

Doxastic Semantics and Ways of Believing*

Johannes Stern
Department of Philosophy
University of Bristol
johannes.stern@bristol.ac.uk

Abstract

Taking a discussion of Soames's argument to the effect that semantic contents cannot be identified with sets of truth-supporting circumstances as a starting point we develop a novel semantics for propositional attitude contexts, and more specifically for belief-reports, that implements contextualist ideas within the framework of possible world semantics.

1 Introduction

The semantics of propositional attitude reports have been a testing ground for compositional semantics ever since Frege's first forays into philosophy of language and semantics. In particular, several puzzles and arguments related to propositional attitude reports have been used to argue for more fine grained individuation of semantic content. For example, to name only two, the problem of doxastic closure has been used as an argument against identifying semantic content with a set of possible worlds and Frege's original puzzle about Hesperus and Phosphorus is an argument for identifying semantic content with so-called Fregean propositions or thoughts. These arguments need to be treated with care, as for any such argument in favor of a more fine grained individuation of semantic content, there is a counterargument to the effect that a more fine grained individuation of semantic content is to be resisted. Frequently, the point of disagreement between theorists is whether specific linguistic intuitions ought to be explained semantically, pragmatically, e.g., in terms of implicatures, or, as argued by Williamson (2020), by means of heuristics. Notice that all semanticists will need to refer to pragmatic explanations (and, possibly, heuristics) at some point in explaining their linguistic intuitions: for example, a theorist who individuates semantic content as interpreted syntax will presumably have to explain why agents reliably adopt the same kind of attitude to differing semantic contents. Similarly, semanticists need to explain our linguistic intuitions about the Fregean puzzles if they endorse the view that names are rigid designators or constant individual concepts. Frequently these types of explanations will invoke pragmatic notions.

*Forthcoming in *Synthese*

The point we wish to emphasize is that linguistic intuitions need to be carefully examined and that not all these intuitions, i.e., all linguistic data needs to be reflected in our semantic theory. This leads to the problem of *overfitting* and *underfitting* of our semantics to the linguistic data. The idea of overfitting, as recently popularized by Williamson (2020), and applied to semantics amounts to complicating our semantic theory to fit all (shared) linguistic intuitions of competent speakers (cf. Williamson, 2020, p.265). The underlying worry is that by accommodating all data the semantics has lost its explanatory power: we have added further parameters to our semantics not for principled reasons but only so we can accommodate the data. In such a case we are *describing* the data rather than *explaining* it. The target of Williamson's ire are hyperintensional semantics that adopt fine grained individuation of semantic content or, alternatively, forsake central tenets of formal semantics.

However, as much as there is overfitting there is underfitting, that is, adopting a semantic theory that fails to explain important pieces of data (cf., e.g., Berto, 2024). Hyperintensionalists may thus push back against Williamson and hold that his proposed semantics is unsatisfactory for reasons of underfitting, that is, because his favoured semantics does not satisfactorily explain linguistic data pertaining to, say, propositional attitude reports. It is not an easy task to adjudicate this matter but ultimately this will depend on the overall package provided by competing semantics and the strategies provided for explaining contravening linguistic data.

This paper focuses on a particular family of arguments that purport to show that truth-conditional semantics suffers from underfitting. By truth-conditional semantics we mean what Fine (2017b) calls *objectual* truth-conditional semantics and, more precisely, any form of inexact truthmaker semantics.¹ The arguments at stake are known as *Soames's argument* (cf. Soames, 1987, 2005). According to Soames his arguments show that semantic content cannot be identified with *sets of truth-supporting circumstances*, i.e., semantic content cannot be conceived as sets of possible worlds, situations, or states. This is a striking conclusion which flies into the face of much of contemporary work in natural language semantics and which requires careful examination.

The first aim of this paper is to present the argument and to critically discuss the various options proponents of truth-conditional semantics have to resist Soames's conclusion. We use this discussion to highlight the costs of some of these options and tentatively conclude that neither a pragmatic explanation nor an explanation in terms of heuristics seems promising. Rather Soames's arguments needs to be defused semantically. This takes us to the second aim of the paper, which is a discussion of a semantics for attitude contexts that—building on ideas from Stern (2021); Goodman and Lederman (2021)—combines possible world semantics with contextual ideas introduced by Crimmins and Perry (1989). On this semantics the attitudinal relation, that is, the accessibility relation will be context sensitive, which can be used to block the counterintuitive conclusion in Soames's arguments. Interestingly, while these ideas have been widely discussed, with the exception of Stern (2021); Goodman and Lederman (2021) they have not been investigated within the context of truth-conditional semantics, that is, possible

¹By truth-conditional semantics we thus understand semantics in which semantic content is explained in terms of truth and truth conditions. In contrast, semantics that stipulate a primitive notion of meaning (semantic content/proposition) such as structured proposition based semantics (see, e.g., King, 2007) and derive the truth conditions of a sentence from its meaning are not classified as truth-conditional semantics.

world semantics (PWS).² The second part of the paper is devoted to developing the formal semantics in some detail and to showing how it blocks Soames's argument. We end the paper by discussing the prospects and limitations of the semantics, and by pointing to future work.

2 Soames's Argument

The label *Soames's argument* is usually applied to a series of arguments which, according to Soames, show that semantic content cannot be sets of truth-supporting circumstances. The structure of Soames's arguments is to show that granted some basic assumptions the identification of semantic content with sets of truth-supporting circumstances leads to unpalatable consequences, that is, consequences that should ultimately lead us to reject the identification of semantic content with sets of truth-supporting circumstances. These basic assumptions are:

Compositionality The semantic content of a sentence is a function from the semantic content of its parts.

DR/CIC-view of names The referent of a name is the same relative to all truth-supporting circumstances.³

Relational analysis Propositional attitude reports express relations between agents and semantic contents.

In addition to these basic assumptions Soames requires at least one more specific assumption regarding the "logic" of the propositional attitude under consideration, that is, for his argument Soames requires that the attitude distributes over conjunction (DoC):

DoC If an individual *i* satisfies 'x V's that P & Q' relative to some context *c*, then *i* satisfies 'x V's that P' and 'x V's that Q' relative to *c*.

DoC allows us to conclude that if, say, 'Maya believes that snow is white and grass is green' is true relative to some context *c*, then "Maya believes that snow is white" is true relative to *c* and so is 'Maya believes that grass is green'.

Soames argues that if these assumptions are granted, we arrive at an unacceptable conclusion. Then either one of the assumptions must be given up or we need to accept that semantic content cannot be identified with sets of truth-supporting circumstances. According to Soames it is the latter. Truth-conditional semantics does not suffice for determining the semantic content of a sentence.

It is time to look at the actual argument. Following Soames (1987) we focus on the attitude verb 'believe' (cf. Soames, 1987, p. 42):

1. The ancients believed that 'Hesperus' refers to Hesperus and 'Phosphorus' refers to Phosphorus.

²Here, we understand possible world semantics in a wide sense and to include versions of situation semantics, but also semantics in which worlds are allowed to be paracomplete and/or paraconsistent.

³This means that a name refers directly to its referent or that the associated individual concept is a constant function that yields the same referent for ever possible world. See, e.g., Abbott (2010) for further discussion.

2. The ancients believed that ‘Hesperus’ refers to Hesperus and ‘Phosphorus’ refers to Hesperus.
3. The ancients believed that ‘Hesperus’ refers to Hesperus and ‘Phosphorus’ refers to Hesperus and for some x , ‘Hesperus’ refers to x and ‘Phosphorus’ refers to x .
4. The ancients believed that for some x , ‘Hesperus’ refers to x and ‘Phosphorus’ refers to x .

The idea is that 1 is uncontroversially true, while 4 is false. Yet, if we accept **Compositionality**, **DR/CIC-view**, **Relational analysis** together with **DoC**, then we are forced to conclude that 4 is true. This seems unacceptable and thus by modus tollens some assumption needs to be given up. According to Soames it should be truth-conditional semantics.

Let us go through the reasoning step by step. As mentioned, 1 seems uncontroversially true. Surely, the ancients believed that ‘Hesperus’ referred to the object they had named ‘Hesperus’, i.e., Hesperus and similarly for ‘Phosphorus’. But then 1 implies 2 by the **DR/CIC-view**. Admittedly, many will already find 2 hard to accept. However, Soames thinks that this intuition needs to be explained pragmatically and that the truth of 2 is just a consequence of direct reference, which is, according to Soames, a basic semantic fact. The inference from 3 to 4 follows from **DoC** and Conjunction Elimination. What about 2 to 3?

To illustrate the reasoning let ‘ $\varphi(h)$ ’ stand for the sentence ‘Hesperus’ refers to Hesperus and ‘Phosphorus’ refers to Hesperus’ where ‘ h ’ is short for ‘Hesperus’.⁴ Furthermore, let $\|\varphi(h)\|$ stand for the set of truth-supporting circumstances in which ‘ $\varphi(h)$ ’ is true, ‘bel’ for the belief-relation and ‘ a ’ for ‘the ancients’.⁵ Then 2 can be semantically represented as

$$(\dagger) \quad a \text{ bel } \|\varphi(h)\|.$$

Now in all truth-supporting circumstances in which $\varphi(h)$ is true there exists an object o , e.g., the referent of ‘ h ’ such that o satisfies $\varphi(x)$. Then given the standard truth conditions for the existential quantifier we obtain

$$\|\varphi(h)\| \subseteq \|\exists x \varphi\|,$$

and thus

$$\|\varphi(h)\| = \|\varphi(h)\| \cap \|\exists x \varphi\|.⁶$$

According to orthodox versions of intensional semantics (see, e.g., Gamut, 1991; Heim and Kratzer, 1998), but also situation semantics (Kratzer, 2023; Barwise and Perry, 1983), that is, inexact truthmaker semantics the semantic content of a conjunction is simply the intersection of the contents of its conjuncts, i.e.,

$$(\star) \quad \|\varphi(h)\| = \|\varphi(h)\| \cap \|\exists x \varphi\| = \|\varphi(h) \wedge \exists x \varphi\|.$$

⁴Accordingly, we take ‘ $\varphi(p)$ ’ to be the embedded sentence of the attitude report 1.

⁵We ignore complication that may arise because of the plural nature of ‘the ancients’.

⁶Notice that up to this point the reasoning holds in both *exact* and *inexact* truthmaker semantics. For a discussion of the truth conditions of quantifiers in exact and inexact truthmaker semantics see, e.g., Fine (2017b) and Picenni (2023).

In combination with (†) this yields

$$a \text{ bel } \|\varphi(h) \wedge \exists x\varphi\|,$$

that is, Line 3 of the argument.

Our above reconstruction shows that the argument will apply to all forms of possible worlds semantics or inexact truthmaker semantics that meet Soames’s basic assumptions *and* that for sentences φ_1, φ_2 and any closed term t :

$$\begin{aligned} (\wedge) \quad & \|\varphi_1 \wedge \varphi_2\| = \|\varphi_1\| \cap \|\varphi_2\| \\ (\exists) \quad & \|\varphi_1(t)\| \subseteq \|\exists x\varphi_1(x/t)\|. \end{aligned}$$

If these conditions are met, the argument will apply independently of whether “worlds”, that is, truth-supporting circumstances are classical, paraconsistent or paracomplete.⁷

2.1 Soames argument in exact truthmaker semantics

Conditions (\wedge) and (\exists) are widely accepted by theorists who conceive of contents as sets of truth-supporting circumstance with one notable exception: in exact truthmaker semantics (\wedge) does not hold (cf. Fine, 2017b, for an overview and discussion of exact truthmaker semantics). As a consequence, (\star) does not hold in exact truthmakers semantics without special assumptions and the reasoning from 2 to 3 in the argument will not go through. In exact truthmaker semantics a truthmaker is required to be “wholly relevant” for the truth of the sentence. From an information-theoretic perspective this means that all of the information provided by the truthmaker (infor) guarantees the truth of the sentence without providing excess information.

