An object-centric solution to Edelberg's puzzles of intentional identity

Eugene Ho

27 July 2022

This is a preprint. Please cite published version.

Abstract

My belief that Socrates was wise, and your belief that Socrates was mortal can be said to have a common focus, insofar as both these thoughts are about Socrates. In Peter Geach's terminology, the objects of our beliefs bear the feature of intentional identity, because our beliefs share the same putative target. But what if it turned out that Socrates never existed? Can a pair of thoughts share a common focus if the object both thoughts are about, does not actually, really exist? Object-centric accounts of intentionality which explain the aboutness or directedness of thought in terms of the intentional object the thought in question is about, contend that thoughts which share a common focus do so in virtue of both thoughts simply being about the same intentional object. However, Alexander Sandgren contends that such theories face difficulties in explaining a puzzle of intentional identity put forward by Walter Edelberg, in which a pair of sentences seem to differ in truth value but are purportedly logically equivalent on the object-centric theory. If this is right, then it seems that any account which explains intentionality with reference to an intentional object is threatened by this result, whether this object be abstract, merely possible, Meinongian, or otherwise. In this paper, I argue that Edelberg's Puzzle is analogous to Frege's Puzzle and the same tools conventionally used to solve Frege's Puzzle can be used to solve Edelberg's Puzzle. I then

propose a new object-centric solution to Edelberg's Puzzle which takes into account modes of presentation and which is able to accommodate all the relevant linguistic data.

1 Introduction

Geach (1967) suggests that the objects of a pair of beliefs bear the feature of intentional identity when the beliefs in question share the same putative target and presents a puzzle regarding formulating statements of intentional identity when a pair of thoughts involve an apparently inexistent target. Object-centric accounts of intentionality explain the phenomenon of intentional identity in terms of the real identity of the intentional objects of the thoughts in question. In my usage, the term "object-centric" is an umbrella term applying to any account of intentionality which explains the aboutness of thought with reference to an intentional object, regardless of the nature of the intentional object in question. And so, paradigmatic examples of such object-centric accounts are Salmon (2005), Priest (2016), and Parsons (1974), but the object-centric approach has also been advocated more recently by Manning (2015), Woodling (2016), Voltolini (2017), and Marchesi (2021).

However, Edelberg (1986) constructs a harder puzzle of intentional identity not addressed by any of the above accounts. Edelberg's Puzzle presents a case of disagreement between two agents such that a pair of statements of intentional identity which intuitively differ in truth value are logically equivalent on the standard object-centric solution to the puzzle. Sandgren (2019) suggests that object-centric accounts face difficulties in accounting for Edelberg's Puzzle. If Sandgren is right, it seems then that essentially any account of intentionality which explains the aboutness or directedness of thought in terms of an intentional object is prima facie threatened by this result since all that is required to generate Edelberg's Puzzle is mere disagreement between two agents about whether their thoughts are indeed about the same object.

In this paper, I first motivate object-centric accounts of intentional identity by presenting some theoretical benefits of adopting their general strategy of

formalizing statements of intentional identity as statements about the same intentional object. Next, I present Edelberg's Puzzle of intentional identity which purportedly presents problems for object-centric accounts. I then survey some attempted solutions to Edelberg's Puzzle which are ultimately unsuccessful, in part because of a version of Edelberg's Puzzle analogous to Frege's Puzzle. Following which, I argue that the same solutions conventionally employed to solve Frege's Puzzle can be used to solve Edelberg's Puzzle, and thus propose a new truth conditional analysis of statements of intentional identity which takes into account modes of presentation, and which can accommodate all the relevant linguistic data.

2 Intentional identity and intentional objects

In this section I present Geach's Puzzle which involves the difficulties in formulating a notional reading for statements of intentional identity. I then explain and motivate the object-centric theorist's general account of statements of intentional identity, which employs a relational reading for attitude reports involving intentional identity.

Geach's Puzzle of intentional identity concerns the puzzle of formulating sentences which report thoughts which share a common focus. To help us visualize such a situation, let us consider the following case:

Agreement

> Smith is found dead with a bullet wound in the chest. Detectives Oneskey and Twoskey inspect the crime scene and infer that someone murdered Smith. The next day, Jones is also found dead with a similar bullet wound in the chest. Oneskey and Twoskey inspect the crime scene and also infer that someone murdered Jones. Reflecting on the similarities between the cases, Oneskey now thinks that the man who murdered Smith is the same person as the man

who murdered Jones. [1]

In this context the following utterance seems true:

CLAIM-A: Twoskey believes that someone murdered Smith, and Oneskey believes that he murdered Jones.

CLAIM-A is a propositional attitude report which has the form of a paradigmatic case of intentional identity. Oneskey's and Twoskey's beliefs share a common focus because both their beliefs share the same putative target. However, how ought we to formalize CLAIM-A?

2.1 Difficulties of quantifying in

If we wish to quantify into intensional contexts, one option might be to consider analyzing beliefs as doxastic modals and the indefinite determiner as expressing an existential quantifier. However, this raises the question of choosing between a relational (or *de re*) reading and a notional (or *de dicto*) reading of the quantifier.

Since the objects of belief can often fail to actually concretely, or really exist, let us first attempt to formalize CLAIM-A using a notional reading which typically does not commit one to the existence of the intentional object of belief. Using 'B_1' and 'B_2' to denote the doxastic modals "Oneskey believes that..." and "Twoskey believes that...", and 'S' and 'J' to denote the one-place predicates "... murdered Smith" and "... murdered Jones" respectively, we could try:

UNBOUND-A: $B_2[\exists x Sx] \wedge B_1[Jx]$

but here, the 'x' which occurs in the scope of Oneskey's belief operator denotes a variable which is not bound to any quantifier, making UNBOUND-A an open formula. While formulas with free variables are sometimes assigned truth

1. The original puzzle of intentional identity is from Geach (1967) but I have adapted some details from Edelberg (1986) whose examples I focus on in the bulk of the paper. I also follow Cumming (2014) in renaming the detectives as Oneskey and Twoskey for mnemonic purposes so that the names introduced are less distracting (than Hob, Nob, Bob, and Cob).

values relative to assignments, this doesn't help the reading because we need a formalization that specifically tracks the truth value of CLAIM-A.

We could bind the variable to another quantifier within the scope of Oneskey's belief operator like so:

UNFOCUSED-A: $B_2[\exists x Sx] \wedge B_1[\exists y Jy]$

but this won't work because this now fails to capture the fact that Oneskey and Twoskey are thinking about the same thing. We could try adding another stipulation that $x = y$:

IDENTITY-A: $B_2[\exists x Sx] \wedge B_1[\exists y(Jy \wedge x = y)]$

but this still doesn't work since $x$is once again unbound. The problem for formalizing notional statements of intentional identity is now clear. For x to be identical to y, the variables x and y must presumably both occur in a context where both variables are bound, but there is no way to do this if the quantifiers $\exists x$ and $\exists y$ occur in the scope of different belief operators. Geach's Puzzle can thus be understood as the puzzle of finding an appropriate formalization for statements of intentional identity, if we want to quantify into intensional contexts.

