in precisely the same dyadic belief relations in w as they do in w' iff $Dox_{a,b}^w = Dox_{a,b}^{w'}$. Given natural assumptions about the space of worlds, we can show that there are worlds w and w' such that $Dox_a^w = Dox_a^{w'}$, $Dox_b^w = Dox_b^{w'}$, but $Dox_{a,b}^w \neq Dox_{a,b}^{w'}$. In that case, a and b will each have precisely the same monadic beliefs in w as they have in w' , but they will not stand in precisely the same dyadic belief relations in w as they do in w' . Relationism will be true, non-relationism false.

To see how this would work, suppose there are worlds w and w' such that:

$$Dox_{a,b}^w = \{(v, v') : \exists x (x \text{ blighted Bob's mare in } v \text{ and } x \text{ killed Cob's sow in } v')\}.$$

$$Dox_{a,b}^{w'} = \{(v, v') : \exists x (x \text{ blighted Bob's mare in } v) \text{ and } \exists y (y \text{ killed Cob's sow in } v')\}.$$

It seems that $Dox_{a,b}^w \neq Dox_{a,b}^{w'}$. For suppose that in v , x alone blighted Bob's mare, and that in v' , y alone killed Cob's sow, where $y \neq x$. Then (v, v') will be in $Dox_{a,b}^{w'}$, but not in $Dox_{a,b}^w$. Thus, a and b will not stand in precisely the same dyadic belief relations in w as they do in w' . But we also have that $Dox_a^w = Dox_a^{w'}$ and $Dox_b^w = Dox_b^{w'}$, which means that a and b each have precisely the same monadic beliefs in w as they have in w' .⁹ Thus, we have a difference in the dyadic belief facts despite no difference in the relevant monadic belief facts.

3.3. Metasemantics

One thing that keeps popping up in our discussion is the idea that that the truth of an intentional identity sentence imposes a common cause requirement. This was one of the main motivations for the causal descriptivist proposal discussed in §2. And Glick (2012, 392) takes pairs of cases like the Connected Case and the Unconnected Case to show that (1) is true iff three conditions obtain: (i) Hob believes that a witch blighted Bob's mare, (ii) Nob believes that someone killed Cob's sow, and (iii) Hob's belief and Nob's belief have a common causal source. Suppose, for the moment, that Glick's claim is true. Then while the present proposal entails that (1) is true only if conditions (i) and (ii) hold, it does not obviously imply (iii). So, again assuming (1) does impose a common cause requirement, this raises the question of where condition (iii) fits into our

⁹To see that $Dox_a^w \subseteq Dox_a^{w'}$, suppose $v \in Dox_a^w$. Then there is a v' and an x such that x blighted Bob's mare in v and x killed Cob's sow in v' . But then there is there is a z that blighted Bob's mare in v and there is a y that killed Cob's sow in v' , for x is such a z and such a y . So $(v, v') \in Dox_{a,b}^{w'}$, which means $v \in Dox_a^{w'}$. To see that $Dox_a^{w'} \subseteq Dox_a^w$, suppose $v \in Dox_a^{w'}$. So there is a v' such that there is an x that blighted Bob's mare in v and such that there is a y that killed Cob's sow in v' . Let u be a world in which x killed Cob's sow (we assume, plausibly, that there is such a world). So x blighted Bob's mare in v and x killed Cob's sow in u . So $(v, u) \in Dox_{a,b}^w$, which means $v \in Dox_a^w$, as desired. The argument that $Dox_b^w = Dox_b^{w'}$ is similar.

analysis. I will first answer this question on the assumption that intentional identity sentences really do impose a common cause requirement, and then return to consider whether this assumption in fact holds.

My view is that the common cause requirement is most naturally understood as a *metasemantic* requirement, something that must obtain in order for Hob and Nob to stand in the dyadic belief relation that (1) says that they stand in. It is not, as the causal descriptivist maintains, something that figures in the *content* of the psychological states reported. To appreciate the point, recall sentence (2):

(2) Kripke believes that Feynman was a physicist.

As we discussed at the outset of the paper, it may be that (2) is true only if Kripke is causally related to Feynman in an appropriate manner. But this fact arguably does not figure in the content of the belief reported, which simply concerns Feynman and one of his properties. Instead, the causal requirement appears to be a metasemantic requirement, a requirement on what must be true of Kripke in order for him to have the property of believing that Feynman was a physicist. Similarly, if (1) does impose a common cause requirement, I suggest this fact does not figure in the content of the state that (1) attributes to Hob and Nob, but is instead a requirement on what must be true of Hob and Nob in order for them to stand in the dyadic belief relation that (1) reports them as standing in. From the present perspective, the causal descriptivist mis-locates the common cause requirement, putting it into the semantics when it is properly understood as a feature of the metasemantics.

All this is assuming that (1) really does impose a common cause requirement. Is that true? And even if that is true of sentence (1) in particular, is it generally true that intentional identity sentences impose common cause requirements? Sarah Moss (p.c.) has suggested to me that this last question should be answered in the negative.¹⁰ For example, imagine that we have two causally disconnected cultures, that, perhaps by chance, have very similar theological beliefs concerning matters like the creation of the universe and the origins of humanity. If we fill in the details in the right way, this might suffice for the truth of (8):

(8) Culture A believes that a supreme being formed humans out of clay, while Culture B believes that he formed them out of fire and water.

And this despite the fact that Culture A's belief does not have the same causal source as Culture B's. If this is possible, then it will not generally be true that intentional identity sentences require for their truth that the corresponding monadic beliefs derive from a common causal source. Perhaps what is driving our judgment that (8) is true in this scenario is the fact that the deity hypothesized by Culture A plays a similar explanatory role to the deity posited by Culture B. Edelberg (1992, §8) takes related cases to indicate that intentional

¹⁰Sandgren (2019, 3682-3683) likewise rejects the common cause requirement.

identity sentences require for their truth that the relevant beliefs have a common explanatory role. Sandgren (2019) offers a different non-causal account of the metasemantic requirement, one which gives an important role to an agent's disposition to judge whether the relevant pair of monadic beliefs is about the same thing or not.