In truthmaker semantics the truth conditions for a conjunction, say, $p \wedge q$ is given by:

$$s \Vdash p \wedge q \text{ iff } \exists t, u (s = t \sqcup u \ \& \ t \Vdash p \ \& \ u \Vdash q),$$

where s, t and u are truth-supporting circumstances and $t \sqcup u$ forms the fusion of t and u . These truth conditions only justify interpreting conjunction as set-theoretic intersection if the following monotonicity principle is assumed: let ψ be an arbitrary sentence, \leq partial ordering relation on the set of truth-supporting circumstances (truthmakers), and s, t truthmakers. Then

$$(\text{MON}) \quad \text{if } s \leq t \ \& \ s \Vdash \psi, \text{ then } t \Vdash \psi.^8$$

However, this principle is precisely what sets exact and inexact truthmaker semantics apart. It does not hold for exact truthmaker semantics as t might not be ‘wholly relevant’ for the truth of ψ . Conjunction cannot be understood as set-theoretic intersection if truth-conditional

⁷Sometimes semantics that allow for paraconsistent “worlds” or, more generally, for “worlds” that are deemed metaphysically impossible are called ‘impossible world semantics’. In this paper we reserve the label ‘impossible world semantics’ for semantics that allow for worlds, which, by design, are not allowed to contribute to the semantic content of any sentence, that is, worlds that are not truth-supporting circumstances.

⁸Assume (Mon), then every truthmaker of $p \wedge q$ is a truthmaker of p and q respectively and so $\|p \wedge q\| \subseteq (\|p\| \cap \|q\|)$. For the converse assume $s \in \|p\| \cap \|q\|$. Then $s = s \cap s \in \|p \wedge q\|$ by the truth conditions of exact truthmaker semantics.

semantics is understood in terms of exact truthmaker semantics. Soames’s argument does not apply in this case.⁹

The observation that Soames’s argument does not apply to exact truthmaker semantics will not come as a surprise to many, since, at least if the consequence relation is assumed to be that of exact entailment the exact truthmaker theorist will not accept **DoC**. In the logic of exact entailment conjunction elimination is not a valid rule of inference (cf., e.g., Fine and Jago, 2019) and, arguably, one should then not expect an agent’s beliefs to be closed under this rule. However, the key message is that the reason that Soames’s argument does not apply to exact truthmaker semantics is independent of whether **DoC** is endorsed. Rather it is due to the fact that (\wedge) does not hold for exact truthmaking: Soames argument does not apply even if the exact truthmaker semanticist endorses **DoC**.

2.2 A Variation on Soames

From the perspective of the exact truthmaker theorist Soames’s argument falls short of showing that semantic content cannot be sets of truth-supporting circumstances. Nonetheless we argue that Soames-style arguments have far reaching consequences even if exact truthmaker semantics is assumed: using a variation on Soames’s argument one can show that exact truthmaker semantics cannot be squared with disquotationalist views of truth according to which a sentence φ is true relative to a truth-supporting circumstance iff the sentence ‘ Tt_φ ’ is true relative to that truth-supporting circumstance whenever t_φ is a name of φ , and ‘ T ’ is the truth predicate of the language (see Stern, 2021, for further discussion). On this widely accepted view φ and Tt_φ express the same semantic content, as they are true at exactly the same truth-supporting circumstances. It turns out that we can run a Soames-style argument in exact truthmaker semantics and PWS more generally without relying on (\wedge) and (\exists) . To this effect we replace **DoC** by

Disquotation Let φ be a sentence and t_φ a name of the sentence. Then $\|\varphi\| = \|Tt_\varphi\|$, i.e., φ and Tt_φ express the same semantic content.

If **Disquotation** is assumed

- (a) Every even number greater than two is the sum of two primes.
- (b) Goldbach’s conjecture is true.

express the same semantic content and, assuming **Compositionality** and **Relational analysis**, we can conclude that the following two sentences

- (c) Max believes that Goldbach’s conjecture is true.
- (d) Max believes that every even number greater than two is the sum of two prime numbers.

⁹Relatedly, exact truthmaker semantics is irreducibly bilateral, that is, in contrast to inexact truthmaker semantics a sentence’s falsemaker cannot be retrieved/determined from the sentence’s truthmaker. Ultimately, this yields a further argument against (\star) as in exact truthmaking semantics content needs to be identified with the ordered pair consisting of a set of truthmakers and a set of falsemakers.

express the same semantic content. Yet, this seems to be in plain contradiction with the following scenario:

Max believes that Goldbach’s conjecture is true. His friend Philip told him so and Philip is a mathematical genius. Max has absolute faith in Philip and believes him even though he has no idea what Goldbach’s conjecture asserts. In fact, he does not believe that every even number greater than two is the sum of two prime numbers.

It seems that relative to this scenario (c) is true but (d) is not. It simply seems wrong to ascribe to Max the belief that every even number greater than two is the sum of two prime numbers.¹⁰

Prima facie the example shows that even exact truthmaker theorists should not remain unfazed by Soames-types argument. Perhaps, disquotationalism about truth is wrong, but should that be a consequence of adopting exact truthmaker semantics or of conceiving of semantic content as a set of truth-supporting circumstances?¹¹

On a more general note our variation on Soames’s argument **also** suggests that attempts of blocking Soames’s argument that merely give up **DoC** may fall short of blocking Soames-style arguments. As discussed, the argument establishing the equivalence of (c) and (d) does not rely on **DoC**. Granted that rejecting **DoC** is not a viable option, what option does a truth-conditional semanticist have to resist Soames’s conclusion that the semantic content of a sentence cannot be identified with a set of truth-supporting circumstances?

3 Resisting Soames’s Argument

According to Soames his argument shows that semantic contents cannot be sets of truth-supporting circumstances, that is, semantic content cannot be obtained from truth conditions. How can Soames’s argument and conclusion be resisted. For one, one might seek to expose Soames’s push for a more fine grained individuation of content as an instance of overfitting and explain our intuitions by non-semantic means, that is, pragmatically or, following Williamson (2020), by means of heuristics. For another, one may of course give up one or more of the assumptions Soames employs in his argumentation, that is, one could give up **Compositionality**, **Relational Analysis**, or **DR/CIC-view**. Indeed, all of these assumptions have been challenged by various theorists and, before turning to the idea of overfitting, we shall discuss the main semantic strategies employed in undermining Soames’s key assumption.¹²

¹⁰Stern (2021) provides further examples to the effect that *believing* and *believing true* may fall apart. For example, by **Disquotation** we may also infer that the two sentences

- Lois believes Superman is strong.
- Lois believes the sentence ‘Clark Kent is strong’ is true.

express the same semantic content.

¹¹Georgi (2019) proposes a related argument purporting to show that there is a clash of **disquotation** and treating propositions as sets of truth-supporting circumstances.

¹²Our discussion should not be understood as an attempt to provide knockout arguments, but rather as highlighting consequences of the proposals we find unappealing.

Contra the DR/CIC-view Elbourne (2010) challenges the view that names but also demonstratives and pronouns are constant individual concepts, and, consequently, also the idea that they are directly referential.¹³ According to Elbourne “*there is a growing consensus among researchers on natural language semantics that pronouns and demonstratives are not (or at least not always) directly referential*” (Elbourne, 2010, p. 105) and that “*the presumption that names are directly referential is surely not one on which it is safe to base an argument*” (Elbourne, 2010, p. 108). If the **DR/CIC-view** is given up, then the inference from 1 to 2 does not go through. Soames’s argument is blocked.

It worth noting that semanticists that are willing to reject the **DR/CIC-view** are not really the target of Soames argument, as these theorists will find 2 unpalatable and there is then no need to continue the reasoning.¹⁴ Rather the argument is directed at staunch proponents of truth-conditional semantics who commit to the **DR/CIC-view**.

Yet, even theorists that are willing to reject the **DR/CIC-view** might need to worry about the variation on Soames, i.e., the Max-example. Is it possible to block the inference from (c) to (d) by rejecting the **DR/CIC-view**? Presumably, the idea would be that ‘Goldbach’s conjecture’ does not rigidly designate the proposition that every even number greater than two is the sum of two prime numbers. However, presumably (c) and (d) still turn out equivalent on the *de re*-reading of (c) and, more general, to make sense of truth ascriptions in belief contexts we, arguably, need to view proposition and/or sentence denoting terms directly referential. So it seems that giving up **DR/CIC-view** does not yield the correct outcome nor is it appropriate when handling the argument from **disquotation**.¹⁵

Contra Compositionality I: Structured Intensions A common strategy for handling the problems attitude contexts pose for intensional semantics is to implicitly or explicitly give up on, or at least substantially weaken, the compositionality requirement, at least for what attitude contexts are concerned: *the semantic content of a propositional attitude report is not determined (merely) by the semantic contents of its parts*. Arguably, the idea traces back to Carnap’s (1947) notion of *intensional isomorphism*. Abstracting away from some specifics of Carnap’s proposal the idea is that propositional attitude reports should not be understood as relations between agents and semantic contents, but as relating to finer individuated entities, entities which are sensitive to the structure of the sentence (*intensional structures*). In Lewis (1970) these entities are identified as *meanings* and Cresswell (1975) more or less explicitly proposes a semantics for attitude reports in which propositional attitudes are identified as relations between agents and such *structured meanings*. If we identify meanings with propositions, then, according to this view, semantic contents and propositions fall apart. The latter but not the former are the objects of attitudes, but the former and not the latter are sets of truth-supporting circumstances.

One problem with such proposals is that they ultimately require a theory of structured intensions and developing a consistent theory may prove more complicated than one may

¹³In his paper, Soames (1987) also gives arguments that do not rely on the use of names, but variables. This is dubious since, as Elbourne (2010) points out, it is difficult to judge the natural language intuitions of speakers for sentences that incorporate vocabulary from artificial languages. By arguing that demonstratives and pronouns are not directly referential Elbourne attempts to further undermine the idea that variables are directly referential.

¹⁴Also, on Soames preferred semantics 2 will come out true.