2.2 The relational reading and Edelberg's puzzle

Given the difficulties in finding a notional reading for such statements of intentional identity, the proponents of object-centric accounts of intentionality such as Salmon, Priest and Parsons, propose to solve Geach's Puzzle by using a relational reading in which a particular (or existential) quantifier takes wide scope over the conjunction of both agent's beliefs, like so (2005, pp. 96–108; 2016, p. 62; and 1974, p. 578):

QUANT-WIDE-A: $\exists x(B_2[Sx] \wedge B_1[Jx])$

This formalization promises certain attractive explanatory benefits for those who are unsatisfied with treating thoughts as unanalyzable, referentially opaque

contexts, while solving (sidestepping) the difficulties of finding an acceptable notional reading. By allowing quantification over intentional objects, it offers a schema for understanding the logical form of statements of intentional identity and captures the intuitive idea that intentional identity just is identity of an intentional object. If thoughts with a common focus are, as the object-centric theorist contends, thoughts about the same intentional object, then the necessary and sufficient conditions for intentional identity are clear. As Voltolini (2017, p. 294) puts it, a pair of thoughts are about the same thing if and only if the object of the first thought is identical to the object of the second thought, and there really is such an object.[2]

2. Voltolini's use of the quantifier "there is" is ontologically committing. However, I depart from Voltolini's criterion in how we read the quantifier. Voltolini identifies a larger set of relations of thought, called "thinking the same", of which intentional identity, the phenomenon of a pair of thoughts being about the same intentional object (the same intentionale) is merely a proper subset of "thinking the same". A pair of thoughts then are "thinking the same" if and only if, either, (1a) the pair of thoughts are coindividuated by the same intentional object (the same intentionale), such that the object of the first and second thoughts satisfy the metaphysical criterion of identity for that thing – that thing belongs to a certain metaphysical kind and things of that kind are identical iff (here follows what states the relevant identity criterion) and (1b) there really is (exists) such an object in the overall ontological inventory. (1a and 1b) jointly yield intentional identity properly speaking), or (2a) a pair of thoughts are coindividuated by the same proposition-like intentional content, and (2b) there fictionally is such an object (we pretend that there exists such an object). (2a and 2b) jointly yield a pair of mock thoughts which are "about the same thing", albeit only fictionally.). And so satisfying either (1a) and (1b) or (2a) and (2b) would both count as "thinking the same", but in my taxonomy, only satisfying (1a) and (1b) is a case of intentional identity properly speaking. I am largely in agreement with Voltolini that intentional identity properly speaking occurs iff a plurality of thoughts are about the same intentionale, but I depart from Voltolini's criterion in some important respects. The criterion I ultimately endorse is that: a pair of thoughts are about the same thing (the objects of their beliefs have intentional identity) iff there is something that both thoughts are about (both thoughts coindividuated by the same intentionale). This seems almost the same, but it is subtly different. My criterion makes no presuppositions about the metaphysical nature of these intentional objects as existent objects. Here, my use of the quantifier "there is" is ontologically neutral—that is, I don't think what we are committed to saying exists, exhausts the domain of our quantification. In my usage, the quantifier allows as a witness anything at all, existent or non-existent. And so, for me, "thinking about the same" is not a disjunctive phenomenon. A pair of thoughts are really about the same thing iff there is something that both thoughts are about, and thoughts about the same non-existent intentionale are genuine cases of intentional identity. But perhaps this difference is merely taxonomical. I am happy to adopt a version of Voltolini's criterion so long as the quantifier is read in an ontologically neutral way.

At this juncture, it is almost customary to object that the intentional object in question might not actually concretely, or even really exist (Sandgren 2018). Suppose that Smith and Jones were not murdered, but were killed in bizarre freak accidents involving stray bullets. In such a case, CLAIM-A would still be a felicitous utterance despite there not being a real murderer responsible for the deaths of Smith and Jones. QUANT-WIDE-A is thus not an acceptable formalization of CLAIM-A, or so the objection goes.

Here, the object-centric theorists typically appeal to considerations of explanatory indispensability and fruitfulness, suggesting that existentially neutral quantification is ubiquitous in natural language, which already frequently quantifies over non-actual, non-existent objects, suggesting that we have strong theoretical and linguistic reasons to admit such entities into our discourse (Manning 2015; Priest 2016; Salmon 2005).

I am convinced by the object-centric theorist of the need to quantify over intentional objects, if only in a deflationary way. And so my intention here in this paper is not to argue that the object-centric theory is preferable to other theories of intentionality, nor do I commit myself here to a theory of the metaphysical nature of these intentional objects. Rather, my aim is to present a truth conditional analysis of statements of intentional identity which addresses a harder puzzle of intentional identity discovered by Edelberg which seems to prove difficult for the object-centric theorist. Consider:

Disagreement

> Smith is found dead with a bullet wound in the chest. Detectives Oneskey and Twoskey inspect the crime scene and infer that someone murdered Smith. The next day, Jones is also found dead with a similar bullet wound in the chest. Oneskey and Twoskey inspect the crime scene and also infer that someone murdered Jones. Reflecting on the similarities between the cases, Oneskey now thinks that the man who murdered Smith is the same person as the man who murdered Jones. Twoskey disagrees. He thinks Smith and Jones were murdered by two different people. The two detectives

argue heatedly about whether the man who murdered Smith is the same person as the man who murdered Jones.

As before, it seems true to utter:

CLAIM-A: Twoskey believes that someone murdered Smith, and Oneskey believes that he murdered Jones.

However, consider this minimal change to CLAIM-A:

CLAIM-B: Oneskey believes that someone murdered Jones and Twoskey believes that he murdered Smith.

I take it that it is clear to any competent speaker of English that CLAIM-B at least sounds false in Disagreement. Twoskey certainly vocally denies believing that the person Oneskey believes murdered Jones, murdered Smith. And so the second conjunct of CLAIM-B must be false. However, formalizing claims A and B using the schema provided above yields:

QUANT-WIDE-A: $\exists x(B_2[Sx] \wedge B_1[Jx])$

QUANT-WIDE-B: $\exists x(B_1[Jx] \wedge B_2[Sx])$

Recall that it seemed uncontroversial that CLAIM-A is true, and CLAIM-B is false. But now note how QUANT-WIDE-A and QUANT-WIDE-B are equivalent by the commutativity of conjunction. How can a schema like QUANT-WIDE be an adequate formalization of claims A and B if CLAIM-A and CLAIM-B diverge in truth value while QUANT-WIDE-A and QUANT-WIDE-B are logically equivalent?

We now see the problem Edelberg's Puzzle presents for object-centric accounts like the one presented above. Even if we grant that we can quantify over entities which do not actually concretely, or really exist, it seems that so long as we rely on the relational reading the object-centric account favours (with the particular quantifier taking wide scope over both conjuncts), we can construct a pair of sentences which intuitively differ in truth value, despite them being logically equivalent on the object-centric formalization. If this means that we can no longer account for intentional identity in terms of the

real identity of the intentional object both thoughts are about, then the debate about whether we should admit intentional objects into our discourse on grounds of explanatory indispensability is beside the point. Regardless of the nature of these intentional objects, be they mythical, merely possible, Meinongian, or otherwise, if they fail to help us explain the aboutness or directedness of thought the moment disagreement arises, then it seems that any account of intentionality which employs intentional objects of some kind or other is threatened by this result. It is thus especially pertinent for any object-centric theorist of intentionality that Edelberg's Puzzle be solved, as any arguments for admitting intentional objects into our discourse fail if these intentional objects are unable to do the explanatory work they were purported to do.

3 Some non-solutions and some lessons

In this section, I consider four intuitive but ultimately unsuccessful proposals which are motivated by an intuition about this puzzle I find compelling and use these findings to motivate my own proposal.

The intuition goes thus: "Clearly the reason why CLAIM-A and CLAIM-B diverge in truth value is because the pronoun "he" in the second conjunct refers to different putative targets in CLAIM-A and CLAIM-B! In CLAIM-A, "he" functions anaphorically on the indefinite "someone" to contribute as its semantic value, the denotation of "Smith's murderer", while in CLAIM-B, the semantic value of "he" contributes the denotation of "Jones' murderer". Since Oneskey does believe that Smith's murderer murdered Jones, but Twoskey does not believe that Jones' murderer murdered Smith, it's no surprise that CLAIM-A and CLAIM-B diverge in truth value!". This intuition seems to get something right. But how should we implement such a solution?

3.1 Suggestion 1: pronoun expresses definite description

When we say that the pronoun contributes as its semantic value, the denotation of 'Smith's murderer' and 'Jones' murderer', do we mean that the pronoun variably abbreviates the definite description "the murderer of Smith" and "the murderer of Jones" in CLAIM-A and CLAIM-B respectively? Applying the usual Russellian analysis of definite descriptions yields (Russell 1905, p. 490):

DESC-PRONOUN-A: $B_2[\exists x(Sx)] \wedge B_1[\exists y(Sy \wedge \forall z(Sz \rightarrow z = y) \wedge Jy)]$

DESC-PRONOUN-B: $B_1[\exists x(Jx)] \wedge B_2[\exists y(Jy \wedge \forall z(Jz \rightarrow z = y) \wedge Sy)]$

Geach and Edelberg both dub this approach the 'pronouns of laziness' reading, because it contends that the pronoun 'he' expresses a definite description determined by the indefinite noun phrase, 'someone', the pronoun is preceded by (1967, p. 630; 1986, p. 14). However, the pronouns of laziness proposal won't work because, as Geach and Edelberg both note, we can easily amend the case such that the definite description embedded in the second conjunct is unavailable to the agents in question. Suppose that Oneskey learns that Smith was shot but is not aware that Smith has succumbed to his injuries. CLAIM-A is still felicitous to assert in the amended case, since Oneskey does seem to be thinking about the same object Twoskey thinks murdered Smith—Smith's shooter—and believing of him, that he murdered Jones. However, it seems wrong to formalize CLAIM-A as DESC-PRONOUN-A because such a formalization attributes to Oneskey the (*de dicto*) belief that the person who murdered Smith, murdered Jones, when this is implausible given that Oneskey is ignorant of the fact that Smith was murdered.