As far as I can see, the relationist *qua* relationist needn't take any particular stand on this issue. From the relationist point of view, the questions raised by this example concern the metasemantic requirements that must be met in order for a pair of subjects to stand in a particular dyadic belief relation. What the relationist has offered is an abstract account of particular dyadic belief relations themselves, not an account of the conditions that must obtain in order for such a relation to be instantiated. The relationist tells us that (1) is true iff Hob and Nob stand in the relation that *a* bears to *b* iff every (w, w') compatible with what (a, b) believe is such that there is a witch who blighted Bob's mare in *w* and who killed Cob's sow in *w'*. It does not tell us what has to be true of Hob and Nob in order for every (w, w') compatible with what (Hob, Nob) believe to satisfy this condition.

Again, it is instructive to compare the present situation with the case of the standard possible worlds semantics for attitude ascriptions. On the standard account, (2), for example, is true iff every world *w* compatible with what Kripke believes is such that Feynman is a physicist in *w*. But the standard account does not, by itself, tell us what has to be true of Kripke in order for every world *w* compatible with what he believes to satisfy this condition.

All that being said, the present discussion might help us to see what information is being encoded by dyadic belief ascriptions, and this might cast light on why we would use dyadic belief ascriptions instead of simply relying on conjunctions of monadic belief ascriptions. For if we generalize a bit from Glick's proposal, we might suppose that (1) is true iff (i) Hob believes that a witch blighted Bob's mare, (ii) Nob believes that someone killed Cob's sow, and (iii) [INSERT FAVORED METASEMANTIC REQUIREMENT HERE]. In that case, intentional identity sentences can be seen as encoding two types of information: they tell us about the relevant subjects' respective monadic beliefs, and they also tell us that the relevant metasemantic requirement—whatever, precisely, it is—has been met. An account of this general form can be given regardless of what precise form the metasemantic requirement takes.

Note that if the preceding (schematic) view is correct, then it suggests that dyadic belief facts *can* be reduced to the corresponding monadic belief facts together with certain other facts, namely whatever other facts constitute the metasemantic requirement. Is this a source of embarrassment for the relationist, who maintains that the dyadic belief facts are not reducible to the monadic belief facts? It is not. For it is no part of the relationist's view that dyadic beliefs are irreducible *tout court*—the relationist, *qua* relationist, need not hold that dyadic beliefs are part of the fundamental furniture of the universe. Even if relationism is true, it may also be true that dyadic belief facts are reducible in some way to something, and if that is so, they are likely to be reducible to the monadic belief facts together with certain additional facts. It might be helpful to compare the

situation here to the case of externalism about psychological states, in the style of Putnam (1973) and Burge (1979). The externalist denies that the monadic belief facts are reducible to certain local functional and physical facts. But that thesis is compatible with the claim that the monadic belief facts are reducible to those local facts *together with* facts about the physical and social environment (Stalnaker, 1984, Ch. 1). The externalist, *qua* externalist, needn't maintain that monadic beliefs are part of the fundamental furniture of the universe. And the same goes, *mutatis mutandis*, for the relationist.¹¹

3.4. Compositional semantics

Our final task is to construct a compositional semantics that predicts our proposed truth-condition for (1). But before we get to that, we first need to generalize our account. For note that there is nothing special about *dyadic* beliefs in particular:

- (9) Hob believes that a witch blighted Bob's mare, Nob believes that she killed Cob's sow, and Joe believes that she stole Janice's tractor.

That suggests that we should speak of an n -ary sequence of worlds (w_1, \dots, w_n) 's being compatible with what an n -ary sequence of individuals (a_1, \dots, a_n) believes. But since nothing in the present phenomenon mandates the order built into these sequences, I propose to use *functions* from agents to worlds instead of sequences. Where A is the set of agents and W the set of worlds, let an *indexed possibility* be a function w mapping agents in A to worlds in W . Note that the range of any given indexed possibility w can be organized into an indexed sequence of worlds $(w_{\text{Hob}}, w_{\text{Nob}}, \dots)$, and it can be useful to visualize an indexed possibility w by picturing the corresponding indexed sequence.¹² Assuming $\{\text{Hob}, \text{Nob}\} \subseteq A$, I propose that (1) is true at a world w iff for all indexed possibilities $w \in W^A$ compatible with what the agents in A believe in w , there is an x such that x is a witch who blighted Bob's mare in w_{Hob} and x killed Cob's sow in w_{Nob} .¹³

Our task now to is to provide a theory that assigns this truth-condition to (1) in a compositional manner. Let's begin by agreeing to regiment (1) as follows:

$$(R1) \mathcal{B}_b \exists x Fx \wedge \mathcal{B}_c Gx$$

Here \mathcal{B}_b translates *Hob believes that*, \mathcal{B}_c translates *Nob believes that*, Fx translates *x is a witch who blighted Bob's mare*, and Gx translates *x killed Cob's sow*. It is not straightforward to compositionally generate our proposed truth-conditions for this sentence assuming this syntax. For note that our proposed truth-conditions essentially have the following form:

$$(T1) \forall w \in \text{Dox}_A : \exists x (Fx \text{ in } w_b \wedge Gx \text{ in } w_c)$$

One problem with getting from our regimentation of (1) to our proposed truth-conditions is that the syntactically free occurrence of x in the second conjunct

¹¹Thanks to Andy Egan and Arc Kocurek for discussion on this issue.

¹²We write w_a for the result of applying function w to individual a .

¹³Here, W^A is the set of all indexed possibilities, i.e. the set of all functions from A to W .

of (R1) appears to correspond to a bound variable in (T1). How does a syntactically free variable get interpreted as a bound variable?