¹⁵See Stern (2021) for further discussion.

think. The issue is that the immediate idea of taking these structured intensions to be tuples (=structures, trees,...) of intensions, i.e., functions from expressions to their extensions runs into problems with respect to iterated attitudes. Presumably, the intension of an attitude verb should be a function from individual concepts and a structured intension to truth values. But, in the case of iterated attitudes this means that the intension of the attitude verb needs to occur in its argument, i.e., in the structured intension. This takes us into paradoxical territory and, assuming standard set theory, the intension of the attitude verb may turn out to be ill-defined (see, e.g., Cresswell, 1975; Kratzer, 2022, for discussion). Cresswell (1975, 1985) proposes some workarounds but ultimately this leads to a form of ad hoc'ness which is in tension with the idea of systematic, compositional semantics. Moreover, if the theory of structured intension can no longer be extracted from truth-conditional semantics in a straightforward way, one may wonder whether this does not invert the explanatory order, i.e., are we still working within a framework in which meanings are obtained from truth conditions? If not, what is the difference between a theory of structured intensions and a theory of structured propositions? Independently of how we answer this question, the take home message is that embracing structured intensions violates a strong version of compositionality, as the semantic content of an attitude report is no longer the output of the semantic contents of its part (via some form of functional application). Therefore, we take it that adopting such a semantics for propositional attitudes amounts to an implicit rejection of compositionality and makes the semantics vulnerable to arguments from overfitting.¹⁶

Returning to Soames's argument the effect of taking structured intensions as objects of the attitudinal relation will have the consequence that the step from 2 to 3 will no longer be valid: while the semantic content of sentences embedded in the clausal complement of 'believe' in 2 is just the semantic content of the sentence embedded in clausal complement of 'believe' in 3, the two sentences have different structures which will lead to different structured intensions/meanings. The inference is no longer valid. Similarly, since the two sentences (a) and (b) have different syntactic structures, they will have different structured intension, which allows us to block the semantic equivalence of (c) and (d).

Contra Compositionality II: Impossible World Semantics Structured intension theorists hold that attitudinal objects are not semantic contents, that is, sets of truth-supporting circumstances but structured entities. Similarly, the impossible world theorist holds that attitudinal objects are not semantic contents conceived of as sets of possible worlds, but sets that may also contain impossible worlds (or potentially tuples consisting of a set of possible worlds and a set of, potentially, impossible worlds). Impossible worlds are worlds that are not truth-supporting circumstances and are usually deemed to be metaphysically and, perhaps, logically, impossible. For example, impossible worlds could be epistemic alternatives (cf. Hintikka, 1962) in which, contra the **DR/CIC-view**, 'Hesperus' and 'Phosphorus' may have different referents at different epistemic alternatives. This would block the step from 1 to 2 in the argument. There

¹⁶Admittedly, this assessment is based on what Yalcin (2014) might call the equation of semantic value and semantic content, and theorists who resist this equation on principled grounds will be unfazed by violations of compositionality of the kind we discussed. At the very least, it amounts to a rejection of *semantic innocence*, that is, the view that expressions should have the same reference and content in all environments (Davidson, 1968; Barwise and Perry, 1983).

are more drastic versions of impossible world semantics in which impossible worlds are not semantically controlled, that is, an impossible world is ultimately just some set of sentences—the sentences that are true in this possible world.¹⁷ For example, at such an impossible world $\varphi(h)$ can be true and $\varphi(h) \wedge \exists x\varphi$ false despite the fact that both sentences express the same semantic content conceived of as set of possible worlds. This would block the step from 2 to 3 in the argument, but ultimately every step in the argument can be blocked by appealing to impossible world semantics. In particular, impossible world semantics can also be used to resist the equivalence of (c) and (d) in the Max-example.

Whilst Impossible World Semantics is arguably a systematic semantics it is not a compositional semantics: the semantic content of a complex expression does not only depend on the semantic contents of its parts, since the content of the belief report depends on impossible worlds, that is, worlds that are not elements of the semantic content of the embedded clause of the attitudinal verb. Notice that if, as a remedy, we allow impossible worlds to figure as elements of the semantic content of a sentence, then $\varphi(h)$ and $\varphi(h) \wedge \exists x\varphi$ will no longer express the same semantic content (as per the above example) in violation of clause (\wedge) and (\exists). To vindicate the compositionality of the semantics alternative clauses to (\wedge) and (\exists) need to be proposed. To our knowledge no convincing clauses have been proposed in this context (*pace* exact truthmaking discussed in Section 2.1).¹⁸ Moreover, in the argument from disquotation the only options available to the impossible world theorists are to reject compositionality or to hold that $\|Tt_\varphi\| \neq \|\varphi\|$, which amounts to rejecting **disquotation**.

Similarly, if we were to include doxastic alternatives in which ‘Hesperus’ and ‘Phosphorus’ may have different referents in the semantic content of $\varphi(h)$ and $\varphi(p)$, we end up explicitly rejecting the **DR/CIC-view**, as $\varphi(h)$ and $\varphi(p)$ will have different semantic contents.¹⁹ In sum we take it that if Impossible World Semantics is deemed attractive, it is as a non-compositional semantics.²⁰

While, at least for some, a non-compositional semantics is unacceptable, we may ask whether the flexibility of the semantics justifies its non-compositionality, and whether introducing impossible worlds is a promising strategy for resisting Soames’s argument. As we see it, Impossible World Semantics that places no restriction on the notion of an impossible world, that is, worlds are conceived of arbitrary collections of sentences should be met with a some scepticism. Having points of evaluation in which ‘everything is possible’ seems to go counter the idea of a principled semantics. It seems to be precisely the kind of manoeuvre Williamson

¹⁷See, e.g., Fagin et al. (1995) for a presentation of impossible world semantics. Importantly, on our understanding of ‘impossible world’ paraconsistency is not sufficient for a world to be impossible (cf. Footnote 7). For example, a form of possible world semantics that requires all worlds to be closed under the logic FDE rather than classical logic is not an impossible world semantics in our terminology. It would also not help with Soames’s argument.

¹⁸Berto and Jago (2009, Ch. 8.5) attempt to recover compositionality within impossible world semantics. However, without further detail on their “worldmaking” language and what kind of (algebraic) operations the usual connectives in this worldmaking language correspond to, their proposal amounts to interpreting the logical operations in terms of their corresponding syntactic operations. To our mind this can hardly be considered satisfactory.

¹⁹The point is not that rejecting **DR/CIC-view** is per se problematic, but that adding ‘impossible’ worlds to the semantic content of a sentence turns those worlds into possible worlds, that is, the idea of distinguishing between possible and impossible worlds gets lost. We are no longer working with impossible world semantics but with possible world semantics for which the **DR/CIC-view** has been given up.

²⁰As per above, this again assuming the equation of semantic value and semantic content.

has in mind when he argues against overfitting: we have aligned the semantics with some important linguistic data, yet we have done so in an fairly ad hoc manner, that is, by adding impossible worlds as parameters to our semantic equations. Admittedly, the issue becomes more subtle when considering ‘impossible’ epistemic or doxastic alternatives, i.e., worlds in which co-referential names may refer to different objects. We do not attempt to answer this question in this paper. It is worth recalling, however, that rejecting the **DR/CIC-view** or, equivalently, allowing for ‘impossible’ epistemic alternatives is arguably not sufficient for handling the Max-example.

Contra Relational Analysis A further option for resisting Soames’s argument is to deny that propositional attitude reports express relations between agents and propositions conceived of as sets of truth-supporting circumstances. However, not all alternatives fit well with the present discussion. For one, taking the attitudinal object to be something different than propositions will lead us straight back to the discussion of compositionality: the semantic content of the attitude report would no longer be a function from the contents of its parts. Arguably, this would affect sententialist proposals or proposal that take attitudinal objects to be interpreted logical forms (Larson and Ludlow, 1993). Another option that does not seem promising in response to Soames are probabilistic approaches, as one can recast the argument in probabilistic terms. We refer to Swanson (2011) for further details and discussion.²¹

A more promising proposal is due to Moltmann (see, e.g., Moltmann, 2020). On her view a belief report ‘*a* believes φ ’ should be understood in the following terms: *a has a belief—which is the attitudinal object—and the content of φ is a partial content of *a*’s belief* where the latter means that for every truthmaker *s* that makes *a*’s belief true (satisfies *a*’s belief) there is a truthmaker s' with $s' < s$ such that s' makes φ true and for every truthmaker *s* of φ there is a truthmaker s' such that $s < s'$ and s' satisfies *a*’s belief. On Moltmann’s view attitudinal objects are not propositions, that is, contents of sentences of the languages but objects such as beliefs, claims, etc. In this sense Moltmann’s view resembles ideas that featured in our discussion of impossible world semantics: in impossible world semantics we end up postulating distinct categories of V-alternatives, which then leads to different attitudinal objects, i.e., sets of V-alternatives. In Moltmann’s framework we end up with a plethora of different attitudinal objects, which are primitives of the semantics. It is not immediate how Soames’s argument could be carried out in this context, indeed, it seems unlikely that it can. However, to properly evaluate the view we would need to know whether and how the beliefs figuring in (c) and (d) differ. For another Moltmann’s proposal is situated within the framework of exact truthmaker semantics which already blocks Soames’s original argument. Although Moltmann’s proposal is very interesting and seems to be immune against the challenge of Soames’s argument, the flexibility of the semantics is, at least to some extent, achieved by populating the semantics by a plentitude of different attitudinal object and by application of the truthmaker relation to attitudinal objects. But how could the latter be recursively defined? And if this is not possible, does that not mean a great loss of explanatory power? One may therefore wonder whether Moltmann’s proposal

²¹Let us also point out that understanding propositional attitudes in terms of quantification over possible worlds will unsurprisingly not counter Soames’s argument either. As should become clear in Section 5.1, against the backdrop of possible world semantics the relational and the quantificational analysis are notational variants.

is an interesting way of *describing* the linguistic data (in an elegant and intriguing way), as opposed to explaining it.

While Moltmann’s approach seems to avoid Soames’s argument against the identification of semantic content with sets of truth-supporting circumstances it has several puzzling and controversial features that make it difficult to accept. Moreover, on our view there is also a more general reason why we deem Moltmann’s approach and, indeed, any approach that assumes the framework of exact truthmaking unsatisfactory from the perspective of truth-conditional semantics: exact truthmaking violates the monotonicity principle (MON) discussed in Section 2.1.²² Recall that (MON) is the following principle

$$\text{if } s \leq t \ \& \ s \Vdash \psi, \text{ then } t \Vdash \psi,$$

where ψ is an arbitrary sentence, \leq a partial ordering relation on the set of all truth-supporting circumstances (truthmakers), and s, t are truthmakers. If we understand the truthmakers s and t in terms of situations along the lines of Barwise and Perry (1983) or Kratzer (1986) $s \leq t$ means that the situation s is part of situation t , i.e., the situation t encodes all the semantic information provided by s . Exact truthmaker semanticists argue that a truthmaker s of, say, ψ should be “wholly relevant” to the truth of ψ . The idea is that if $s \leq t$, then t may provide us with a surplus of semantic information, which means t may no longer be wholly relevant to the truth of ψ and may disqualify t as a truthmaker of ψ . However, this idea sits ill with the fundamentals of truth-conditional semantics. If a situation s makes a sentence ψ true, why should a situation that encodes all the information provided by s not make ψ true? Is there some semantic information provided by t that stops ψ from being true? If so, it seems that our truth conditions are off; we should not have deemed ψ true in the first place: s is not a truthmaker of ψ . The idea of being “wholly relevant” may play an important role if we adopt, say, a psychological perspective: plausibly, we sometimes update our attitude towards a proposition in case we receive further salient semantic information. A non-monotonic semantics thus seems to be desirable from this perspective. But truth conditions are not attitudes, and we should be careful in keeping them apart. On our view the monotonicity principle (MON) and, more generally, the monotonicity of the semantics is a cornerstone of truth-conditional semantics: more semantic information leads to more sentences being made true and false respectively. In sum, exact truthmaker semantics might be an interesting semantics, but it is not a truth-conditional semantics as we understand it.²³

3.1 Overfitting: non-semantic explanations

The conclusion of Soames’s argument seems wrong, that is, Sentence 4 does not seem true but false. However, this does not show that it *is* false. Maybe there are pragmatic or other effects at work that just give us the impression that the sentence is false, whilst in actual fact it is true. So perhaps we should ultimately simply accept the conclusion of Soames’s argument,

²²To be clear, exact truthmaking violates (MON) by design and proponents of exact truthmaking such as Fine (2017a) take this to be a major advantage of exact truthmaking over inexact truthmaking.