3.2 Suggestion 2: pronoun expresses speaker's referent

Perhaps we should just take the pronoun to contribute as its semantic value, the speaker's referent of the preceding indefinite noun phrase—that is, the referent the speaker has in mind in uttering the preceding indefinite 'someone', like so:

REF-PRONOUN-A: $B_2[\exists x(Sx)] \wedge B_1[J\alpha]$
REF-PRONOUN-B: $B_1[\exists x(Jx)] \wedge B_2[J\beta]$

Where 'α' and 'β' refer to the entity the speaker has in mind in uttering the indefinite 'someone' in CLAIM-A and CLAIM-B respectively. This reading essentially assumes the truth of the Donnellan–Kripke account of referential pronouns, which suggests that when an anaphoric pronoun is preceded by an indefinite noun phrase, its semantic value is the speaker's referent of the preceding indefinite noun phrase (1966, p. 137; 1977, p. 275). But since Smith's murderer and Jones' murderer might both not exist, let 'α' and 'β' respectively refer to the intentional objects *Smith's murderer* and *Jones' murderer* of Twoskey's and Oneskey's mistaken theorizing instead. REF-PRONOUN-A and REF-PRONOUN-B thus say that 'he' refers to *Smith's murderer* in CLAIM-A but refers to *Jones' murderer* in CLAIM-B.

But this can't be right, because, as Cumming notes, REF-PRONOUN A and B do not require that Twoskey and Oneskey have thoughts about the same putative target (2014, p. 378). Suppose that the speaker mistakenly believes that Twoskey and Oneskey both have in mind Hannibal, when in reality, Twoskey believes that Dolarhyde murdered Smith, and Oneskey believes that Hannibal murdered Jones. REF-PRONOUN-A predicts that CLAIM-A is true even when Twoskey and Oneskey have thoughts about different targets, since the second conjunct is still true in such a case. However, CLAIM-A is not plausibly true when Twoskey and Oneskey have thoughts about different targets. REF-PRONOUN-A is thus not plausibly an adequate formalization of CLAIM-A.

3.3 Suggestion 3: indefinite expresses speaker's referent

To reinstate intentional identity between Oneskey's and Twoskey's beliefs, one might attempt to specify the referents of both the indefinite and the pronoun as the speaker's referent. If one finds the stipulation that both the indefinite and the pronoun are read referentially to be ad hoc (a view which I am sympa-

thetic to), we can alternatively think of the indefinite as being referential and the pronoun as being anaphoric on the indefinite as this still yields the same formalization. Given that Fodor and Sag (1982) argue for the possibility of a referential reading of indefinite noun phrases, I think that the suggestion that indefinite noun phrases sometimes have referential readings are at least prima facie plausible, and thus worth considering, even if they are (foreshadowing a little) not ultimately successful. In any case, the referential reading of the indefinite yields:

REF-INDEF-A: $B_2[S\alpha] \wedge B_1[J\alpha]$

REF-INDEF-B: $B_1[J\beta] \wedge B_2[S\beta]$

where 'α' and 'β' again respectively refer to the intentional objects *Smith's murderer* and *Jones' murderer* of Twoskey's and Oneskey's theorizing.

REF-INDEF only works if the indefinite noun phrase has a referential reading, but Cumming and King highlight some pertinent linguistic evidence, which, I think, demonstrates that the indefinite does not behave as REF-INDEF suggests, regardless of how strongly one intends one's utterance of an indefinite noun phrase to be referential (2014, p. 378; 1988, p. 419). To illustrate, let us consider an extension of Disagreement.

Deception

> Twoskey has fooled the speaker into thinking that he (Twoskey) believes that Dolarhyde is responsible for the murder of Smith and Oneskey believes the same person (Dolarhyde) murdered Jones. In reality, Twoskey thinks a different criminal, Hannibal, murdered Smith and Oneskey thinks the same person, Hannibal, murdered Jones.

In this situation, if the speaker utters CLAIM-A, thinking of Dolarhyde:

CLAIM-A: Twoskey believes that someone murdered Smith and Oneskey believes that he murdered Jones.

the referential theory predicts that CLAIM-A should be rendered as:

REF-INDEF-A*:$B_1[J(Dolarhyde)] \wedge B_2[S(Dolarhyde)]$

but it's not plausible that CLAIM-A should be formalized as REF-INDEF-A*, because REF-INDEF-A* is false, but CLAIM-A is technically true. Contrary to what the speaker believes, Twoskey does not believe that Dolarhyde murdered Smith and Oneskey does not believe that Dolarhyde murdered Jones either. So REF-INDEF-A* predicts that CLAIM-A is false. However, Twoskey does believe that someone (namely Hannibal) murdered Smith and Oneskey does also believe that that same someone (Hannibal) murdered Jones. The person whom Oneskey and Twoskey have in mind is just not whom the speaker had in mind. So CLAIM-A is true. But since the truth value of CLAIM-A diverges from that of REF-INDEF-A, REF-INDEF-A cannot be the correct analysis of CLAIM-A.

I find Cumming's objection to the referential reading of the indefinite compelling.

3.4 Suggestion 4: indefinite denotes attributively

Given that the intuitions pointed out by Cumming give us strong reason to think that indefinites just don't work the way the referential reading predicts, it seems that we are forced back into a reading of the indefinite in which its semantic value is not determined by mere referential intent.

A natural suggestion might then be to say that the indefinite expresses the definite descriptions "the murderer of Smith" and "the murderer of Jones" in CLAIM-A and CLAIM-B respectively. But then since the variable is bound to the quantifier in the scope of the first belief attributee's belief operator, there is, as we have seen in Section 1.1, no way for such a variable to appear in the scope of the second belief attributee's belief operator. One might thus naively think that for the pronoun to co-denote with the indefinite, the pronoun itself must also express the same definite description that the indefinite noun phrase expresses. This yields:

DESC-INDEF-A: $B_2[\exists x(Sx \wedge \forall y(Sy \rightarrow y = x))] \wedge B_1[\exists x((Sx \wedge \forall y(Sy \rightarrow y = x)) \wedge Jx)]$

DESC-INDEF-B: $B_1[\exists x(Jx \wedge \forall y(Jy \rightarrow y = x))] \wedge B_2[\exists x((Jx \wedge \forall y(Jy \rightarrow y = x)) \wedge Sx)]$

but this immediately falls prey to the same criticisms which applies to the pronouns of laziness view. It seems that CLAIM-A is still felicitous even if Oneskey is ignorant that Smith has succumbed to his injuries but nonetheless thinks that Smith's assailant (whom, unbeknownst to Oneskey, is such that Twoskey believes of him that he murdered Smith) murdered Jones. But DESC-INDEF-A is false in such a situation—Oneskey is ignorant of Smith's death and thus *a fortiori* cannot believe (*de dicto*) that the murderer of Smith murdered Jones. This means that DESC-INDEF-A cannot be the correct translation of CLAIM-A (since CLAIM-A is true but DESC-INDEF-A is false).

Cumming (2014, pp. 379–380) considers an ingenious modification to the non-referential reading, which he calls the attributive reading. On the attributive reading, the semantic contribution of the indefinite is determined by the linguistic context in a manner analogous to the denotation of an attributive definite description, while the pronoun co-denotes with the indefinite, but without the relevant description appearing in the scope of either attitude verbs. This yields:

ATTR-INDEF-A: $B_2[S\alpha] \wedge B_1[J\alpha]$ (where α is an intentional object that Twoskey thinks murdered Smith).