This general challenge is familiar from other cases of ‘cross-clausal anaphora’:

(10) A man is walking in the park, and he is whistling a cantata.

(11) If a farmer owns a donkey, he feeds it.

To see the difficulty posed by sentence (10), for example, imagine that we regiment it as:

(R10) $\exists x Mx \wedge Wx$

Since (10) seems to say that there is a man walking in the park and whistling a cantata, it would be natural to think of its truth-conditions as instead corresponding to something like:

(T10) $\exists x (Mx \wedge Wx)$

Thus, the question arises as to how we provide a semantics for (R10) that makes it equivalent to (T10).

There are a number of well-known solutions to this last problem. For example, *Dynamic Predicate Logic* (DPL) (Groenendijk and Stokhof, 1991) solves this problem by appeal to a dynamic treatment of existential quantification and conjunction. The account adopted here builds on DPL, extending it to a language that contains belief operators, though that extension turns out to be non-trivial, as we shall see. But before I present that account, I should mention that my aim here is simply to show that *there is some way* of compositionally implementing our relationist proposal, rather than to argue that this particular way of doing so is superior to potential alternatives. I suspect our relationist insight could be implemented in frameworks other than DPL, such as alternative versions of dynamic semantics or even in a suitably sophisticated static semantic framework. I have chosen to give a concrete implementation of the proposal in an extension of DPL simply because that theory is reasonably well-known and comparatively simple.¹⁴

¹⁴For static approaches to cross-clausal anaphora, see Rothschild (2017) and Mandelkern (2022). For other dynamic approaches, see Heim (1982) and Kamp (1981), among many others. Particularly relevant here are dynamic accounts of modal subordination, such as Roberts (1989) and Brasoveanu (2010). We might ultimately want to implement our account in one of these latter frameworks in order to unify our treatments of intentional identity and modal subordination, but I leave that as a matter for future research.

As an anonymous reviewer observes, I am not the first to apply dynamic semantics to intentional identity sentences. For example, Asher (1987) offers an account of intentional identity in a version of Discourse Representation Theory; see Asher (1986) for background. But Asher’s approach appears to be rather different from the one offered here, built as it is on a philosophy of mind that treats belief as a relation between an agent and a special kind of syntactic object (what he calls a *delineated DRS*). A sentence like (1) is true only if a certain syntactic object (a *reference marker*) used to represent Hob’s state of mind stands in a suitable relation to a similar syntactic object used to represent Nob’s state of mind (Asher, 1987, 151-154). I leave comparing the present account with Asher’s for future work.

The full details of our account are presented in an appendix, but we can informally outline some of the basic ideas here by indicating how the account predicts that our proposed truth-condition for (1) is necessary for the truth of that sentence, i.e. by sketching how it predicts that $(\mathcal{B}_b \exists x Fx \wedge \mathcal{B}_c Gx)$ is true at a world w only if every indexed possibility w' compatible with what the agents in A believe in w is such that there is an individual that is F in w'_b and G in w'_c .

We begin with some terminology. Let us say that a *point* is a pair (w, g) consisting of an indexed possibility w and a variable assignment g . And let us say that an *accessibility relation* R is a binary relation between a world w_a and a point (w', g) meeting two conditions:

- (i) $w_a R(w', g)$ only if w' is compatible with what the relevant agents in A believe in world w_a , and
- (ii) if w' is compatible with what the relevant agents in A believe in world w_a , then there is some g such that $w_a R(w', g)$.

These two conditions ensure that two different accessibility relations will only differ with respect to what variable assignments can be accessed from world w_a , not with respect to what indexed possibilities can be accessed—every accessibility relation can access from w_a all and only the indexed possibilities compatible with what the agents A believe in w_a .

The interpretation of a clause ϕ is given relative to an indexed possibility w and an agent a , and this denoted by $\llbracket \phi \rrbracket^{w,a}$. In standard DPL, the interpretation of a clause relates an ‘input’ variable assignment to a possible ‘output’ variable assignment. DPL updates are non-deterministic, so that a single input may be related to multiple possible outputs; our account is similarly non-deterministic. On the official version of the theory presented in the appendix, the interpretation of a clause relates an input pair (g, R) consisting of a variable assignment g and an accessibility relation R to an output pair (g', R') . But we can focus here on how belief clauses relate an input accessibility relation R to an output accessibility relation R' , since, in our system, clauses of the form $\mathcal{B}_b \phi$ do not update the input variable assignment.

We’ll say that a sentence ϕ is *true at a world* w_a just in case $\llbracket \phi \rrbracket^{w,a}$ relates an arbitrary input accessibility relation R to some output accessibility relation R'' . So let us suppose that $(\mathcal{B}_b \exists x Fx \wedge \mathcal{B}_c Gx)$ is true at world w_a ; thus, where R is any input accessibility relation, we have that $\llbracket \mathcal{B}_b \exists x Fx \wedge \mathcal{B}_c \rrbracket^{w,a}$ relates R to some R'' . In DPL, a conjunction denotes the composition of the relations denoted by the conjuncts; in our system, this means that since $\llbracket \mathcal{B}_b \exists x Fx \wedge \mathcal{B}_c \rrbracket^{w,a}$ relates R to R'' , there is an R' such that the first conjunct $\llbracket \mathcal{B}_b \exists x Fx \rrbracket^{w,a}$ relates R to R' and the second conjunct $\llbracket \mathcal{B}_c \rrbracket^{w,a}$ relates R' to R'' .