²³Admittedly, our take on exact truthmaking may be controversial and the central role we ascribe to the monotonicity principle (MON) requires further discussion, yet such a discussion is left for another paper.

i.e., perhaps we need to simply accept that 4 is true and that any attempt to defuse Soames's argument semantically leads to overfitting.

Pragmatic Explanations For example, one could try to employ the two-dimensional model proposed by Stalnaker (Stalnaker, 1978) to explain why we deem the sentence false. On Stalnaker's view Sentence 2 seems false to us because, arguably, the proposition (semantic content) expressed by $\varphi(h)$ is different from the proposition the ancients took $\varphi(p)$ to express. By the **DR/CIC-view** we know that $\varphi(p)$ and $\varphi(h)$ express in fact the same proposition. However, the ancients thought of the names 'Hesperus' and 'Phosphorus' as referring to different objects and thus of $\varphi(p)$ and $\varphi(h)$ as expressing different propositions: this explains why we judge the belief reports 1 and 2 differently.

To this effect Stalnaker (1978) introduces the notion of the *diagonal proposition*. Roughly put, the diagonal proposition of $\varphi(p)$ is the proposition that $\varphi(p)$ would express if the world were such as the ancients imagined it to be, i.e., if 'Hesperus' and 'Phosphorus' were to refer to different celestial objects. Now, while $\varphi(p)$ and $\varphi(h)$ express the same proposition (=semantic content), they have different diagonal propositions. The differing diagonal propositions explain why we judge 1 true, but 2 false—it is not that 2 is in fact false. $\varphi(p)$ and $\varphi(h)$ coincide semantically, but because of their differing diagonal proposition they have different pragmatic implicatures and this is the reason why we judge 1 true and 2 false.²⁴

Of course, this is an overly simplistic presentation of Stalnaker's account and questions may be asked of whether the strategy is promising (cf. Soames, 2005, for a critical discussion). But let us grant that the model can be used for explaining the differing truth-value intuitions regarding 1 and 2. Presumably, the same type of explanation should then be available for explaining why we deem 4 unacceptable. At the outset, it is not obvious how the strategy can be put to work for it seems that the ancients understood $\exists x\varphi$ the way we do. If it is true that we understand $\exists x\varphi$ in the same way as the ancients, then the proposition and the diagonal proposition coincide and it is unclear how to employ Stalnaker's account. It seems that we ought to take the intuition that 4 is false seriously. However, following Soames's reasoning $\varphi(h) \wedge \exists x\varphi$ implies $\exists x\varphi$ and perhaps one might then try to piggyback on the explanation why we deem 2 and, as consequence 3, false: perhaps we can explain our intuition that 4 is false by the fact that 4 has been derived from premises that we deemed false.²⁵ The problem with this strategy is that it also works in the opposite way: if the conclusion seems false to us and we cannot provide a convincing pragmatic explanation that accounts for this intuition, doesn't that mean that either the premise of the argument was wrong or one of the assumptions we used in the derivation was mistaken? Soames would argue that the step from 2 to 3 was mistaken since semantic content should not be conceived as truth-supporting circumstance, but alternatively one could also argue that, e.g., 2 is indeed false and one of Soames's assumption is not quite correct. Our argumentation by no means shows that a pragmatic explanation for our intuition regarding 4 is unattainable, but it does suggest that coming up with a good explanation is not straightforward. Moreover, given that the intuition that 4 is false seems rather robust, we think

²⁴See Stalnaker (1987) for an argument along these lines.

²⁵This reasoning is of course logically fallacious because logical reasoning is usually not falsity preserving but truth preserving.

that one should not yet give up on looking for a good semantic explanation.

Williamson on heuristics Recently, Williamson (2020) has argued that that our intuition that 4 is false should not be explained pragmatically but in terms of *heuristics*. According to Williamson heuristics are “*fast and frugal*’ or, in less positive terms, ‘*quick and dirty*’ methods that tend to give the right answers under normal conditions, but the wrong one in unfavourable cases” (Williamson, 2020, p. 4). According to Williamson the role of heuristics is to “to explain how speakers recognize the truth value of the proposition in practice” (Williamson, 2020, p. 5). Applied to propositional attitudes and, in particular, belief reports the idea is that a typical heuristic for assessing whether a belief report is true, is to query whether the agent believes the proposition under the *guise* of the sentence in the attitudinal clause. For example,

- (e) The ancients believed that Hesperus is Hesperus.
- (f) The ancients believed that Hesperus is Phosphorus.
- (g) The ancients believed that Hesperus is Hesperus under the guise of the sentence ‘Hesperus is Hesperus’.
- (h) The ancients believed that Hesperus is Hesperus under the guise of the sentence ‘Hesperus is Phosphorus’.
- (i) The ancients believed that Hesperus is Phosphorus under the guise of the sentence ‘Hesperus is Hesperus’.
- (j) The ancients believed that Hesperus is Phosphorus under the guise of the sentence ‘Hesperus is Phosphorus’.

according to Williamson (e) and (f) express the same semantic content and are both true. However, the heuristics for determining whether the two sentences are true will arguably differ in the two cases: whereas the standard heuristic for (e) will lead us to query whether (g) is true, in the case of (f) we tend to use (j) instead of (i). Of course (j) is false: the ancients did not believe Hesperus is Phosphorus under the guise of the sentence ‘Hesperus is Phosphorus’. We find ourselves in an unfavourable case in which our heuristic for determining whether a belief report is true leads us astray. Following this line of thought, Williamson would presumably argue that when it comes to Soames’s argument we find ourselves in such an unfavourable case and that the naive heuristics for determining the truth value of belief reports breaks down for Sentence 2 to 4. In particular, the ancients do not believe that something is both the referent of ‘Hesperus’ and of ‘Phosphorus’ under the guise of the sentence ‘something is both the referent of ‘Hesperus’ and of ‘Phosphorus’’. Presumably, according to Williamson it is simply a semantic fact that they hold that belief and it might not be possible to give a ‘quick and dirty’ explanation of this fact. The truth of 4 is simply a consequence of our semantics and, ultimately, needs to be judged against the background of the theoretical framework that is offered in way of understanding language and how it is used by speakers. The fact that we intuit 4 false needs to be evaluated against the bigger theoretical picture, that is, against the background of the entire theoretical package that is employed to explaining language and how it is used.

It's important to stress that Williamson's story about heuristics and believing under a guise is not a semantic story and that he would presumably accept all of Soames's underlying assumptions. Yet, on his account the fact that 4 is true does not create a problem: it is simply a case in which our 'quick and dirty' heuristics went wrong. This sounds great but perhaps a bit too great. Indeed, Williamson's use of 'heuristics' has the air of a 'get out jail free'-card that can be used whenever semantic and pragmatic explanations fall short. In particular, if a semantics systematically falls short of explaining a certain type of data, one should arguably investigate whether this data can be accommodated by modifying the semantics. Of course, as we discussed in the Introduction there is a balance to be struck between over- and underfitting of the semantics, but we don't think all avenues have been sufficiently explored when it comes to the semantics of propositional attitude reports. In particular, we think that the idea of guises under which we believe, if somewhat generalized, can be fruitfully integrated into our semantics.

4 Reconsidering the Relational Analysis

On our view there seem to be two take home messages of our discussion of Soames argument so far:

- (I) if **Disquotation** is accepted, rejecting the **DR/CIC**-view is not sufficient;
- (II) if **DR/CIC**-view is accepted, then resisting Soames argument seems to come at the cost of distinguishing between attitudinal objects and semantic content, and thus a violation of compositionality (at least on an orthodox understanding of the latter).

Is there a way to resist Soames argument for theorists who want to avoid committing to this distinction and who feel that acknowledging such a distinction ultimately amounts to accepting Soames's argument against truth-conditional content? That is, can we find truth conditions for attitude reports that are monotone, identify attitudinal objects with semantic content, and are compatible with

- compositionality;
- the DR/CIC-view;
- the relational analysis;
- DoC, and
- disquotation?

Our discussion up to this point suggests otherwise, but so far we have not considered the option of introducing further parameters to the logical form of attitude reports. Focussing on belief reports the idea would be that the logical form of a belief report

a believes that φ

is not exhausted by

$$a \text{ bel } \|\varphi\|.$$

The idea is that whilst conceiving of the attitudinal relation as a relation between an agent and a proposition is on the right track, further analysis suggests that we might need to slightly alter this idea. For example, the attitudinal relation could depend on a further parameter that is contextually supplied (cf. Crimmins and Perry, 1989).²⁶ There is another popular option that stems from the analysis of so-called de re belief reports originating from the works of Quine (1956); Kaplan (1968); Lewis (1979) and which has been proposed in different versions by Cresswell and Von Stechow (1982); Percus and Sauerland (2003); Charlow and Sharvit (2014) and many others. Glossing over some details the belief report

(★) Ralph believes Ortcutt is spy.

is true at a world w iff

- there is some salient acquaintance relation A which relates Ralph with the individual Ortcutt at w , i.e., $A(\text{Ralph}, \text{Ortcutt}, w)$, and
- if $A(\text{Ralph}, d, w')$, then d is a spy at w' for all $d \in D$ and all of Ralph's doxastic alternatives w' .²⁷

Strictly speaking, Ralph's acquaintance relations are functions that for every of his doxastic alternatives deliver the unique object that Ralph is acquainted with given the specific acquaintance relation.

If one applies the analysis of belief reports to Soames's argument, then the argument will break down at step 1 to 2. Presumably, the premise of Soames's argument, i.e. 1, is true since

- the ancients are acquainted with Venus via (at least) two acquaintance relations: *the brightest star on the evening firmament* and *the brightest star on the morning firmament*. Let's call the former A_1 and the latter A_2 . Then
- if $A_1(a, d', w')$ and $A_2(a, d'', w')$, then 'Hesperus' refers to d' is true at w' and 'Phosphorus' refers to d'' is true at w' for all $d', d'' \in D$ and all doxastic alternatives w' .

However, it is not the case that for A_i with $i \in \{1, 2\}$:

- if $A_i(a, d, w')$, then 'Hesperus' refers to d is true at w' and 'Phosphorus' refers to d is true at w' for all $d \in D$ and all doxastic alternatives w' .

So, the inference from 1 to 2 fails and Soames's argument breaks down.²⁸

²⁶Indeed, this idea will form the basis of the semantics we develop in the next section of the paper.

²⁷That is, acquaintance relations relate objects of the domain to the agent relative to a world.

²⁸At least on the account of Percus and Sauerland (2003) things might not be as simple as we make it seem, as they allow for every *occurrence* of a name to be associated with a different acquaintance relation, that is, using their terminology a different concept generator (see below). Accordingly, the two occurrences of 'Hesperus' in 2 could be associated with different concept generators and 2 would then follow from 1 on this reconstruction provided that both acquaintance relations (concept generators) that are salient /contextually provided for the context of 1 are also available in 2. Whether Soames's argument is successful on this reconstruction is a subtle issue. We will not pursue this further since the initial formal rendering of 2 seems to be the more intuitive and salient one. We wish to thank an anonymous referee for pointing out this subtlety to us.