ATTR-INDEF-B: $B_1[J\beta] \wedge B_2[S\beta]$ (where β is an intentional object that Oneskey thinks murdered Jones).

The relevant description is then determined by the linguistic context—more specifically, it (attributively) denotes the referent of the indefinite in the first conjunct, and so the description in CLAIM-A is "an intentional object that Twoskey thinks murdered Smith" while the description in CLAIM-B is "an intentional object that Oneskey thinks murdered Jones", which intuitively are *Smith's murderer* and *Jones' murderer* respectively.

This seems to get the right results for the cases so far, given that since the description no longer appears in the scope of the attitude verb, it no longer becomes susceptible to the counterexamples levied against the pronouns of laziness readings, and it also successfully accounts for the divergence in truth value between CLAIM-A and CLAIM-B, provided that the witnesses for CLAIM-A and CLAIM-B really are *Smith's murderer* and *Jones' murderer* respectively.

However, there are two problems with the attributive readings.

First, Cumming concedes that *Jones' murderer* is not the only candidate that satisfies the description in ATTR-INDEF-B. *Smith's murderer* is also a viable candidate for being the object of Oneskey's thoughts in the indefinite (2014, p. 380). But if the intentional object is determined in a manner analogous to an attributive use of a definite description, then contrary to our intuitions, CLAIM-B has a true reading since ATTR-INDEF-B gives us no reason for thinking that the object in question is really *Jones' murderer* and not *Smith's murderer*. But perhaps this is a strength of the account, in that it allows us to predict the availability of an alternative reading in which CLAIM-B is true.

However, the second problem is that the attributive readings, and in fact all the readings we have considered so far, are susceptible to a variation of Edelberg's Puzzle analogous to Frege's Puzzle. To illustrate, note first that technically, the case description does not require that Smith's murderer and Jones' murderer both fail to actually concretely exist in order to generate the puzzle. In fact, Edelberg's puzzle (the puzzle of why a pair of statements of intentional identity diverge in truth value) arises even in cases whereby Twoskey is right that there are two murderers, or, more importantly, in cases whereby Oneskey is right that there is just one murderer. There are thus at least three different situations compatible with Disagreement:

1. Situation 1: Twoskey is right. There are two murderers.

2. Situation 2: Oneskey is right. There is only one murderer.

3. Situation 3: Both are wrong. There is no murderer.

ATTR-INDEF seems to have no problems with situation 1 and situation 3 (provided that *Smith's murderer* and *Jones' murderer* are distinct in situation 3), but let us consider situation 2 for a moment and suppose that Hannibal actually murdered both Smith and Jones. Assuming as the attributive reading proposes that the indefinite contributes the denotation of "Smith's murderer" in CLAIM-A and the denotation of "Jones' murderer" in CLAIM-B, when both terms denote Hannibal, the ATTR-INDEF readings yield:

ATTR-INDEF-A1: $B_2[S(Hannibal)] \wedge B_1[J(Hannibal)]$

ATTR-INDEF-B1: $B_1[J(Hannibal)] \wedge B_2[S(Hannibal)]$

But this is problematic since, contrary to intuition, CLAIM-A and CLAIM-B no longer diverge in truth value, which means Edelberg's Puzzle remains unsolved.

3.5 Final Solution: wide scope quantifier with guises

It may already have occurred to the reader that the reason why CLAIM-A and CLAIM-B differ in truth value despite both conjuncts denoting the same object in Disagreement situation 2, is the same reason why a pair of belief reports with co-referential terms differ in truth value in Frege's Puzzle. As such, we can easily construct a variation of Frege's Puzzle which is structurally analogous to that of Edelberg's Puzzle. Consider:

CLAIM-C: Hammurabi believes that the morning star is visible in the morning and Frege believes that it is visible in the evening.

CLAIM-D: Frege believes that the evening star is visible in the evening and Hammurabi believes that it is visible in the morning.

It seems also that both conjuncts in CLAIM-C are about the morning star, and both conjuncts in CLAIM-D are about the evening star, because the noun phrase determines what the anaphoric pronoun is about. Assuming that Frege is aware of the identity fact that the morning star is the evening star, but Hammurabi is not, I take it that CLAIM-C is uncontroversially true but

CLAIM-D is intuitively false. Since Frege is aware of the fact that the terms "the morning star" and "the evening star" co-refer, the second conjunct of CLAIM-C comes out true. By contrast, Hammurabi does not think of the evening star under a "the morning star"-y mode of presentation and, thus, the second conjunct of CLAIM-D is false.

It seems that the same explanation applies to CLAIM-A and CLAIM-B. While Twoskey thinks of Hannibal under the "Smith's murderer"-y mode of presentation, as being responsible for the murder of Smith, Twoskey does not think of Hannibal under the "Jones' murderer"-y mode of presentation, as being responsible for the murder of Smith. But since Oneskey thinks of Hannibal under both the "Smith's murderer"-y mode of presentation and under the "Jones' murderer"-y mode of presentation as being responsible for the murder of Jones, this makes CLAIM-A true but CLAIM-B false if we have good reason to think the guises invoked in CLAIM-A are that of the "Smith's murderer"-y mode of presentation, while the guises invoked in CLAIM-B are that of the "Jones' murderer"-y mode of presentation.

But if this is so, then a natural response would be to just apply some of the theoretical machinery philosophers have already proposed to deal with Frege's Puzzle to deal with the possibility of Oneskey and Twoskey inadvertently having in mind the same intentional object. I don't commit myself to a particular theory of guises here because it seems to me that there are several options that will function just as well for our purposes. Minimally, all I require from a theory of guises here is just that 1) it is not irrational for an agent to affirm and deny the same proposition under different guises, and 2) that the guise invoked for the indefinite noun "someone" in the first conjunct corresponds to the guise invoked for the pronoun "he" in the second conjunct. And so any theory of guises that delivers these results will work for the purposes of solving Edelberg's Puzzle. For an example on how to develop an account of such guises, see Salmon (1986, 1989), Crimmins and Perry (1989), and Zalta (1989, 2001).

As a simple proof of concept, I shall just take belief to be a ternary relation between agents, propositions, and propositional guises, whereby propositional

guises represent ways that an agent can grasp or conceive of a given proposition. I will follow Salmon in calling this ternary relation "BEL". These propositional guises are then constituted by modes of presentation of ordinary objects and properties which represent ways that an agent can grasp or conceive of a given ordinary object or property. These propositional guises are not inherently linguistic but since they are ways in which we grasp propositions, which are the meanings of sentences, they are typically expressible in terms of the sentences the agent would use to express the proposition grasped. And so, if we are right in suggesting that Twoskey thinks of Hannibal under the "Smith's murderer"-y mode of presentation, as being responsible for the murder of Smith, while Twoskey does not think of Hannibal under the "Jones' murderer"-y mode of presentation, as being responsible for the murder of Smith, then spelling out the guises involved should be quite straightforward:

QUANT-GUISED-A: $\exists x(BEL(Twoskey, Sx, f) \wedge BEL(Oneskey, Jx, g))$

QUANT-GUISED-B: $\exists x(BEL(Oneskey, Jx, g') \wedge BEL(Twoskey, Sx, f'))$

Where:

1. f: the "Smith's murderer murdered Smith"-y mode of presentation,
2. g: the "Smith's murderer murdered Jones"-y mode of presentation,
3. f': the "Jones' murderer murdered Smith"-y mode of presentation,
4. g': the "Jones' murderer murdered Jones"-y mode of presentation.

QUANT-GUISED-A is true in Disagreement situation 2 because there is an intentional object which is both such that Twoskey believes of it that it murdered Smith, under a "Smith's murderer murdered Smith"-y mode of presentation, and such that Oneskey believes of it that it murdered Jones, under a "Smith's murderer murdered Jones"-y mode of presentation—namely, the actual murderer of Smith, Hannibal. However, QUANT-GUISED-B is false because its second conjunct is false, given that according to the case, there is nothing that is such that Twoskey believes of it that it murdered Smith,

under a "Jones' murderer murdered Smith"-y mode of presentation. This is because, as before, Twoskey affirms a two-killer theory which contends that Jones' murderer is not responsible for murdering Smith. QUANT-GUISED thus preserves the results we want when there is only one murderer involved.