To appreciate our account of the first conjunct, $\mathcal{B}_b \exists x Fx$, it will help to adopt some more notation and terminology. Given a world w_a and an accessibility relation R , we have a set $R(w_a)$ consisting of all and only the points (w', g) such that $w_a R(w', g)$. For any variable α , we’ll say that a point (w', g') is an α -variant of a point (w, g) just in case $w' = w$ and for all variables α' other than α , $g'(\alpha') = g(\alpha')$. Then what $\llbracket \mathcal{B}_b \exists x Fx \rrbracket^{w,a}$ essentially does is it takes the

input set $R(w_a)$ and replaces each point (w', g) in that set with one or more x -variants (w', g') of (w', g) such that $g'(x)$ is F in w'_b . That is, $\llbracket \mathcal{B}_b \exists x Fx \rrbracket^{w, a}$ relates R to R' iff $R'(w_a)$ is the result of such a series of replacements.¹⁵ Thus, if $\llbracket \mathcal{B}_b \exists x Fx \rrbracket^{w, a}$ relates R to R' , at each point $(w', g') \in R'(w_a)$, x will denote something that is F at w'_b . Now, since we are assuming that $\llbracket \mathcal{B}_b \exists x Fx \rrbracket^{w, a}$ does relate R to R' , we have:

(A) Each point $(w', g') \in R'(w_a)$ is such that $g'(x)$ is F at the world w'_b .

We now turn to the second conjunct, $\mathcal{B}_c Gx$. In our system, $\llbracket \mathcal{B}_c Gx \rrbracket^{w, a}$ will relate R' to R'' just in case $R'(w_a) = R''(w_a)$ and each point (w', g) in this set $R'(w_a) = R''(w_a)$ is such that $g(x)$ is G in w'_c . In effect, $\llbracket \mathcal{B}_c Gx \rrbracket^{w, a}$ is just ‘checking’ that at each point (w', g) in the set $R'(w_a) = R''(w_a)$, x denotes something that is G in world w'_c — $\llbracket \mathcal{B}_c Gx \rrbracket^{w, a}$ isn’t updating the points in the input set the way $\llbracket \mathcal{B}_b \exists x Fx \rrbracket^{w, a}$ did. Thus, since we are assuming that $\llbracket \mathcal{B}_c Gx \rrbracket^{w, a}$ does relate R' to R'' , we have that $R'(w_a) = R''(w_a)$ and:

(B) Each point $(w', g') \in R''(w_a) = R'(w_a)$ is such that $g'(x)$ is G at the world w'_c .

Note that claims (A) and (B) imply (C):

(C) Each point $(w', g') \in R'(w_a)$ is such that $g'(x)$ is F at w'_b and $g'(x)$ is G at w'_c .

So the first conjunct examines the set $R(w_a)$ and replaces each point (w', g) in that set with an x -variant (w', g') at which x denotes something that is F in world w'_b . This yields a new set of points $R'(w_a)$. This new set is then fed into the second conjunct which simply checks that at each point (w', g') in the set, x denotes something that is G in world w'_c . Thus, if this procedure yields an output, what it leaves us with in the end is a set of points (w', g') at which x denotes something that is F in w'_b and G in w'_c .

Now suppose w' is compatible with what the agents in A believe in w_a . Then since R' is an accessibility relation, there will be some g such that (w', g) is in $R'(w_a)$; this simply follows from the definition of an accessibility relation given above. Then by (C), $g(x)$ will be F at w'_b and G at w'_c . Thus, there will be an individual o that is F in w'_b and G in w'_c , for $g(x)$ will be such an o . Since w' was an arbitrary indexed possibility compatible with what the agents in A believe in w_a , this holds for all such indexed possibilities—for each indexed possibility w' compatible with what the agents in A believe in w_a , there is an individual o that is F in w'_b and G in w'_c . This gives us the left-to-right direction of our proposed truth-condition.

The full details of our proposal can be found in the appendix, which also demonstrates that our proposed truth-condition is sufficient for the truth of (1).

¹⁵For example, suppose $R(w_a)$ contains just two points, (w', g) and (w'', h) , and that Alma is the only F in w'_b and Betty is the only F in w'_c . Then $\llbracket \mathcal{B}_b \exists x Fx \rrbracket^{w, a}$ will relate R to R' just in case $R'(w_a)$ also contains (exactly) two points: (w', g') , the x -variant of (w', g) such that $g'(x)$ is Alma, and (w'', h') , the x -variant of (w'', h) such that $h'(x)$ is Betty.

4. Conclusion

I have argued that our relationist account of intentional identity avoids many of the problems facing non-relationist accounts (§§2 and 3.2). I have also argued that there is direct motivation for relationism, arising out of pairs of cases like the Connected and Unconnected Cases (§3.1). Furthermore, once we adopt the relationist perspective, a number of things seem to fall neatly into place. We obtain a natural account of the truth-conditions of sentences like (1) (§3.2), and get a cleaner separation between the semantics and metasemantics of intentional identity (§3.3). Finally, I have just now been indicating how we can provide a compositional semantic theory that assigns to sentences like (1) our proposed truth-conditions. Our theory extends a familiar theory of cross-clausal anaphora, and thus brings the phenomenon of intentional identity into closer dialogue with contemporary formal semantics.¹⁶

Appendix

This appendix presents our semantic theory in detail, offers some explanatory remarks on its treatment of belief operators, and then shows how the theory assigns our proposed truth-condition to (1).

Definition 1. Given a non-empty set of agents A , we define a language \mathcal{L}_A . The vocabulary of this language consists of n -ary relation symbols, variables, \neg , \wedge , $\exists x$, and, for each $a \in A$, a belief operator \mathcal{B}_a . The definition of the formulas of \mathcal{L}_A can be gleaned from the recursive semantics below.

Definition 2. A *model* for \mathcal{L}_A is a tuple $M = (W, D, B, I)$ where:

- (1) W is a non-empty set, whose elements we call *worlds*,
- (2) D is a non-empty set, whose elements we call *individuals*,
- (3) B is a relation between worlds $w \in W$ and elements $w \in W^A$ (i.e. $B \subseteq W \times W^A$), and
- (4) I is a function which maps an n -ary relation symbol and a world to a subset of D^n .