From the linguistic perspective one may wonder how these truth conditions are derived from the belief report. Without going into too much details there are (at least) two popular proposals. The first is due to Cresswell and Von Stechow (1982) and stipulates that the *res*, i.e., the subject of the embedded clause, is moved outside the clause (leaving a trace) and becomes an argument of the attitudinal verb. An attitudinal verb such as belief thus no longer expresses a binary relation relating the agent with propositions, but a ternary relation which instead of propositions takes individuals and properties as arguments. Since the *res* appears explicitly as an argument of the attitudinal relation it is available for replacement of objects that are in the suitable acquaintance relation at the relevant world. The second proposal due to Percus and Sauerland (2003) allows for a replacement of the *res* in situ. Their idea is that the logical form of (de re) attitude reports contains unarticulated constituents (or *silent items* in their terminology), so-called concept generators, which are functions from individuals to individual concepts. This means that the context of utterance supplies an individual concept relative to the subject of the belief report, i.e., relative to the relevant agent.²⁹ In a nutshell, the concept generator provides us with the way under which the agent is acquainted with the *res* of the attitude report.³⁰ Given a suitable concept generator G for (\star), the report will be true at w iff $G(\text{Ralph})(w')$ is a spy at w' for all of Ralph's doxastic alternatives at w .

From a formal point of view this strategy amounts to making the embedded clause context sensitive: concept generators map worlds onto propositions relative to the subject of the attitude report. At each of Ralph's doxastic alternatives the sentence

Ortcutt is a spy

may be mapped onto a different proposition, that is, a different set of words. The proposition expressed by the embedded clause will depend on the guise the *res* is presented to the agent. The strategy thus seems to fit with Soames's desiderata but blocks his arguments. The caveat is that the semantic strategy is meant to apply to de re belief reports only: what happens in the case of de dicto reports, e.g., reports for which there is no salient *res*. Indeed, Soames might hold that his argument pertains to de dicto as opposed to de re beliefs.³¹

This said, if ways (and/or guises) under which specific objects are presented to us play a prominent role in the semantics of de re belief reports, why shouldn't we, contra Williamson, take guises under which propositions are presented to us play a prominent role in the semantics of de dicto belief reports. From a methodological point of view it seems worth exploring whether one can exploit guises in the semantics of de dicto belief reports. Yet, both strategies we just discussed are not immediately applicable to de dicto belief reports: without a designated *res* analysing the attitudinal relation as ternary relation between agents, the subject of the embedded clause, and a property just seems unjustified. Similarly, the silent item of the embedded clause was stipulated for de re- (and de se-) reports only and since in the case of de

²⁹See Aloni (2005) for a related proposal that includes a fully developed formal semantics.

³⁰As indicated in Footnote 28 this applies not only to a particular *res* but any DP occurring in the embedded clause that is constructed de re.

³¹The semantics for de re beliefs might be better placed to handle the Max-example if sentence (d) is construed as de re belief and where the t_φ denotes Goldbach's conjecture qua set of possible worlds. The acquaintance relation would then relate Max to propositions, i.e., sets of possible worlds and thus, depending on Max's doxastic alternatives, Tt_φ will be mapped onto different propositions.

dicto reports the embedded clause is typically viewed as a black box, i.e., as opaque and non-transparent, it is not obvious that we could use any silent items occurring within the embedded clause in the semantics of de dicto reports. To make room for guises in the logical form of (de dicto) attitude reports an immediate idea is then to stipulate that guises affect the attitudinal relation rather than the proposition expressed by the embedded clause. This means that the attitudinal relation and, in particular, the belief relation is not a binary relation between agents and propositions, but a ternary relation applying to agents, propositions, and guises (or *ways of believing* as we shall call them).³²

Such a view thus amounts to a variant of the contextualist view proposed by Crimmins and Perry (1989). Of course, Crimmins and Perry (1989) take propositions to be structured entities and not sets of truth-supporting circumstances, yet the idea can also be implemented in various forms of possible world semantics. On such a view, 4 come out as false. Indeed, Soames’s argument would already break down at the inference from 1 to 2. On our picture $\|\varphi(h)\| = \|\varphi(p)\|$ would still hold, but the proposition might only be believed under a way of believing associated with $\varphi(p)$ and not under the way of believing associated with $\varphi(h)$. Let x be the way of believing associated with 1 and y the way of believing associated with 2. Then, according to our picture, we would have

$$\text{bel}(a, \|\varphi(h)\|, x)$$

but

$$\text{not bel}(a, \|\varphi(h)\|, y)$$

and thus also $\text{not bel}(a, \|\varphi(p)\|, y)$, which blocks Soames’s argument. Employing a similar reasoning will also block our variation of Soames’s argument, that is, the Max-example.

In the remainder of the paper we seek to develop this view in some detail within a version of possible world semantics.³³ Similar to Percus and Sauerland (2003), and Crimmins and Perry (1989) we stipulate that the context contributes “unarticulated constituents” to the proposition expressed by the belief report, which can be roughly described as what we label ways of believing. However, at this point we wish to remain neutral as to whether these unarticulated constituents, i.e., ways of believing are syntactically realized or whether the context-sensitivity of the attitudinal relation is a consequence of the lexical entry of the attitude verb under consideration.³⁴ We shall assume the attitudinal relation and in particular the belief relation to be a ternary relation that holds of three-place tuples consisting of an agent, a way of believing, and a proposition qua semantic content. In doing so we build on work by Stern (2021) and Goodman and Lederman (2021). Indeed, our semantics generalizes the main ideas of Stern’s (2021) semantics to arbitrary belief reports, whereas the latter semantics was designed to handle truth-ascription in belief contexts and specifically to dissolve problems like the Max-example.

³²On our view, ways of believing are not syntactic structures or sentences but, rather reflecting the talk of acquaintance relation above, ways under which the agent is acquainted with a proposition or the agent’s specific mode of presentation of the propositions.

³³More precisely, we work in a Stalnaker-style two-dimensional semantics, but where the “worlds” are allowed to be neither complete nor consistent.

³⁴This means that, following Goodman and Lederman’s (2021) terminology, we wish to remain neutral between *hidden indexicalism* and *verbalism*.

Goodman and Lederman (2021) develop their contextualist “semantics” in a more abstract setting and do not explicitly commit to forms of possible world semantics. However, they do not require content to be structured and their paper can be considered as providing general guidelines for developing contextualist semantics for attitude reports within the possible world framework. We shall discuss their ideas in more detail when evaluating our semantics.

At this point our ambition is to develop a precise version of doxastic semantics within a possible world semantics (broadly) understood that embraces the idea of a context-sensitive attitudinal relation. We claim that this semantics meets all the assumption stipulated by Soames, but at the same time does not suffer from overfitting in the sense of Williamson (2020): the semantics is compositional and semantic content can be conceived as sets of possible worlds. On our view the introduction of the additional parameter is well-motivated and should not be used to dismiss the semantics. Our hope is that our semantics may find application as a natural language semantics for attitude reports, but for this purpose more work needs to be done. For one, we would need to provide more details on the nature of the alleged context sensitivity of the attitude verb and provide further details on how to conceive of ways of believing, but, for another, we would need to provide empirical data in support of our view rather than merely conceptual and methodological argumentation. For the time being we view our semantics as an exploration of the prospects and limitations of possible world semantics that allows for a context-sensitive attitudinal relation.

5 Semantics for Attitudes

Within possible world semantics the binary belief relation may be defined using a doxastic accessibility relation:

$$\text{bel}_w(a, \|\varphi\|) : \leftrightarrow \forall v(wR_a v \Rightarrow v \in \|\varphi\|).$$

This means that an agent a believes the proposition $\|\varphi\|$ at w iff φ is true in all worlds v that are doxastically accessible for the agent a from w . How ought we implement a ternary belief relation in this framework that applies to tuples consisting of agents, propositions, and ways of believing? The obvious choice is to parametrize the accessibility relation, which in fact amounts to introducing a function $\text{Dox}_a : \text{WB} \rightarrow W \times W$ where WB is a set of ways of believing and W the set of truth-supporting circumstances under consideration. So Dox_a is a function, which applied to a way of believing, yields a doxastic accessibility relation for the agent a . It is in this sense that ways of believing simply amount to additional parameters in the semantic interpretation of propositional attitude reports. Of course, the question arises whether we can say anything more specific about these ways of believing or what they depend on. In setting up the semantics we wish to remain as general and neutral as possible, although the explanatory power of the semantics will depend on adding further substance to the notion of a way of believing. For now we conceive of ways of believing as depending on:

- the agent;
- the embedded clause;

- the context;
- the model, i.e., the interpretation function.

Ways of believing are thus the output of a function that takes a context, a formula, an interpretation function and an assignment function as an argument.³⁵ This allows for a lot of flexibility and should allow us to accommodate different views and ideas about how to conceive of the context sensitivity of attitudinal relation. In particular, following Goodman and Lederman (2021) we could also take ways of believing as sets of sentences, i.e., as a set of guises available to us at the given context (cf., Section 6). Or, alternatively, they could be taken to encode the QUD or topic of the belief report.³⁶ Following Stalnaker (1978) we take contexts to be worlds or, more generally, truth-supporting circumstances. This leads to a two-dimensional modal semantics: formulas will be evaluated relative to a tuple consisting of two worlds one being the circumstance of evaluation and the other the context of utterance. We leave it open what these worlds are, that is, what the truth-supporting circumstances are: they could be paracomplete, complete, consistent, or paraconsistent. In our framework they are simply parameter in the semantic interpretation. However, our semantics will be a monotone semantics and truth-supporting circumstances should be conceived of as inexact truthmakers.³⁷

5.1 Formal Semantics

We now begin setting up our formal semantics for belief. We make the simplifying assumption that there is only one agent and as a consequence we can drop reference to the agent from our definitions. However, it is worth remembering that the belief operator, the doxastic accessibility relation as well as the function that outputs ways of believing are thought to be indexed by a particular agent.

First, we introduce the syntax of our formal language. Terms t_0, t_1, t_2, \dots of our languages \mathcal{L} are either variables x_0, x_1, x_2, \dots or individual constants c_0, c_1, c_2, \dots . \mathcal{L} has countable predicate constants P_0, P_1, P_2, \dots of arbitrary finite arity, the identity predicate $=$ and the propositional constant \perp . The primitive logical constants of \mathcal{L} are $\neg, \wedge, \text{B},$ and \forall . All other logical constants will be considered as defined by their usual definitions. Well-formed formulas of \mathcal{L} are specified using the Backhus-Naur form:

$$\varphi ::= s = t \mid P s_1, \dots, s_n \mid \perp \mid \neg \varphi \mid \varphi \wedge \varphi \mid \text{B} \varphi \mid \forall v \varphi$$

where s, t and s_1, \dots, s_n are terms of \mathcal{L} and v is a variable.

³⁵We have included assignment functions for handling formulas (as opposed to sentences) of the language.

³⁶This would provide an interesting connection to proposals by, e.g., Berto (2024); Yalcin (2018); Hoek (2022) and Mayr and Schmitt (2024).