The QUANT-GUISED readings also get the right results in cases whereby two murderers are involved such as Disagreement situation 1, the case in which Twoskey is correct that there are two different murderers. Let's call Smith's murderer 'Dolarhyde' and Jones' murderer 'Hannibal'. On the object-centric theory, QUANT-GUISED-A is true in Disagreement because there is something which is both such that Twoskey believes of it that it murdered Smith, under a "Smith's murderer murdered Smith"-y mode of presentation, and such that Oneskey believes of it that it murdered Jones, under a "Smith's murderer murdered Jones"-y mode of presentation, namely, Dolarhyde. However, QUANT-GUISED-B is false because its second conjunct is false, given that its implausible that there is something that is such that Twoskey believes of it that it murdered Smith, under a "Jones' murderer murdered Smith"-y mode of presentation. Neither Dolarhyde nor Hannibal will do because, as before, Twoskey affirms a two-killer theory which contends that Jones' murderer is not responsible for murdering Smith. QUANT-GUISED thus also preserves the results we want in cases whereby two murderers are involved.

What about Disagreement situation 3, the case whereby no murderers are involved? The first thing to note is that even if we grant that situation 3 is a case whereby Smith and Jones were not murdered, that is not the same as granting that Oneskey and Twoskey don't have in mind the same putative target when we report their beliefs in CLAIM-A. Ostensibly, Oneskey and Twoskey are still both thinking about Smith's murderer in CLAIM-A—Twoskey believes that "Smith's murderer murdered Smith", and Oneskey believes of the very same putative object, that "Smith's murderer murdered Jones". The only difference between situation 1 and 3 is that the expression "the murderer of Smith" does not denote a real, flesh and blood murderer. However, depending on one's analysis of the semantics of empty terms, this may not prove to be a threat to the QUANT-GUISED account. I

can think of two possible approaches to situation 3.

The first approach assumes a kind of realism about abstracta. Salmon suggests that acts of mistaken theorizing result in the existence of abstract mythical objects which serve as the intentional objects of the thoughts involving these mistakenly postulated entities (2005). If so, then depending on how we spell out the details of Oneskey's and Twoskey's mistaken theorizing when they come to the conclusion that Smith and Jones were murdered, and how Oneskey and Twoskey both come to acquire the terms "Smith's murderer" and "Jones' murderer", we can apply the QUANT-GUISED account mutatis mutandis to situation 3, retaining the same truth conditions as before.

The second approach involves a denial of the Quinean thesis that the quantifier "there is" is existence entailing by distinguishing between the particular quantifier ($\exists x$) from the existence predicate ($E!$). By reading the particular quantifier in an ontologically neutral way and quantifying over both existents and non-existents alike, this approach grants that Smith's murderer and Jones' murderer both do not exist, but still contends that the non-existent object, Smith's murderer, is nonetheless the object of both Oneskey's and Twoskey's beliefs in CLAIM-A. The general picture I have in mind is that the introduction of the terms "Smith's murderer" and "Jones' murderer" functions as something like semantic baptisms, and Oneskey's and Twoskey's dispositions to use the same terms to describe their thoughts suggest that subsequent uses of the same terms trace their reference to the initial baptism (Priest 2016, pp. 207-215). [3] This allows us to straightforwardly account for the relevant data in situation 3, by retaining essentially the same truth conditions as before.[4]

3. One thing to note is that the problem of reference for non-actual objects is not unique to my account—any possibilist account also needs to contend with such a problem. And I think that the same tools used by possibilists to solve this problem (or some analogous counterpart) may very well be adopted here. My characterization of the causal picture is quite rough for I concern myself here mostly with *the truth conditions of attitude reports* themselves rather than the mechanism of reference for non-actual objects more generally.

4. In section 3.1, I generalize the schema for token statements of intentional identity, suggesting that statements of the form "A believes that some F is G and B believes that it is H" is true iff $\exists x(BEL(A, Fx \wedge Gx, f) \wedge BEL(B, Hx, g))$. This approach thus formalizes the QUANT-GUISED readings in essentially the same way as before but $\exists x$ is an outer quantifier that quantifies over both existents and non-existents alike. $\exists x(BEL(A, Fx \wedge$

In this paper, I refrain from committing myself to one approach rather than the other, but I think that the fact that the QUANT-GUISED readings are compatible with both accounts of the metaphysical nature of intentional objects, existent or non-existent, is a strength of the account. The QUANT-GUISED account thus also yields the right results when no real murderers are involved.[5]

4 Objections and replies

4.1 Belief reports and modes of presentation

At this point, one might reasonably object that the counterexamples to the pronouns of laziness readings already show us that having thoughts about the same thing doesn't require thinking about the same thing under the exact same description. And if thinking of an object under a guise is analogous to thinking of an object under a description, then it's unclear how deploying guises avoids criticisms levied against the pronouns of laziness readings.

Let's take a look. I take it that minimally, what we would need to craft Edelberg's anti-laziness modifications would be something like:

Ignorance

> Smith has been shot. Oneskey and Twoskey go and interview Smith. Oneskey leaves to investigate a new incident. Twoskey

$Gx, f) \wedge BEL(B, Hx, g))$ is thus read: There is some (*existent or non-existent*) thing which A believes that it is G, under f: the "the F is G"-y mode of presentation, and B believes that it is H, under g: the "the F is H"-y mode of presentation.

5. The two approaches I propose are not exhaustive, but it seems to me that most object-centric accounts either propose some kind of realism with regard to the existence of abstract objects (adopt the first approach), or propose that quantifying over non-existent objects is philosophically innocuous (adopt the second approach). Of course, one may very well deny both the existence of abstracta, as well as our ability to quantify over non-existents, but if so, it's hard to see how one can still call oneself an object-centric theorist of intentionality when many thoughts are *prima facie* about objects which do not concretely exist in the actual world. Recall that my goal here in this paper is not to persuade the opponent of the object-centric account, but to simply demonstrate that Edelberg's Puzzle poses no special threat to the object-centric account.

> then witnesses Smith succumbing to his injuries. Oneskey arrives at the scene to find that Jones has been shot to death. Reflecting on the similarities in the cases, Oneskey believes that the man who shot Smith is the man who shot (and murdered) Jones.

As before, the following utterance seems true:

CLAIM-A: Twoskey believes that someone murdered Smith, and Oneskey believes that he murdered Jones.

but QUANT-GUISED-A is false because Oneskey does not think of anyone under the "Smith's murder"-y mode of presentation and thus does not affirm any proposition under the "Smith's murderer murdered Jones"-y propositional guise. Doesn't this count as a counterexample to the QUANT-GUISED readings? I don't think so. The reason is that CLAIM-A is not equivalent to QUANT-GUISED-A in the anti-laziness case. CLAIM-A is equivalent to QUANT-GUISED-A in Disagreement because a natural way of spelling out the guises invoked in CLAIM-A in Disagreement would go:

DIS-CLAIM-A: Twoskey believes that someone (Smith's murderer) murdered Smith, and Oneskey believes that he (Smith's murderer) murdered Jones.

but we wouldn't spell out the guises in CLAIM-A in Ignorance the same way because we know Oneskey doesn't think of the same thing Twoskey thinks of as Smith's murderer, under the "Smith's murderer"-y mode of presentation, but by the other mode of presentation, the "Smith's shooter"-y mode of presentation. And so we would instead say:

IGN-CLAIM-A: Twoskey believes that someone (Smith's shooter) murdered Smith, and Oneskey believes that he (Smith's shooter) murdered Jones.

but if this is so then all we have to do is apply the relevant changes, mutatis mutandis like so:

QUANT-GUISED-A*: $\exists x(BEL(Twoskey, Sx, f*) \wedge BEL(Oneskey, Jx, g*))$

Where:

1. $f*$: the "Smith's shooter murdered Smith"-y mode of presentation

2. $g*$: the "Smith's shooter murdered Jones"-y mode of presentation

but QUANT-GUISED-A* is true! And so if CLAIM-A is equivalent to QUANT-GUISED-A* instead of QUANT-GUISED-A in Ignorance, then the objections to the pronouns of laziness cases pose no threat to the QUANT-GUISED readings.

Of course, this then raises the question of how the guises involved in CLAIM-A are determined. I will attempt to spell out a general schema for formalizing attitude reports involving an indefinite and an anaphoric pronoun which attribute beliefs with a common focus to a pair of agents.