The definitions that follow are given relative to a fixed model $M = (W, D, B, I)$ for a fixed language \mathcal{L}_A .

Definition 3. A *variable assignment* is a (total) function from the variables of \mathcal{L}_A into D . We use G to denote the set of variable assignments. If g and h are variable assignments and α a variable, we say that h is an α -variant of g , $h[\alpha]g$, iff for all variables α' other than α , $h(\alpha') = g(\alpha')$.

¹⁶Earlier versions of this material were presented at PhLiP 2022 (Tarrytown, NY) and at the 23rd Amsterdam Colloquium. Thanks to these audiences and to Justin Bledin, Andy Egan, Arc Kocurek, Matthew Mandelkern, Sarah Moss, Craige Roberts, and Frank Veltman. I am also grateful to the Editors of *Mind* and to two anonymous referees.

Definition 4. A binary relation $R \subseteq W \times (W^A \times G)$ is an *accessibility relation* iff: (i) $wR(\mathbf{w}, g)$ only if $wB\mathbf{w}$, and (ii) if $wB\mathbf{w}$, then $wR(\mathbf{w}, g')$, for some g' . We let $R(w) = \{(\mathbf{w}, g) : wR(\mathbf{w}, g)\}$

As noted in §3.4, if R and R' are different accessibility relations, then $R(w)$ and $R'(w)$ will only differ with respect to what variable assignments they determine, not with respect to what indexed possibilities they determine. That is, we have:

$$\{\mathbf{w}' \in W^A : (\mathbf{w}', h) \in R(w), \text{ for some } h\} = \{\mathbf{w}' \in W^A : (\mathbf{w}', h) \in R'(w), \text{ for some } h\}$$

for any R, R' and w . On the other hand, we also have:

$$\text{If } R \neq R', \text{ then } \{h \in G : (\mathbf{w}', h) \in R(w), \text{ for some } \mathbf{w}'\} \neq \{h \in G : (\mathbf{w}', h) \in R'(w), \text{ for some } \mathbf{w}'\}$$

for any R, R' , and w .

In our system, the semantic value of a formula, relative to an indexed possibility and an agent, will be a relation between pairs (g, R) consisting of a variable assignment g and an accessibility relation R . We will write $(g, R) \llbracket \phi \rrbracket^{\mathbf{w}, a}(g', R')$ to mean that $((g, R), (g', R')) \in \llbracket \phi \rrbracket^{\mathbf{w}, a}$. The complexity added by our account to standard DPL will only really be exploited in the clause for the belief operator, which we discuss below.

Definition 5. We now define the semantic value of a formula ϕ relative to an indexed possibility $\mathbf{w} \in W^A$ and an agent $a \in A$. For any $(g, R), (g', R')$:

- (1) $(g, R) \llbracket Fx_1, \dots, x_n \rrbracket^{\mathbf{w}, a}(g', R')$ iff $g = g', R = R'$, and $(g(x_1), \dots, g(x_n)) \in I(F, \mathbf{w}_a)$
- (2) $(g, R) \llbracket \neg\phi \rrbracket^{\mathbf{w}, a}(g', R')$ iff $g = g', R = R'$, and there is no (h, Q) such that $(g, R) \llbracket \phi \rrbracket^{\mathbf{w}, a}(h, Q)$
- (3) $(g, R) \llbracket \phi \wedge \psi \rrbracket^{\mathbf{w}, a}(g', R')$ iff there is an (h, Q) such that $(g, R) \llbracket \phi \rrbracket^{\mathbf{w}, a}(h, Q)$ and $(h, Q) \llbracket \psi \rrbracket^{\mathbf{w}, a}(g', R')$
- (4) $(g, R) \llbracket \exists x\phi \rrbracket^{\mathbf{w}, a}(g', R')$ iff $R = R'$ and there is an h such that $h[x]g$ and $(h, R) \llbracket \phi \rrbracket^{\mathbf{w}, a}(g, R')$
- (5) $(g, R) \llbracket \mathcal{B}_b\phi \rrbracket^{\mathbf{w}, a}(g', R')$ iff $g = g'$ and there is a Q such that:
 - (a) for all $(\mathbf{w}', h') \in Q(\mathbf{w}_a)$, there is an h such that $(\mathbf{w}', h) \in R(\mathbf{w}_a)$ and $(h, Q) \llbracket \phi \rrbracket^{\mathbf{w}', b}(h', R')$, and
 - (b) for all $(\mathbf{w}', h) \in R(\mathbf{w}_a)$, there is an h' such that $(\mathbf{w}', h') \in Q(\mathbf{w}_a)$ and $(h, Q) \llbracket \phi \rrbracket^{\mathbf{w}', b}(h', R')$.

The first four clauses here are essentially the standard DPL clauses lifted into our system. Clause (5) is new, since DPL is defined over a language that lacks belief operators. The easiest way to understand what this clause is saying is by looking at examples. To that end, we will examine how it handles the two conjuncts of (1). We start with the first conjunct and note the following fact:

Lemma 1. $(g, R) \llbracket \mathcal{B}_b \exists x Fx \rrbracket^{w,a}(g', R')$ iff $g = g'$ and

- (a) for all $(w', h') \in R'(w_a)$, there is an h such that $(w', h) \in R(w_a)$, $h'[x]h$ and $h'(x) \in I(F, w'_b)$, and
- (b) for all $(w', h) \in R(w_a)$, there is an h' such that $(w', h') \in R'(w_a)$, $h'[x]h$ and $h'(x) \in I(F, w'_b)$.