³⁷In the framework we develop in Section 5.1 this amounts to the following claim: let I, J be two interpretation functions, F a frame, $w \in W$ and φ a sentence of the language. Then

$$\text{if } I \leq J \text{ and } (F, I), w \Vdash \varphi, \text{ then } (F, J), w \Vdash \varphi,$$

where

$$I \leq J \text{ iff } [I^+(w, P) \subseteq J^+(w, P) \text{ and } I^-(w, P) \subseteq J^-(w, P) \text{ for all } w \in W \text{ and } P \in \text{Pred}].$$

We refer to Section 5.1 for the relevant definitions.

As for standard possible world semantics the central notion of the semantics is that of a frame. However, we start by introducing the notion of a preframe.

Definition 1 (Preframe). *A preframe G is a tuple $\langle W, D, U \rangle$ such that $W \neq \emptyset$ is a set of circumstances, $U \neq \emptyset$ is the universe of the frame and $D : W \rightarrow \mathcal{P}(U)$ is a function from circumstances to their respective domain such that*

$$\bigcup_{w \in W} D(w) = U \neq \emptyset.$$

Our semantics is a so-called inner/outer domain semantics (see, e.g., Fitting and Mendelsohn, 1998). In this kind of semantics the domain of quantification is world relative, but the interpretation of a term may be picked from the universe of the frame, that is, an object that might not be in the domain of quantification of the world under consideration. Inner/outer domain semantics leads to a positive free logic, that is, a logic where all quantifier-free logical truths are preserved, but for which Universal Instantiation and Existential Generalization need to be restricted to terms which denote an object in the domain of quantification. Nothing hinges on our choice of inner/outer domain semantics: the idea is simply to allow for maximal flexibility.

Next we define the notion of an interpretation function for \mathcal{L} on a preframe G .

Definition 2 (Interpretation). *An interpretation I on preframe G is a function such that for all $w \in W$*

- $I(w, c) \in U$, if c is an individual constant, and for all $v \in W$:

$$I(w, c) = I(v, c);$$

- $I(w, P) = \langle I^+(w, P), I^-(w, P) \rangle \in \underbrace{\mathcal{P}(U \times \dots \times U)}_{n\text{-times}} \times \underbrace{\mathcal{P}(U \times \dots \times U)}_{n\text{-times}}$, if P is an n -place predicate constant.³⁸

The set of all interpretation over a preframe G is denoted by Int_G .

We assume that individual constants are interpreted rigidly, so our semantics embraces the **DR/CIC-view**. As the world parameter is irrelevant for the interpretation of individual constants, we omit it and simply write $I(c)$ instead of $I(w, c)$. Notice however that the referent of a constant may not be in the domain of that world, that is, it may not “exist” at that world. In contrast to classical possible world semantics the interpretation of predicates consist of an extension and an antiextension which may or may not exhaust the universe of the frame (nor the domain of the world).

Definition 3 (Variable assignment). *A function $\beta : \text{Var}_{\mathcal{L}} \rightarrow U$ is called a variable assignment. The assignment $\beta(x : d)$ is an assignment function such that for all $v \in \text{Var}_{\mathcal{L}}$:*

$$\beta(x : d)(v) := \begin{cases} d, & \text{if } v \doteq x \\ \beta(v), & \text{otherwise.} \end{cases}$$

³⁸Notice that at this stage we make no assumption on whether extension and antiextension of a predicate exhaust the universe, or whether they are allowed to overlap. Nothing we say will depend on that.

It's time to introduce the central notion of our doxastic semantics, that is, the notion of a belief frame.

Definition 4 (Belief Frame). *A belief frame F is a tuple $\langle G, \text{Dox}, \mathcal{B} \rangle$ such that*

- G is a preframe;
- \mathcal{B} a function that outputs ways of believing conceived of as ordinals relative to interpretation $I \in \text{Int}_G$ such that $\mathcal{B}_I : W \times \text{Frml}_{\mathcal{L}} \times {}^{\text{Var}U} \rightarrow \text{WB}_F$ where $\text{WB}_F \subseteq \text{ON}$ is the smallest initial segment of the ordinals such that

$$|\text{WB}_F| \geq |W \times \text{Frml}_{\mathcal{L}} \times {}^{\text{Var}U}|;^{39}$$

- $\text{Dox} : \text{WB}_F \rightarrow W \times W$ is a function that takes ordinals as inputs and outputs a doxastic accessibility relation in the orthodox sense.⁴⁰

The function \mathcal{B} specifies the way of believing relative to an interpretation. As mentioned, relative to an interpretation the way of believing is meant to depend on a context (a world), the linguistic (syntactic) information, and the assignment function, i.e., which objects are assumed to be the value of the respective variables. Arguably, one should impose constraints on \mathcal{B} . For example, if φ is a sentence, then the way of believing should not depend on the particular variable assignment and for all $w \in W$ and $\beta, \beta' \in {}^{\text{Var}U}$:

$$\mathcal{B}(w, \varphi, \beta) = \mathcal{B}(w, \varphi, \beta').$$

Further constraints on \mathcal{B} may be imposed and we pick up the discussion about particular constraints later in this paper (see, especially, Appendix A).⁴¹

It remains to specify how to evaluate terms and formulas of the language over a belief frame relative to an interpretation.

Definition 5 (Term Denotation). *Let $F = \langle G, \text{Dox}, \mathcal{B} \rangle$ be a belief frame and t be a term of \mathcal{L} . Then the denotation of t relative to an interpretation $I \in \text{Int}_G$ under a variable assignment β is defined as follows:*

$$t^{I, \beta} := \begin{cases} \beta(t), & \text{if } t \in \text{Var}_{\mathcal{L}} \\ I(t), & \text{if } t \in \text{Const}_{\mathcal{L}} \end{cases}$$

We now introduce the truth conditions for the formulas of the language, i.e., we specify under which conditions a formula φ is true relative to an interpretation at a pair of worlds (w, v) under a variable assignment β where w is the circumstance of evaluation and v the context.

³⁹This means that ways of believing are represented by ordinals in our semantics.

⁴⁰We leave it open which properties the doxastic accessibility relation needs to satisfy.

⁴¹Importantly, we do **not** impose the following constraints for **all** sentences φ, ψ :

$$\text{if } \varphi \neq \psi, \text{ then } \mathcal{B}(w, \varphi, \beta) \neq \mathcal{B}(w, \psi, \beta).$$

In other words, we **do not require** \mathcal{B} to be injective relative to the syntactical argument position. This is important for otherwise we would make our proposal vulnerable to the so-called translation argument (cf. Church, 1950; Kratzer, 2022, for discussion).

Definition 6 (Truth in a Model). Let $F = \langle G, \text{Dox}, B \rangle$ be a belief frame and $I \in \text{Int}_G$ an interpretation. Then a formula φ is true in the belief model $M = (F, I)$ at a world w and context v under assignment β iff

- (i) $M, (w, v) \Vdash s = t[\beta]$ iff $s^{I, \beta} = t^{I, \beta}$
- (ii) $M, (w, v) \Vdash s \neq t[\beta]$ iff $s^{I, \beta} \neq t^{I, \beta}$
- (iii) $M, (w, v) \Vdash Pt_1, \dots, t_n[\beta]$ iff $\langle t_1^{I, \beta}, \dots, t_n^{I, \beta} \rangle \in I^+(w, P)$
- (iv) $M, (w, v) \Vdash \neg Pt_1, \dots, t_n[\beta]$ iff $\langle t_1^{I, \beta}, \dots, t_n^{I, \beta} \rangle \in I^-(w, P)$
- (v) $M, (w, v) \Vdash \neg \neg \psi[\beta]$ iff $M, (w, v) \Vdash \psi[\beta]$
- (vi) $M, (w, v) \Vdash \psi \wedge \chi[\beta]$ iff $M, (w, v) \Vdash \psi[\beta]$ & $M, (w, v) \Vdash \chi[\beta]$
- (vii) $M, (w, v) \Vdash \neg(\psi \wedge \chi)[\beta]$ iff $M, (w, v) \Vdash \neg \psi[\beta]$ or $M, (w, v) \Vdash \neg \chi[\beta]$
- (viii) $M, (w, v) \Vdash B\psi[\beta]$ iff for all $u \in W(\text{Dox}_{B(v)}(w, u) \Rightarrow M, (u, v) \Vdash \psi[\beta])$
- (ix) $M, (w, v) \Vdash \neg B\psi[\beta]$ iff for some $u \in W(\text{Dox}_{B(v)}(w, u) \& M, (u, v) \Vdash \neg \psi[\beta])$
- (x) $M, (w, v) \Vdash \forall v \psi[\beta]$ iff for all $d \in D : M, (w, v) \Vdash \psi[\beta(v : d)]$
- (xi) $M, (w, v) \Vdash \neg \forall v \psi[\beta]$ iff for some $d \in D : M, (w, v) \Vdash \neg \psi[\beta(v : d)]$

$B(v)$ in (viii) and (ix) is short for $B_I(v, \psi, \beta)$ and $B_I(v, \neg \psi, \beta)$ respectively.

Let us connect these truth conditions to the relational analysis and show how (viii) may be used to define a ternary belief relation between an agent, a proposition and a way of believing. In the present framework a sentence expresses a proposition relative to a context. Let φ be a sentence of \mathcal{L} and v a context. Then the proposition $\|\varphi\|_v$ expressed by φ in context v is defined as follows:

$$\{w \mid M, (w, v) \Vdash \varphi\}$$

and we may define the ternary belief relation at a world between an agent, a proposition, and a way of believing. In fact, since we are only considering one agent, we may drop explicit mention of the agent, that is, we end up with a binary relation between propositions and ways of believing:

$$\text{bel}_w(\|\varphi\|_v, B_I(v, \varphi, \beta)) : \leftrightarrow \forall u (wR_{B_I(v, \varphi, \beta)} u \Rightarrow M, (u, v) \Vdash \varphi[\beta]).$$

5.2 Soames argument reconsidered

With the definition of the belief relation in place we may evaluate Soames's argument in the context of our semantics. First, we check whether our semantics meets all the conditions we laid out at the beginning of Section 4. We claim that our semantics still respects **Relational analysis**: we merely have to be attentive to role of contextual parameters when applying the relational analysis. Since the interpretation of an individual constant does not depend on worlds, the semantics also aligns with the **DR/CIC**-view.

Compositionality The semantics is compositional. The semantic content of a belief report is composed out of the semantic contents of its parts: let φ, ψ be sentences of \mathcal{L} , $\alpha \in \text{WB}_F$, and $w, v \in W$. Then

$$(C) \quad \text{if } \|\varphi\|_v = \|\psi\|_v \text{ and } \text{bel}_w(\|\varphi\|_v, \alpha), \text{ then } \text{bel}_w(\|\psi\|_v, \alpha).$$

This is of course trivial. The trick that guarantees that compositionality is preserved is the introduction of the way of believing parameter. Since we are not guaranteed that for a given context v

$$B_I(v, \varphi, \beta) = B_I(v, \psi, \beta)$$

even if $\|\varphi\|_v = \|\psi\|_v$, (C) does not imply

$$M, (w, v) \Vdash B\varphi \text{ iff } M, (w, v) \Vdash B\psi.$$

Some may object that the semantics is, perhaps, compositional in a technical sense, but ultimately a quotational (and thus non-compositional) approach in spirit, as the way of believing is allowed to depend on the words chosen to report a belief. To fully address this complaint more will ultimately need to be said about the particular type of context-sensitivity and the resolution of the contextual parameter proposed. Yet, while, on our account, the way of believing depends on the way a belief is reported it is not the case that different ways of reporting always lead to different ways of believing (cf. Footnote 41). Moreover, if our semantics were just a quotational approach, then the version of the problem of necessary truths affecting our semantics should not arise (see Section 6).