For a token attitude report with a common focus:

TOKEN-1: Agent A believes that some F is G and Agent B believes that it is H.

TOKEN-1 is to be analyzed as:

QUANT-GUISED-TOKEN-1: $\exists x(BEL(A, Fx \wedge Gx, f) \wedge BEL(B, Hx, g))$

Where:

1. f: the "the F is G"-y mode of presentation,

2. g: the "the F is H"-y mode of presentation.

It should be noted that the Edelberg sentences don't follow the exact surface grammar as TOKEN-1 since they omit a description from the indefinite noun phrase (using "someone" instead of "some F"). In which case then the guise under which the agent ought to take the object in question is not explicitly specified and is to be determined contextually. I take this to be in line with Perry's and Crimmins' (1989) intuition that belief reports have an unarticulated constituent which is determined contextually by the most relevant candidate mode of presentation of the target object.

And so for the second kind of attitude reports in which no description is embedded in the indefinite:
TOKEN-2: Agent A believes that something is G and Agent B believes that it is H.

TOKEN-2 is to be analyzed as:

QUANT-GUISED-TOKEN-2: $\exists x(BEL(A, Gx, f) \wedge BEL(B, Hx, g))$

Where:

1. f: the "... is G"-y mode of presentation,
2. g: the "... is H"-y mode of presentation.

and the "..." represents the mode of presentation under which x is thought of. I think that what we intuitively fill in the blanks with is contextually determined, and obeys the constraint that the mode of presentation invoked in the pronoun must be the same mode of presentation that is invoked in the indefinite.

Notice how the guises in TOKEN-2 don't seem to be doing much work since the mode of presentation of the target object is not yet specified. If we just look within the scope of either belief operator, x is bound to a quantifier outside the scope of the belief operator. And so we don't yet know how the x which is the target of both beliefs is conceived of by both agents. All we know is that the guise of the object in question invoked in the indefinite will be the same as the objectual guise invoked in the pronoun, whatever that may be.

To motivate this reading, let's take stock of some data we have not covered just yet. Cumming notes that there are less salient readings of the Edelberg sentences in which CLAIM-B is true (2014, p. 380).[6]

6. Cumming provides two paraphrases to help us get to this less salient reading:
(28) There is someone that Onesky thinks murdered Jones; Twosky thinks he murdered Smith.
(i) There is someone that Onesky (now) thinks murdered Jones (namely, Smith's murderer); Twosky thinks he murdered Smith.
As a minimal change, I retain the phrasing of CLAIM-B but simply specify the referent of the indefinite noun phrase 'someone' as 'Smith's murderer'.

STRAINED-B: Oneskey believes that someone (Smith's murderer) murdered Jones and Twoskey believes that he murdered Smith.

I can (with some effort), see how STRAINED-B might be a possible reading of CLAIM-B, given norms about charitably interpreting what your interlocutor is saying as being true. But this poses no threat to my proposal. I think that what is going on in STRAINED-B, is that we are stipulating the guises invoked in CLAIM-B, instead of letting the linguistic context naturally determine the guises invoked. Which is why when primed with Cumming's strained true reading of CLAIM-B, we are able to get a reading in which QUANT-GUISED-B is true.

Applying the QUANT-GUISED analysis yields:

STRAINED-GUISED-B: $\exists x(BEL(Oneskey, Jx, g) \wedge BEL(Twoskey, Sx, f))$

Where

1. g: the "Smith's murderer murdered Jones"-y mode of presentation,
2. f: the "Smith's murderer murdered Smith"-y mode of presentation.

which is true. The flexibility of the QUANT-GUISED readings thus allow us to account for the possibility of less salient readings if we keep in mind the guises being used in the attitude reports.

4.2 Intentional Identity and Belief Reports

At this point, one might again reasonably object that even in Ignorance, surely Twoskey's belief about Hannibal under the "Smith's murderer"-y mode of presentation and Oneskey's belief about the very same thing, Hannibal, under the "Smith's shooter"-y mode of presentation have a common focus. And thus, if the guise invoked in the indefinite is not "Smith's shooter" but "Smith's murderer", then QUANT-GUISED-A predicts CLAIM-A is false when it is intuitively true, which poses a *prima facie* challenge to my account.

In reply, I deny that CLAIM-A remains true if we stipulate that the guise invoked in the indefinite is that of "Smith's murderer", and I account for this data by distinguishing between attitude reports attributing beliefs about an intentionally identical object, and the phenomenon of intentional identity itself.

I agree that Twoskey's belief that Smith's murderer murdered Smith and Oneskey's belief that Smith's shooter murdered Jones have a common focus, because their beliefs are actually about the same object, Hannibal. It thus follows that the object of both their beliefs has the feature of intentional identity. But here I concern myself with a truth conditional analysis of attitude reports of the structure "A believes that some F is G and B believes that it is H", and when considering the linguistic behaviour of such reports, I think it is necessary that the guise of the target object invoked in the belief report of the second conjunct is the same as the guise of the target object invoked in the belief report of the first conjunct.

Recall that I suggested that CLAIM-A and CLAIM-B were structurally analogous to CLAIM-C and CLAIM-D. If it were possible that the guise of the target object invoked in the conjunctive belief report differs between each conjunct, then CLAIM-D should come out true:

CLAIM-D: Frege believes that the evening star is visible in the evening and Hammurabi believes that it is visible in the morning.

This is because Hammurabi's belief that the morning star is visible in the morning is about the very same object that Frege's belief is about—Venus. But unless we are willing to explain away our intuitions that CLAIM-D is false, as being misled by pragmatic features of attitude reports as many direct reference theorists are so eager to do, it seems to me that the contextualist stance which takes seriously our pre-theoretical judgements about the truth conditions of such attitude reports, provides some modest motivation for analyses which do not force CLAIM-D to come out true.

And so, going back to Ignorance, if we stipulate that the guise invoked in the first conjunct of CLAIM-A is a "Smith's murderer"-y mode of presentation

like so:

IGN-CLAIM-A*: Twoskey believes that someone (Smith's murderer) murdered Smith, and Oneskey believes that he murdered Jones.

then, I contend that this is false on its most natural reading, because the second conjunct is false. There is nothing such that Oneskey thinks of it, under a "Smith's murderer"-y mode of presentation in the Ignorance case and thus *a fortiori*, Oneskey does not believe of anything under a "Smith's murderer"-y mode of presentation, that it murdered Jones.

Granted, there is a true simple *de re* reading of IGN-CLAIM-A*, so long as one does not take into account the guises invoked on the most natural reading of the sentence in question, but for reasons already discussed above, I think the invocation of guises is called for given the possibility of inadvertent Frege Puzzles.

Thus, the true reading of attitude reports which attribute beliefs with a common focus to both agents requires that the same guise mobilized in the indefinite is also mobilized in the pronoun, and I think that this feature of statements of intentional identity which employ indefinites and anaphoric pronouns is to be distinguished from the phenomena of intentional identity more generally in which the objects of a pair of beliefs have intentional identity if and only if the beliefs in question are about the same object. In short, I grant that Twoskey's belief that Smith's murderer murdered Smith and Oneskey's belief that Smith's shooter murdered Jones have a common focus, but I deny that the guise under which Twoskey's belief is grasped is thereby the relevant guise which is being used in felicitous utterances of CLAIM-A.

4.3 Cumming's solution to Edelberg's puzzle

Sandgren (2019) contends that object-centric accounts of intentionality have more difficulty accounting for Edelberg's Puzzle, and hopefully I have shown that that is not the case—if I am right, Edelberg's Puzzle presents no special challenge for the object-centric account of intentionality over and above that

of the challenge presented by Frege's Puzzle.

However, Cumming (2014) actually already presents an object-centric solution to Edelberg's Puzzle, making my proposal a rival theory.

Cumming assumes a cognitive architecture of private mental files (2014, p. 382). Some of these files have their referents grounded in causal information links, while some other files are descriptive. Each file denotes a thought-object (2014, p. 382). Two files of different mental agents denote the same thought-object whenever the agents regard those files as coreferential, which can occur even when the files in question are both descriptive but have different rubric conditions (2014, p. 383). Each descriptive file has a distinguished rubric condition that the agent typically believes only one individual satisfies, and the descriptive file may contain further "non-rubric" conditions (2014, p. 383).