This **Lemma** essentially re-states something we discussed earlier in §3.4. For what (b) says is that each point (w', h) in the input set $R(w_a)$ has an x -variant (w', h') at which x denotes something that is F in w'_b , and that one or more of these x -variants can be found in the output set $R'(w_a)$. What (a) says is that $R'(w_a)$ *only* contains points that can be derived from a point in $R(w_a)$ in this way—no other points are added to $R'(w_a)$. Thus, we can say that if $\llbracket \mathcal{B}_b \exists x Fx \rrbracket^{w,a}$ updates R to R' , $R'(w_a)$ is the result of replacing each point in $R(w_a)$ with one or more appropriate x -variants.

Let's see why this **Lemma** holds. The clause for belief tells us that $(g, R) \llbracket \mathcal{B}_b \exists x Fx \rrbracket^{w,a}(g', R')$ holds iff $g = g'$ and there is an accessibility relation Q meeting two conditions:

- (i) for all $(w', h') \in Q(w_a)$, there is an h such that $(w', h) \in R(w_a)$ and $(h, Q) \llbracket \exists x Fx \rrbracket^{w',b}(h', R')$, and
- (ii) for all $(w', h) \in R(w_a)$, there is an h' such that $(w', h') \in Q(w_a)$ and $(h, Q) \llbracket \exists x Fx \rrbracket^{w',b}(h', R')$.

Now as the reader may verify, the clauses for the existential quantifier and atomic formulas tell us that:

$(h, Q) \llbracket \exists x Fx \rrbracket^{w',b}(h', R')$ holds iff $Q = R'$ and h' is an x -variant of h and $h'(x)$ is F in w'_b .

So what $\llbracket \exists x Fx \rrbracket^{w',b}$ essentially does is it updates the input assignment h to an output x -variant h' such that x denotes something that is F at the point (w'_b, h') . Thus, (i) and (ii) become:

- (i') for all $(w', h') \in Q(w_a)$, there is an h such that $(w', h) \in R(w_a)$ and $Q = R'$ and h' is an x -variant of h and $h'(x)$ is F in w'_b , and
- (ii') for all $(w', h) \in R(w_a)$, there is an h' such that $(w', h') \in Q(w_a)$ and $Q = R'$ and h' is an x -variant of h and $h'(x)$ is F in w'_b .

Note that there will be a Q meeting conditions (i') and (ii') iff R' is such a Q . Thus, $(g, R) \llbracket \mathcal{B}_b \exists x Fx \rrbracket^{w,a}(g', R')$ holds iff $g = g'$ and:

- (i'') for all $(w', h') \in R'(w_a)$, there is an h such that $(w', h) \in R(w_a)$ and h' is an x -variant of h and $h'(x)$ is F in w'_b , and
- (ii'') for all $(w', h) \in R(w_a)$, there is an h' such that $(w', h') \in R'(w_a)$ and h' is an x -variant of h and $h'(x)$ is F in w'_b .

which is what our **Lemma** says.¹⁷

Another point discussed in §3.4 is related to a second lemma:

Lemma 2. $(g', R') \llbracket \mathcal{B}_c Gx \rrbracket^{w,a}(g'', R'')$ iff $g' = g''$, $R'(w_a) = R''(w_a)$, and for all $(w', h) \in R'(w_a)$, $h(x) \in I(G, w'_c)$.

What this means is that $\llbracket \mathcal{B}_c Gx \rrbracket^{w,a}$ is essentially ensuring that the set $R'(w_a) = R''(w_a)$ only contains points (w', h) at which x denotes something that is G in w'_c .

To see why this holds, note that our semantics for belief tells us that $(g', R') \llbracket \mathcal{B}_c Gx \rrbracket^{w,a}(g'', R'')$ holds iff $g = g'$ and there is a Q such that:

- (i) for all $(w', h') \in Q(w_a)$, there is an h such that $(w', h) \in R'(w_a)$ and $(h, Q) \llbracket Gx \rrbracket^{w',c}(h', R'')$, and
- (ii) for all $(w', h) \in R'(w_a)$, there is an h' such that $(w', h') \in Q(w_a)$ and $(h, Q) \llbracket Gx \rrbracket^{w',c}(h', R'')$.

And the clause for atomic formulas tells us that:

$(h, Q) \llbracket Gx \rrbracket^{w',c}(h', R'')$ holds iff $h = h'$, $Q = R''$, and $h'(x)$ is G in w'_c .

Thus, (i) and (ii) become:

- (i') for all $(w', h') \in Q(w_a)$, there is an h such that $(w', h) \in R'(w_a)$ and $h = h'$, $Q = R''$, and $h'(x)$ is G in w'_c , and
- (ii') for all $(w', h) \in R'(w_a)$, there is an h' such that $(w', h') \in Q(w_a)$ and $h = h'$, $Q = R''$, and $h'(x)$ is G in w'_c .

So there is Q meeting (i') and (ii') iff R'' is such a Q . So $(g', R') \llbracket \mathcal{B}_c Gx \rrbracket^{w,a}(g'', R'')$ holds iff $g = g'$ and

- (i'') for all $(w', h') \in R''(w_a)$, there is an h such that $(w', h) \in R'(w_a)$ and $h = h'$ and $h'(x)$ is G in w'_c , and
- (ii'') for all $(w', h) \in R'(w_a)$, there is an h' such that $(w', h') \in R''(w_a)$ and $h = h'$ and $h'(x)$ is G in w'_c .

Given the inclusion of $h = h'$ in each of these clauses, these further reduce to:

¹⁷Note the existential quantification over accessibility relations ('...there is an accessibility relation Q ...') in our clause for the belief operator. Since in the above example $\exists xFx$ does not itself update the input accessibility relation, it may look as though this quantification over accessibility relations isn't doing much work. It's true that this isn't doing much work *in this example*, nor does it do much work with respect to our second conjunct, since Gx does not update the input accessibility relation either. But we include it here because we need to allow for cases $\mathcal{B}_b\phi$ in which ϕ *does* update the input accessibility relation. (One such case would be an iterated belief ascription like $\mathcal{B}_b\mathcal{B}_b\exists xFx$.) This is similar to the reason our DPL-inspired clause for the existential quantifier includes the metalanguage existential quantifier over variable assignments '...there is an h ...' (Groenendijk and Stokhof, 1991, 45-46).