DoC As it stand our semantics does not satisfy **DoC**, that is, inferences of the following form:⁴²

$$B(\varphi \wedge \psi) \Rightarrow B\varphi \wedge B\psi.$$

This can be rectified by adding the following constraint to our semantics:

(\wedge D) for all $\varphi_1, \varphi_2 \in \mathcal{L}$ and $v \in W$:

$$\text{Dox}_{B_I(v, \varphi_1 \wedge \varphi_2, \beta)} \subseteq \text{Dox}_{B_I(v, \varphi_i, \beta)}$$

for $i \in \{1, 2\}$.

Disquotation Our language does not contain a truth predicate and without the truth predicate we cannot check whether our view is compatible with **Disquotation**. Of course, if we work in classical possible world semantics we know that, as a consequence of Tarski's undefinability theorem,

$$\|Tt_\varphi\| = \|\varphi\|$$

⁴²We think of \Rightarrow as a sequent arrow. Since we are working in a non-classical semantics we are not guaranteed that we have a conditional that satisfies the deduction theorem at our disposition. This means that we cannot move from inferences to principles/schemata and back.

cannot hold for every sentence φ . The situation changes if we move to a version of possible world semantics where worlds are non-classical, i.e., paracomplete and/or paraconsistent. In that case we can apply Kripke's theory of truth (Kripke, 1975) to introduce a disquotational truth predicate to \mathcal{L} relative to our semantics, that is, we have $\|Tt_\varphi\|_v = \|\varphi\|_v$ for every sentence φ of the extended language and every context v . We refer to Stern (2021) for further details of the construction.

Monotonicity Our semantics is monotone. In our semantics, worlds are not strictly speaking truthmakers, as they do not convey semantic information in their own right. Rather semantic information is bestowed upon them by the interpretation function under consideration. For this reason, as mentioned in Footnote 37 monotonicity is not expressed by (MON) but by the following principle which can be easily verified in our semantics: let I, J be two interpretation function, F a belief frame, $w, v \in W$ and φ a sentence of the language. Then

$$\text{if } I \leq J \text{ and } (F, I), (w, v) \Vdash \varphi, \text{ then } (F, J), (w, v) \Vdash \varphi$$

where

$$I \leq J \text{ iff } [I^+(w, P) \subseteq J^+(w, P) \text{ and } I^-(w, P) \subseteq J^-(w, P) \text{ for all } w \in W \text{ and } P \in \text{Pred}].$$

Soames's argument After having convinced ourselves that all of Soames's assumption are met it remains to check whether Soames's argument is indeed blocked in our semantics. As before, let $\varphi(p)$ stand for 'Hesperus' refers to Hesperus and 'Phosphorus' refers to Phosphorus'.⁴³ Now let $@$ be the actual world and v the context of utterance (plausibly we could also have $@$ as the relevant context of utterance). Then 1 tells us that

$$\text{bel}_@ (\|\varphi(p)\|_v, \mathcal{B}(v, \varphi(p)))$$
⁴⁴

and assuming the DR/CIV-view we may infer

$$\text{bel}_@ (\|\varphi(h)\|_v, \mathcal{B}(v, \varphi(p))).$$

However 2 is true, if

$$\text{bel}_@ (\|\varphi(h)\|_v, \mathcal{B}(v, \varphi(h))),$$

and the latter is only guaranteed if

$$\mathcal{B}(v, \varphi(p)) = \mathcal{B}(v, \varphi(h)).$$

But unless special conditions on ways of believing are imposed the identity will not hold and so the inference from 1 to 2 will break down. From the perspective of our possible world semantics the agent considers different doxastic alternatives due to the different ways of believing. As a

⁴³We assume that the English sentences have have been a formalized in a suitable way, i.e., we take $\varphi(p)$ to be a sentence of \mathcal{L} .

⁴⁴We omit the assignment function: since $\varphi(p)$ is sentence the assignment will not have an effect on how $\|\varphi(p)\|_v$ is believed.

consequence, not only will the step from 1 to 2 no longer go through, but also the inference from 2 to 3 is blocked. Moreover, the same mechanism will also block the equivalence of (c) and (d), that is, the Max-example.

The doxastic semantics proposed offers a lot of flexibility, arguably too much flexibility. For one, as mentioned ultimately more needs to be said about ways of believing. For another, the semantics does not give rise to an interesting logic of belief. Indeed, Crimmins and Perry (1989) argued that “*there is little possibility of an interesting logic of belief*” (cf. Crimmins and Perry, 1989, p. 711). While Crimmins and Perry (1989) are certainly right if we consider the unconstrained version of our contextualist semantics, we can add constraints that will enable us to maintain certain regularities about belief. As a matter of fact, we have already seen one such constraint, i.e. ($\wedge D$), which guaranteed that belief would distribute over conjunctions. We present a number of constraints in Appendix A and show that a surprising amount of “logic” can be recovered without falling prey to Soames’s argument.

6 Discussions: Limitations and Comparison

We mentioned in Section 4 that Goodman and Lederman (2021) can be considered as providing guidelines for developing contextualist semantics within a possible world framework and the question arises how their proposal compares to ours. Their idea is that agents believe a proposition under some perspective, which, in turn, depends on the context of utterance. On their view a perspective is a set of mentalese sentences, that is, a set of guises under which our beliefs qua propositions are presented to us in a given context.⁴⁵ Perhaps the best way of comparing the two semantics is to conceive of their contexts as Ways of Believing and to associate a perspective with a set of sentences of \mathcal{L} . To this effect we let $\pi : \text{WB}_F \rightarrow \mathcal{P}(\text{Sent}_{\mathcal{L}})$ be a function that yields perspectives relative to every element of WB_F , that is, relative to every context in the sense of Goodman and Lederman (2021). A further crucial notion of their semantics is that of an agent’s belief-box, that is, the set of sentences believed by an agent (at a given world). We let $\mathcal{D} : W \rightarrow \mathcal{P}(\text{Sent}_{\mathcal{L}})$ be the function that specifies the agent’s belief box relative to a world. With this in place Goodman and Lederman’s truth conditions for belief reports can be stated as follows:⁴⁶

$$M, (w, v) \models \text{B}\psi \text{ iff for some } \chi \in \pi(\mathcal{B}(\psi, v)) : \|\psi\|_v = \|\chi\|_v \text{ and } \chi \in \mathcal{D}(w).$$

A major difference between their truth conditions and ours is that the agent’s belief-box is not semantically controlled, but given by an ad hoc stipulation. Of course, this provides a lot of flexibility for accommodating a wide array of data, but it also makes their semantics more vulnerable to objections from overfitting: without further information and constraints an agent’s

⁴⁵Goodman and Lederman (2021) understand contexts in the sense of Kaplan (1989).

⁴⁶We made two simplifying assumption in stating the truth conditions:

- we assume that that worlds are total and no clause for negated belief reports is required;
- we only provide truth conditions for sentences, i.e., we do not consider formulas with free variables in belief contexts. If we were to give a systematic semantics we need to specify how to handle such formulas, which presumably means making perspective not only dependent on contexts but also on assignment functions.

belief-box remains a black box without semantic justification. Indeed, in their paper Goodman and Lederman (2021) discuss constraints that can be imposed on their semantics (similar to our discussion in the appendix to the paper) and which have the consequence of asserting some semantic control over the belief set. These constraints do not, in their own right, provide semantic justification for an agent’s belief, but they might open up the possibility of such semantic justification in the framework of possible world semantics. If a specific combination of constraints are imposed, full semantic control over the belief set can be exerted and we can recast the truth conditions as follows:

$M, (w, v) \Vdash B\psi$ iff for some $\chi \in \pi(\mathcal{B}(\psi, v)) : \|\psi\|_v = \|\chi\|_v$ and for all u s.t. $wRu: M, (u, v) \Vdash \chi$.⁴⁷

Let’s call this Goodman and Lederman’s *worldly semantics*. Qua worldly semantics Goodman and Lederman’s semantics can be viewed as a version of Awareness semantics (Fagin et al., 1995), but where the Awareness set/operator is contextually controlled. Ultimately, this observation points to a key difference between their and our semantics. According to our semantics it is possible that

- ★ The ancients believed that ‘Hesperus’ refers to Hesperus.
- † The ancients believed that ‘Hesperus’ did not refer to Phosphorus.

are both true at a world—importantly, two different ways of believing are at play.⁴⁸ This will not be possible on the worldly semantics (nor on versions of the semantics in which the belief-box is required to express a consistent set of propositions), as an agent’s belief-box is not affected by context. What is altered by the context is which propositions in the belief-box can be grasped, i.e., the agent is aware of. This means that on the worldly semantics the best we can say is that the ancients did not believe that ‘Hesperus’ refers to Phosphorus but nothing more.

Yet, a version of the problem just sketched also plagues our semantics: it was possible for ★ and † both to be true at a world because the way of believing associated with ★ led to a doxastic accessibility relation which selects for worlds in which ‘Hesperus’ refers to Hesperus’ was true whereas the way of believing associated with † led to a doxastic accessibility relation which selects for worlds in which ‘Hesperus’ refers to Hesperus’ was false. This works well for sentences that express contingent propositions, but it won’t work for sentences that express propositions that are true at every world, as the negation of such a proposition will not be true in any world. This implies that

=1 The ancients believed that Hesperus is Phosphorus.

is true at every world in our semantics and we cannot find a world in which

=2 The ancients believed that Hesperus is not Phosphorus.

⁴⁷Perhaps there is a more subtle way of attaching a semantic rationale to an agent’s belief-box, but within the framework of possible world semantics this seems to be an obvious way, and what we say does not really depend on the particulars.

⁴⁸Notice, however, that unless we allow for paraconsistent worlds it will not be true that:

- The ancients believed that ‘Hesperus’ refers to Hesperus *and* ‘Hesperus’ did not refer to Phosphorus.

comes out true since ‘Hesperus is Phosphorus’ is true at every possible world.⁴⁹

What moral should we draw from this, admittedly counterintuitive, observation? First, it is worth noting that it puts us precisely in the same position as structured proposition theorists like Soames (1987) or King (2007), who employ pragmatic explanations to account for the contravening intuition.⁵⁰ This means that Soames won’t be able to use the observation to undermine truth-conditional semantics. Rather, this sort of problem is, to a certain extent, to be expected if we accept the **DR/CIC-view**. Second, as mentioned, pragmatic explanations have been proposed to handle these so-called Frege cases and these explanations can be used to supplement our semantics. In a nutshell, pragmatic explanations typically hold that agents believe that *Hesperus is Hesperus* and that *Hesperus is Phosphorus* in different ways, which, so the story goes, is why we are squeamish about accepting =1 (but happy to accept =3 below). In our semantics the fact that we believe the same proposition in different ways in =1 and

=3 The ancients believed that Hesperus is Hesperus.

is encoded into our semantics. If a pragmatic explanation that exploits different ways of believing is at all successful, it thus surely is against the backdrop of our semantics.