On Cumming's view, when an indefinite noun phrase that precedes an anaphoric pronoun is embedded in a statement of intentional identity of the form "A believes something is G, and B believes that it is H", there are two ways to read such statements of intentional identity (2014, p. 386).

On the *de dicto* reading, the indefinite contributes as its semantic value, the thought object denoted by a descriptive mental file, with the description "x is G" as a rubric condition, in a manner analogous to attributive definite descriptions, while the pronoun co-denotes with the indefinite (although this co-denotation does not require using the same rubric condition or definite description).

On the *de re* reading the thought-object in question contributed by the indefinite is distinct from the denotation of the file that has the rubric condition "x is G" but nonetheless is believed by agent A to be G (thus the file denoting the thought-object in question has "x is G" as a non-rubric condition) and is such that B believes of it that it is H.

Cumming then contends that when Oneskey forms the belief that Smith's murderer murdered Jones, the file that denotes *Smith's murderer* has the description "x murdered Jones" not as a rubric condition, but as a non-rubric condition (2014, p. 384), and that the most natural reading of Edelberg's Puzzle (in which CLAIM-A is true and CLAIM-B is false) is the *de dicto*

reading:

CUMMING-DICTO-A: Twoskey believes β murdered Smith, and Oneskey believes β murdered Jones. (Where β is the denotation of Twoskey's file with the rubric condition "x murdered Smith").

CUMMING-DICTO-B: Oneskey believes α murdered Smith, and Twoskey believes α murdered Jones. (Where α is the denotation of Oneskey's file with the rubric condition "x murdered Jones").

While Cumming's strained readings (in which CLAIM-A is false and CLAIM-B is true) correspond to the *de re* readings:

CUMMING-DE-RE-A: There exists a thought object $x \neq \beta$, such that Twoskey thinks x murdered Smith and Oneskey thinks x murdered Jones. (Where β is the denotation of Twoskey's file with the rubric condition "y murdered Smith").

CUMMING-DE-RE-B: There exists a thought object $x \neq \alpha$, such that Oneskey thinks x murdered Jones and Twoskey thinks x murdered Smith. (Where α is the denotation of Oneskey's file with the rubric condition "y murdered Jones").

CUMMING-DICTO-A is true because Twoskey's file with "x murdered Smith" as a rubric condition denotes the thought-object *Smith's murderer* meaning that Twoskey must believe that *Smith's murderer* murdered Smith, and Oneskey also believes that *Smith's murderer* murdered Jones, since Oneskey believes that Smith's murderer murdered Jones and takes his belief that Smith's murderer murdered Jones to be about the same thing as Twoskey's belief that Smith's murderer murdered Smith.

CUMMING-DICTO-B is false because Oneskey's file with "x murdered Jones" as a rubric condition denotes the thought object *Jones' murderer* and not *Smith's murderer*, but it is implausible to think that Twoskey believes that *Jones' murderer* murdered Smith in any way of spelling out the details of the case.

CUMMING-DE-RE-A is false because it seems that there is nothing such

that Twoskey thinks it murdered Smith which is distinct from the denotation of Twoskey's file which has "y murdered Smith" as a rubric condition. *Smith's murderer* will not do since *Smith's murderer* is the denotation of Twoskey's file which has "y murdered Smith" as a rubric condition. And *Jones' murderer* will not do because Twoskey does not believe that Jones' murderer, whoever he is, murdered Smith.

CUMMING-DE-RE-B is true because there is something which is distinct from the denotation of Oneskey's file which has "y murdered Jones" as a rubric condition, and which is such that Oneskey thinks that it murdered Jones and Twoskey thinks that it murdered Smith, namely *Smith's murderer*.

As an aside, while the *de re / de dicto* distinction is standardly taken to be a simple scope distinction, it is interesting to note that Cumming's *de dicto* reading entails the standard de re reading, while Cumming's *de re* reading requires some sophisticated account of mental files and their denotations, which complicates the standard picture for the *de re / de dicto* distinction. That being said, this is not yet an objection to Cumming's view, and it seems that Cumming is able to get the right results for both the natural reading in which CLAIM-A is true and CLAIM-B is false, and the strained reading in which CLAIM-A is false and CLAIM-B is true.

But what about situation 2, the case of the inadvertent Frege Puzzle? Suppose Oneskey is correct that the person who murdered Smith is the person who murdered Jones. As before, let's call him Hannibal. Then instead of thinking of a merely intentional object, it seems that so long as we allow mental reference to be externalist, as Cumming himself admits contemporary accounts of mental files are (2014, p. 382), the denotation of Twoskey's file with the rubric condition "y murdered Smith" and the denotation of Oneskey's file the the rubric condition "y murdered Jones" is identical. They both denote the flesh and blood murderer, Hannibal, not a mere thought-object. This spells a *prima facie* problem for both Cumming's *de dicto* readings which now become equivalent, and *de re* readings which now become both false.

This might be avoided if, as Cumming seems sympathetic to in some passages, Cumming follows Frege in suggesting that terms do not denote their

ordinary denotations when embedded in intensional contexts but instead denote some intermediate entity akin to a Fregean sense, *even in cases whereby the ordinary denotation exists*.

In which case, then a thought-object is not a mere intentional object (not a mere object of thought) but is a pseudo-Fregean entity existing in an intermediate ontological tier between language and the world (2014, p. 374). According to this framework, noun phrases refer to thought-objects in intensional contexts (2014, p. 375) and these thought-objects then determine a unique ordinary object in the actual world, which is the ordinary denotation of the noun phrase when not embedded in intensional contexts, provided that the unembedded noun phrase has a referent in the actual world (2014, p. 374). If we are committed to quantifying over these intermediate entities, it seems that Frege Puzzles don't directly threaten the analysis since, just like how "Hesperus" and "Phosphorus" denote the same object but have a different Fregean sense, "Smith's murderer" and "Jones' murderer" denote Hannibal in extensional contexts but denote the different thought-objects, *Smith's murderer* and *Jones' murderer* in intensional contexts. This allows CLAIM-A and CLAIM-B to diverge in truth value once more.

But now note how these are substantive and non-trivial philosophical commitments about the nature of intentional objects and attitude reports, requiring some form of Fregeanism to be true. In contrast, my account suspends judgement on the ontological status of the intentional objects in question and accommodates naïve Russellian intuitions about the objects of thought. All I assume is that these intentional objects, like ordinary objects, can be thought of under modes of presentation, and I am suggesting that taking these modes of presentation into account as we already do in the Frege Puzzle cases can solve Edelberg's Puzzle.

In summary, my account keeps the simple *de re / de dicto* distinction as a straightforward matter of scope ambiguity, does not require a separate solution to solve Frege's Puzzle, and does not require intentional objects to be objects of a certain kind (eg. intermediate entities such as Edelbergian thought-objects) which, I think, count as points in my favour.

4.4 Accounting for intentional identity in general

Suppose what I've said above convinces you that my proposal is able to adequately account for the puzzling Edelberg sentences. There is, I think, still a question about when a pair of thoughts bear the feature of intentional identity, not just when a statement of the form "A believes that some F is G and B believes that it is H" is true. To answer this, I gesture back to the slogan introduced in section 1: A pair of thoughts are about the same thing if and only if the object of the first thought is identical to the object of the second thought.[7]

However, consider Oneskey's belief that Jones' murderer murdered Jones, and Twoskey's belief that Smith's murderer murdered Smith. Do the objects of this pair of thoughts have the feature of intentional identity? Oneskey says yes, but Twoskey says no. But if intentional identity is nothing over and above the identity of an intentional object, then surely there has to be a matter of fact as to who is right. And indeed there is. If Oneskey's one-killer theory is right, then Oneskey is right that the pair of thoughts share a common focus and if Twoskey's two-killer theory is right, then Twoskey is right that the pair of thoughts do not. But what if Oneskey and Twoskey are both wrong, and there really is no murderer responsible for either Smith's or Jones' deaths? If your metaphysics of accidental fictions commits you to there being some kind of intentional object playing the role of Smith's murderer or Jones' murderer, it would come down to how the intentional objects *Smith's murderer* and *Jones' murderer* are individuated. If they're distinct then Twoskey is right, and if they're not, then Oneskey is right. So it seems that whether Oneskey and Twoskey are in a position to know whether the objects of their thoughts bear the feature of intentional identity, is beside the point. The object-centric theory already provides us clear truth conditions for intentional identity as nothing over and above thinking about the same thing.