(i''') for all $(w', h') \in R''(w_a)$, $(w', h') \in R'(w_a)$ and $h'(x)$ is G in w'_c , and

(ii''') for all $(w', h) \in R'(w_a)$, $(w', h) \in R''(w_a)$ and $h(x)$ is G in w'_c .

Claims (i''') and (ii''') will hold iff $R'(w_a) = R''(w_a)$, and each point (w', h) in $R'(w_a)$ is such that $h(x)$ is G in w'_c . This yields **Lemma 2**.

As we'll now demonstrate, **Lemmas 1** and **2** entail that (1) has our proposed truth-condition, given the clause for conjunction and an appropriate definition of truth-at-a-world. For although we are working in a dynamic system, our real interest is assigning truth-conditions to sentences like (1). So we adopt the following definition of truth-at-a-world:

Definition 6. A formula ϕ is *true at a world w* iff: for any w and a such that $w_a = w$, and any (g, R) , there is a (g', R') such that $(g, R) \Vdash^w \phi^{w,a}(g', R')$.

We are now in a position to show how the present account assigns our proposed truth-condition to (1).

Proposition 1. $(\mathcal{B}_b \exists x Fx \wedge \mathcal{B}_c Gx)$ is true at a world w iff for all indexed possibilities w' compatible with what the agents in A believe in w , there is an $o \in D$ such that $o \in I(F, w'_b)$ and $o \in I(G, w'_c)$.

Proof. Left-to-right.

Suppose $(\mathcal{B}_b \exists x Fx \wedge \mathcal{B}_c Gx)$ is true at w . Then where w and a are such that $w_a = w$, and (g, R) are arbitrary, it follows that there is a (g'', R'') such that $(g, R) \Vdash^w [\mathcal{B}_b \exists x Fx \wedge \mathcal{B}_c Gx]^{w,a}(g'', R'')$. By the clause for conjunction (clause (3) of **Definition 5**), it follows that there is a (g', R') such that:

(a) $(g, R) \Vdash^w [\mathcal{B}_b \exists x Fx]^{w,a}(g', R')$, and

(b) $(g', R') \Vdash^w [\mathcal{B}_c Gx]^{w,a}(g'', R'')$,

Given **Lemma 1a**, (a) implies:

(A) for all $(w', h') \in R'(w_a)$, $h'(x) \in I(F, w'_b)$.

And given **Lemma 2**, (b) implies:

(B) for all $(w', h') \in R'(w_a)$, $h'(x) \in I(G, w'_c)$.

Notice that (A) and (B) together imply:

(C) for all $(w', h') \in R'(w_a)$, $h'(x) \in I(F, w'_b)$ and $h'(x) \in I(G, w'_c)$.

To see how (C) yields the left-to-right direction of our biconditional, let w' be an indexed possibility compatible with what the agents in A believe in w , and recall that $w = w_a$. Since R' is an accessibility relation, it follows from the definition of an accessibility relation that there is an h' such that $(w', h') \in R'(w_a)$, since $w_a B w'$. From claim (C) it follows that $h'(x) \in I(F, w'_b)$ and that $h'(x) \in I(G, w'_c)$. Thus, there is an $o \in D$ such that $o \in I(F, w'_b)$ and $o \in I(G, w'_c)$, for $h'(x)$ is such an o . Since w' was an arbitrary indexed

possibility compatible with what the agents in A believe in w , the *left-to-right* direction of our target claim follows.

Right-to-left.

Suppose that for all indexed possibilities w' compatible with what the agents in A believe in w , there is an $o \in D$ such that $o \in I(F, w'_b)$ and $o \in I(G, w'_c)$. We need to show that there is a (g', R'') such that $(g, R) \llbracket \mathcal{B}_b \exists x Fx \wedge \mathcal{B}_c Gx \rrbracket^{w,a} (g', R'')$, where (g, R) is arbitrary and w and a are any indexed possibility and agent such that $w_a = w$. We show that g is such a g' and that we get an appropriate R'' by defining it as follows:

Definition of R'' .

For any w' and (w', h') :

- (i) if $w' = w_a$, then $(w', h') \in R''(w')$ iff there is an h s.t. $(w', h) \in R(w')$, $h'[x]h$, and $h'(x) \in I(F, w'_b)$ and $h'(x) \in I(G, w'_c)$.
- (ii) if $w' \neq w_a$, then $(w', h') \in R''(w')$ iff $w' B w$.

We leave it to the reader to verify that R'' so defined is indeed an accessibility relation.

Given clause (3) of **Definition 5**, to show that $(g, R) \llbracket \mathcal{B}_b \exists x Fx \wedge \mathcal{B}_c Gx \rrbracket^{w,a} (g, R'')$, we must show that there is a (g', R') such that:

- (1) $(g, R) \llbracket \mathcal{B}_b \exists x Fx \rrbracket^{w,a} (g', R')$, and
- (2) $(g', R') \llbracket \mathcal{B}_c Gx \rrbracket^{w,a} (g, R'')$

We show that (g, R'') is such a (g', R') .

Given **Lemma 1**, we can establish that $(g, R) \llbracket \mathcal{B}_b \exists x Fx \rrbracket^{w,a} (g, R'')$ holds by noting that $g = g$ and showing:

- (a) for all $(w', h') \in R''(w_a)$, there is an h such that $(w', h) \in R(w_a)$, $h'[x]h$ and $h'(x) \in I(F, w'_b)$, and
- (b) for all $(w', h) \in R(w_a)$, there is an h' such that $(w', h') \in R''(w_a)$, $h'[x]h$ and $h'(x) \in I(F, w'_b)$.

Claim (a) follows immediately from part (i) of the definition of R'' . For (b), let (w', h) be an element of $R(w_a)$. It follows that $w_a B w'$. So by our right-to-left hypothesis, there is an $o \in D$ such that $o \in I(F, w'_b)$ and $o \in I(G, w'_c)$. Let j be that x -variant of h such that $j(x) = o$. Then $(w', h) \in R(w_a)$, $j[x]h$, and $j(x) \in I(F, w'_b)$ and $j(x) \in I(G, w'_c)$. It follows from clause (i) of the definition of R'' that $(w', j) \in R''(w_a)$. Thus, there is an h' such that $(w', h') \in R''(w_a)$, $h'[x]h$ and $h'(x) \in I(F, w'_b)$, for j is such an h' .

Given **Lemma 2**, we can show that $(g, R'') \llbracket \mathcal{B}_c Gx \rrbracket^{w,a} (g, R'')$ holds by noting that $g = g$, $R''(w_a) = R''(w_a)$, and that every $(w', h') \in R''(w_a)$ is such that $h'(x) \in I(G, w'_c)$. Note that the claim that every $(w', h') \in R''(w_a)$ is such that $h'(x) \in I(G, w'_c)$ follows immediately from clause (i) of the definition of R'' . \square

References

- Asher, Nicholas 1986, 'Belief in Discourse Representation Theory', in *Journal of Philosophical Logic* 15, 127–189
- 1987, 'A Typology for Attitude Verbs and their Anaphoric Properties', in *Linguistics and Philosophy* 10, 125–197
- Azzouni, Jody 2013, 'Hobnobbing with the Nonexistent', in *Inquiry* 56, 340–358
- Brasoveanu, Adrian 2010, 'Decomposing Modal Quantification', in *Journal of Semantics* 27, 437–527
- Braun, David 2012, 'Hob, Nob, and Mythical Witches', in Joseph Keim Campbell, Michael O'Rourke, and Harry S. Silverstein (eds.), *Reference and Referring* (Cambridge, MA: MIT Press), pp. 149–87
- Burge, Tyler 1979, 'Sinning Against Frege', in *Philosophical Review* 88, 398 – 432
- Cooper, Robin 1979, 'The Interpretation of Pronouns', in F. Heny and H. Schnelle (eds.), *Syntax and Semantics 10* (New York: Academic Press), pp. 61–92
- Cumming, Samuel 2014, 'Indefinites and Intentional Identity', in *Philosophical Studies* 168, 371–395
- Edelberg, Walter 1986, 'A New Puzzle about Intentional Identity', in *Journal of Philosophical Logic* 15, 1–25
- 1992, 'Intentional Identity and the Attitudes', in *Linguistics and Philosophy* 15, 561–596
- 1995, 'A Perspectivalist Semantics for the Attitudes', in *Noûs* 29, 316–342
- Elbourne, Paul 2005, *Situations and Individuals* (Cambridge, MA: MIT Press)
- Evans, Gareth 1977, 'Pronouns, Quantifiers, and Relative Clauses (1)', in *Canadian Journal of Philosophy* 7, 467–536
- Fine, Kit 2009, *Semantic Relationism* (Oxford: Blackwell)
- Geach, Peter T 1967, 'Intentional Identity', in *Journal of Philosophy* 64, 627–632
- Glick, Ephraim 2012, 'A Modal Approach to Intentional Identity', in *Noûs* 46, 386–399
- Groenendijk, Jeroen and Stokhof, Martin 1991, 'Dynamic Predicate Logic', in *Linguistics and Philosophy* 14, 39–100

- Heim, Irene 1982, *The Semantics of Definite and Indefinite Noun Phrases*, Ph.D. thesis, University of Massachusetts, Amherst
- 1990, ‘E-Type Pronouns and Donkey Anaphora’, in *Linguistics and Philosophy* 13, 137–177
- Hintikka, Jaakko 1962, *Knowledge and Belief* (Ithaca, N.Y.: Cornell University Press)
- Kamp, Hans 1981, ‘A Theory of Truth and Semantic Representation’, in J.A.G. Groenendijk, T.M.V. Janssen, and M.B.J. Stokhof (eds.), *Formal Methods in the Study of Language* (Amsterdam: Mathematisch Centrum), pp. 277–322
- King, Jeffrey C 1993, ‘Intentional Identity Generalized’, in *Journal of Philosophical Logic* 22, 61–93
- Lanier, William 2013, *Intentional Identity*, Ph.D. thesis, Oxford University
- 2014, ‘Intentional Identity and Descriptions’, in *Philosophical Studies* 170, 289–302
- Mandelkern, Matthew 2022, ‘Witnesses’, in *Linguistics and Philosophy* 45, 1091–1117
- Putnam, Hilary 1973, ‘Meaning and Reference’, in *Journal of Philosophy* 70, 699–711
- Roberts, Craige 1989, ‘Modal Subordination and Pronominal Anaphora in Discourse’, in *Linguistics and Philosophy* 12, 683–721
- Rothschild, Daniel 2017, ‘A Trivalent Account of Anaphora and Presupposition’, in *Proceedings of the 21st Amsterdam Colloquium*, pp. 1–13
- Salmon, Nathan 1998, ‘Nonexistence’, in *Noûs* 32, 277–319
- 2002, ‘Mythical Objects’, in *Metaphysics, Mathematics, and Meaning: Philosophical Papers, Volume I* (Oxford: Oxford University Press)
- Sandgren, Alexander 2018, ‘Which Witch is Which? Exotic Objects and Intentional Identity’, in *Synthese* 195, 721–739
- 2019, ‘A Metarepresentational Theory of Intentional Identity’, in *Synthese* 196, 3677–3695
- Stalnaker, Robert 1984, *Inquiry* (Cambridge, MA: MIT Press)