7. This formalization is unacceptable to the agent-centric account of intentionality, as such accounts suggest that the defining feature of intentional identity is whether a pair of thoughts have a common focus, rather than whether or not there is an object at the center of that focus. While object-centric theorists say that these are the same, agent-centric theorists deny this. But it is a slogan after all.

But now, consider the scenario in which Twoskey is right. Does Oneskey's belief that Jones' murderer murdered Jones and Oneskey's belief that Smith's murderer murdered Smith bear the feature of intentional identity? Despite what Oneskey himself thinks, the object-centric theory says no. Intentional identity just is the identity of intentional objects. But Oneskey's belief that Jones' murderer murdered Jones and Oneskey's belief that Smith's murderer murdered Smith are about two distinct objects, in the scenario whereby Twoskey is right. Therefore, these beliefs don't really share the feature of intentional identity. But one may reasonably find this to be odd. I take it that the intuition goes something like this: "Intentional identity is the phenomenon of a pair of thoughts being about the same putative target, but surely I cannot be mistaken about whether any pair of my own thoughts share the same putative target!".

But if this is the intuition driving their response, then it seems that we have enough reasons to judge that it is decidedly false, given that we can be and often are mistaken about the contents of our thoughts. In the scenario in which Twoskey is right, Oneskey's belief that Smith's murderer murdered Smith is about Smith's murderer, and Oneskey's belief that Jones' murderer murdered Jones is about Jones' murderer. But Smith's murderer is not identical to Jones' murderer according to the case. Oneskey thinks this pair of thoughts have a common focus only because of his false belief that Smith's murderer is Jones' murderer. But it seems that once Oneskey discovers that his belief was false, Oneskey would proceed to judge that the two thoughts never had a common focus after all. What has a common focus is not the pair <Oneskey's belief that Smith's murderer murdered Smith, Oneskey's belief that Jones' murderer murdered Jones>, but the other pair <Oneskey's belief that Smith's murderer murdered Smith, Oneskey's belief that Smith's murderer murdered Jones>. But even though Oneskey doesn't distinguish between the pair <Oneskey's belief that Jones' murderer murdered Jones, Oneskey's belief that Smith's murderer murdered Jones>, these thoughts do seem to be different after all since they are about different people!

Perhaps this response is question begging. It is, after all, using the object-

centric theory to explain why the intuitions which follow from the object-centric theory are true. But my goal here is not to persuade the agent-centric theorists to adopt an object-centric approach to intentionality, but instead, simply show that the object-centric theory is still viable in light of these purported problems. Hopefully, I have been successful in what I have set out to do.

5 Conclusion

One of my motivations for writing this paper is Sandgren's 2019 paper, in which Sandgren claims that the object-centric account's inability to account for the divergence of truth values of the Edelberg sentences counts as a point against the viability of such theories of intentional identity (2019, 3684). If Sandgren is right, then Edelberg's Puzzle threatens all object-centric accounts of intentionality, so long as they explain the intentionality of thought involving empty terms using intentional objects, regardless of whether these intentional objects are mythical, merely possible, Meinongian, or otherwise. However, my solution to Edelberg's Puzzle which makes no assumptions on the nature of the intentional objects in question is congenial to both possibilism and neo-meinongianism. If my solution is correct, we do not have to give up quantification into intensional contexts, and Edelberg's Puzzle presents no special threat to object-centric accounts of intentional identity.

6 Acknowledgements

I would like to thank everyone at the Department of Philosophy at the National University of Singapore for all the support they have given me. In particular, I would also like to thank Ben Blumson, Bob Beddor, Daniel Waxman, Lavinia Picollo, Tang Weng Hong, the Graduate Research Seminar at the National University of Singapore, and two anonymous referees at Synthese for their comments and for the helpful discussion which has improved the paper greatly.

References

Crimmins, Mark, and John Perry. 1989. The prince and the phone booth: reporting puzzling beliefs. *The Journal of Philosophy* 86, no. 12 (December): 685. https://doi.org/10.2307/2027014.

Cumming, Samuel. 2014. Indefinites and intentional identity. *Philosophical Studies* 168, no. 2 (March): 371–395. Accessed March 16, 2021. https://doi.org/10.1007/s11098-013-0131-9.

Donnellan, Keith. 1966. Reference and definite descriptions. Publisher: [Duke University Press, Philosophical Review], *The Philosophical Review* 75 (3): 281–304. https://doi.org/10.2307/2183143.

Edelberg, Walter. 1986. A new puzzle about intentional identity. Publisher: Springer, *Journal of Philosophical Logic* 15 (1). https://doi.org/10.1007/BF00250547.

Fodor, Janet Dean, and Ivan A. Sag. 1982. Referential and quantificational indefinites. *Linguistics and Philosophy* 5 (3): 355–398. https://doi.org/10.1007/BF00351459.

Geach, Peter. 1967. Intentional identity. Publisher: Journal of Philosophy Inc, *Journal of Philosophy* 64 (20). https://doi.org/10.2307/2024459.

King, Jeffrey C. 1988. Are indefinite descriptions ambiguous? Publisher: Springer, *Philosophical Studies* 53 (3). https://doi.org/10.1007/BF00353515.

Kripke, Saul. 1977. Speaker's reference and semantic reference. *Midwest Studies in Philosophy* 2:255–276. https://doi.org/10.1111/j.1475-4975.1977.tb00045.x.

Manning, Luke. 2015. No identity without an entity: no identity without an entity. *Pacific Philosophical Quarterly* 96, no. 2 (June): 279–305. https://doi.org/10.1111/papq.12074.

Marchesi, Andrea. 2021. A radical relationist solution to the problem of intentional inexistence. *Synthese* 199, no. 3 (December): 7509–7534. https://doi.org/10.1007/s11229-021-03126-3.

Parsons, Terence. 1974. A prolegomenon to meinongian semantics. Publisher: Journal of Philosophy, Inc. *The Journal of Philosophy* 71 (16): 561–580. https://doi.org/10.2307/2025232.

Priest, Graham. 2016. *Towards non-being.* Oxford University Press, August 11, 2016. ISBN: 978-0-19-878359-6. https://doi.org/10.1093/acprof:oso/9780198783596.001.0001.

Russell, Bertrand. 1905. On denoting. Publisher: [Oxford University Press, Mind Association], *Mind* 14 (56): 479–493. https://doi.org/10.1093/mind/xiv.4.479.

Salmon, Nathan. 1986. *Frege's puzzle.* Vol. 96. Issue: 3. Ridgeview.

———. 1989. Illogical belief. Publisher: [Ridgeview Publishing Company, Wiley], *Philosophical Perspectives* 3:243–285. https://doi.org/10.2307/2214270.

———. 2005. Mythical objects (2002). In *Metaphysics, mathematics, and meaning.* Oxford: Oxford University Press. https://doi.org/10.1093/0199284717.003.0004.

Sandgren, Alexander. 2018. Which witch is which? exotic objects and intentional identity. *Synthese* 195, no. 2 (February 1, 2018): 721–739. https://doi.org/10.1007/s11229-016-1237-3.

———. 2019. A metarepresentational theory of intentional identity. *Synthese* 196, no. 9 (September): 3677–3695. Accessed March 16, 2021. https://doi.org/10.1007/s11229-017-1609-3.

Voltolini, Alberto. 2017. (mock-)thinking about the same. *Organon F: Medzinárodný Časopis Pre Analytickú Filozofiu* 24:282–307.

Woodling, Casey. 2016. The indispensability and irreducibility of intentional objects: *Journal of Philosophical Research* 41:543–558. https://doi.org/10.5840/jpr201683093.

Zalta, Edward N. 1989. Singular propositions, abstract constituents, and propositional attitudes. In *Themes from kaplan,* edited by J. Almog, J. Perry, and H. Wettstein. Oxford University Press.

———. 2001. Fregean senses, modes of presentation, and concepts. *Noûs* 35 (s15): 335–359. https://doi.org/https://doi.org/10.1111/0029-4624.35.s15.15.