The fact that we can try to handle these identity puzzles pragmatically and that the puzzles are a problem for many if not most alternative semantics, does not mean that we should disregard them. Are their appealing options for handling these puzzles semantically? We take it that semanticists who finds =1 unacceptable and who, in addition, wants =2 to be true in their semantics will ultimately have three options:⁵¹

- reject **DR/CIC-view**, or
- reject compositional semantics for belief reports, or, similarly,
- reject semantic control over an agent’s belief set (cf. our discussion of Goodman and Lederman’s proposal at the beginning of the section).

Considering these options we take it that proponents of the **DR/CIC-view** should opt for explaining the intuitions regarding =1 and =2 pragmatically. Our semantics provides them with a good starting point for doing so.

7 Conclusion

We started the paper with a discussion of the problem of overfitting vs underfitting of semantics, and it seems natural to evaluate the proposed semantics against this background. From

⁴⁹More generally, our contextualist semantics will not help with the problem of necessary truth/falsity. Rather for dealing with this problem one arguably needs to assume partial and, perhaps, paracomplete worlds.

⁵⁰Lederman (2021) also points out that a similar problem affects the semantics for de re beliefs advanced by Percus and Sauerland (2003). For reasons of space we omit a detailed comparison with the latter semantics. Roughly put, their semantics is slightly more liberal when it comes to belief ascriptions, but more restrictive when it comes to establishing disbelief.

⁵¹Semanticists who merely wish to reject =1 but don’t require that =2 is true have a further option, namely, that \mathcal{B} is a partial function, that is, contexts do not produce ways of believing for every sentence of the language.

our perspective the semantics shows how to resist Soames-style arguments that aim to establish that truth-conditional semantics necessarily suffers from underfitting. Importantly, the semantics is monotone and compositional in the strong sense and, more general, compatible with all of Soames's initial assumptions. Admittedly, on this semantics the attitudinal relation depends on a contextual parameter and may thus be reconstructed as a ternary rather than a binary relation. Yet, the analysis of attitude reports is most certainly in alignment with the idea of the relational analysis of propositional attitude reports. Still the introduction of a contextual parameter to the semantics of attitude reports is contested and may be claimed to be an instance of overfitting, but it may also be claimed to be a symptom of underfitting. Both overfitting- but also underfitting-theorists will argue that the additional parameter was an ad hoc manoeuvre to account for data that should not be explained semantically according to the overfitting theorist or, according to the underfitting theorists, that should be explained by assuming a more fine grained individuation of content. In contrast, we think that Section 4 makes a good *prima facie* case for invoking contextual parameters in the analysis of attitude reports and that Crimmins and Perry's (1989) idea deserves further investigation in the context of natural language semantics. However, more details need to be provided about the alleged context sensitivity of the attitude verb and of how to conceive of our contextual parameters, that is, our ways of believing to fully evaluate the semantics.

Putting aspects pertaining to natural language semantics aside, our semantics is arguably the first fully developed context-sensitive possible world semantics for belief. The semantics should thus be of interest to researchers working in doxastic and epistemic logic, who are interested in the flexibility of approaches that build on hyperintensional accounts of content, but who prefer working within possible world semantics.

Acknowledgements

I wish to thank the editors of this special issue for giving me the opportunity to contribute to this issue. Without their invitation the paper would most likely have never been written. I would also like to thank the audiences at the *Hyperintensional Formal Epistemology Workshop* in Berlin and the *Institutskolloquium* of the Department of Philosophy at the University of Tübingen for helpful and constructive discussions. Finally, I owe a lot to Poppy Mankowitz, Anthony Everett as well as two anonymous referees of this journal for very helpful comments on earlier versions of the paper.

Initial work on this paper was supported by the ERC Starting Grant *Truth and Semantics* (TRUST, Grant n° 803684).

References

- Abbott, B. (2010). *Reference*, volume 2. Oxford University Press, Oxford.
- Aloni, M. (2005). Individual concepts in modal predicate logic. *Journal of Philosophical Logic*, 34:1–64.
- Barwise, J. and Perry, J. (1983). *Situations and attitudes*. MIT Press.

- Berto, F. (2024). Hyperintensionality and overfitting. *Synthese*, 203(4):117.
- Berto, F. and Jago, M. (2009). *Impossible Worlds*. Oxford University Press, Oxford, UK.
- Carnap, R. (1947). *Meaning and Necessity*. University of Chicago Press, Chicago.
- Charlow, S. and Sharvit, Y. (2014). Bound ‘de re’ pronouns and the LFs of attitude reports. *Semantics and Pragmatics*, 7:3–1.
- Church, A. (1950). On Carnap’s analysis of statements of assertion and belief. *Analysis*, 10:97–99.
- Cresswell, M. J. (1975). Hyperintensional logic. *Studia Logica*, 34(1):25–38.
- Cresswell, M. J. (1985). *Structured meanings: The semantics of propositional attitudes*. MIT press.
- Cresswell, M. J. and Von Stechow, A. (1982). “de re” belief generalized. *Linguistics and Philosophy*, pages 503–535.
- Crimmins, M. and Perry, J. (1989). The prince and the phone booth: Reporting puzzling beliefs. *The Journal of Philosophy*, 86(12):685–711.
- Davidson, D. (1968). On saying that. In *Inquiries into Truth and Interpretation*, pages 93–108. Oxford. 2001.
- Elbourne, P. (2010). Why propositions might be sets of truth-supporting circumstance. *Journal of Philosophical Logic*, 39:101–111.
- Fagin, R., Halpern, J. Y., Moses, Y., and Vardi, M. Y. (1995). *Reasoning about Knowledge*. MIT Press.
- Fine, K. (2017a). A theory of truthmaker content i: Conjunction, disjunction and negation. *Journal of Philosophical Logic*.
- Fine, K. (2017b). Truthmaker semantics. In Hale, B., Wright, C., and Miller, A., editors, *A Companion to the Philosophy of Language*. Wiley, Chichester, UK.
- Fine, K. and Jago, M. (2019). Logic for Exact Entailment. *The Review of Symbolic Logic*, 12(3):536–556.
- Fitting, M. and Mendelsohn, R. L. (1998). *First-Order Modal Logic*. Kluwer Academic Publishers.
- Gamut, L. (1991). *Logic, Language, and Meaning, Volume 2: intensional logic and logical grammar*. University of Chicago Press.
- Georgi, G. (2019). Propositions, representation, and truth. *Synthese*, 196:1019–1043.
- Goodman, J. and Lederman, H. (2021). Perspectivism. *Noûs*, 55(3):623–648.
- Heim, I. and Kratzer, A. (1998). *Semantics in Generative Grammar*. Blackwell Publishing, Oxford.

- Hintikka, J. (1962). *Knowledge and Belief*. Cornell University Press, Ithaca and London.
- Hoek, D. (2022). Questions in action. *Journal of Philosophy*, 3:113–143.
- Kaplan, D. (1968). Quantifying in. *Synthese*, 19:178–214.
- Kaplan, D. (1989). Demonstratives. An essay on the semantics, logic, metaphysics, and epistemology of demonstratives and other indexicals. In Almog, J., Perry, J., and Wettstein, H., editors, *Themes from Kaplan*, pages 481–564. Oxford University Press.
- King, J. C. (2007). *The nature and structure of content*. Oxford University Press.
- Kratzer, A. (1986). Conditionals. *Chicago Linguistics Society*, 22(2):1–15.
- Kratzer, A. (2022). Attitude ascriptions and speech reports. In Altshuler, D., editor, *Linguistics Meets Philosophy*, pages 17–50. Cambridge University Press, New York.
- Kratzer, A. (2023). Situations in Natural Language Semantics. In Zalta, E. N. and Nodelman, U., editors, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Fall 2023 edition.
- Kripke, S. (1975). Outline of a theory of truth. *The Journal of Philosophy*, 72:690–716.
- Larson, R. K. and Ludlow, P. (1993). Interpreted logical forms. *Synthese*, 95(3):305–355.
- Lederman, H. (2021). Fine-grained semantics for attitude reports. *Semantics and Pragmatics*.
- Lewis, D. (1970). General semantics. *Synthese*, 22(1/2):18–67.
- Lewis, D. (1979). Attitudes de dicto and de se. *The philosophical review*, 88(4):513–543.
- Mayr, C. and Schmitt, V. (2024). Non-de dicto construals as a unified phenomenon. In *Semantics and Linguistic Theory*, pages 401–421.
- Moltmann, F. (2020). Truthmaker semantics for natural language: Attitude verbs, modals, and intensional transitive verbs. *Theoretical Linguistics*, 46(3-4):159–200.
- Percus, O. and Sauerland, U. (2003). On the LFs of attitude reports. In *Proceedings of sinn und bedeutung*, volume 7, pages 228–242.
- Picenni, S. (2023). *Exact Truthmaking, Hyperintensionality, and Paradoxes*. PhD thesis, University of Bristol.
- Quine, W. V. O. (1956). Quantifiers and propositional attitudes. *The Journal of Philosophy*, 53.
- Soames, S. (1987). Direct reference, propositional attitudes, and semantic content. *Philosophical Topics*, 15(1):47–87.
- Soames, S. (2005). *Reference and Description*. Princeton University Press, Princeton, New Jersey.
- Stalnaker, R. (1987). Semantics for belief. *Philosophical Topics*, 15(1):177–190.

- Stalnaker, R. C. (1978). Assertion. In Partee, B. and Portner, P., editors, *Formal Semantics*, pages 147–161. Blackwell Publishing.
- Stern, J. (2021). Belief, truth, and ways of believing. In Nicolai, C. and Stern, J., editors, *Modes of Truth: The Unified Approach to Modality, Truth, and Paradox*, pages 151–181. Routledge.
- Swanson, E. (2011). Propositional attitudes. In Maienborn, C., von Stechow, K., and Portner, P., editors, *Semantics: An international handbook of natural language meaning*, volume 2, pages 1538–1560. Walter de Gruyter.
- Williamson, T. (2020). *Suppose and Tell: The semantics and heuristics of conditionals*. Oxford University Press, Oxford.
- Yalcin, S. (2014). Semantics and metasemantics in the context of generative grammar. In Burgess, A. and Sherman, B., editors, *Metasemantics: New essays on the foundations of meaning*, volume 17. Oxford University Press Oxford.
- Yalcin, S. (2018). Belief as question-sensitive. *Philosophy and Phenomenological Research*, 97(1):23–47.

A Constraints: Towards a Logic of Belief

We discussed that by adding the constrain $(\wedge D)$ we can force belief to distribute over conjunction. However, $(\wedge D)$ is a particular instance of a more general constraint. Before we can discuss such more general constraints we introduce the notion of a *positive decomposition*. We note that in this paper and appendix we focus on the propositional logic and structure of belief and put the quantificational structure aside.

Definition 7. Let $\varphi \in \mathcal{L}$. Then the positive decomposition of φ , $\text{Pos}(\varphi)$, is defined by the following recursion:

1. if φ is a literal, or of the form $\forall x\psi$ or $\neg\forall x\psi$, then

$$\text{Pos}(\varphi) = \{\varphi\};$$

2. if $\varphi \doteq \neg\neg\psi$, then

$$\text{Pos}(\varphi) = \{\varphi\} \cup \text{Pos}(\psi);$$

3. if $\varphi \doteq B\psi$, then

$$\text{Pos}(\varphi) = \{\varphi\} \cup \text{Pos}(\psi);$$

4. if $\varphi \doteq \neg B\psi$, then

$$\text{Pos}(\varphi) = \{\varphi\} \cup \text{Pos}(\neg\psi);$$

5. if $\varphi \doteq \psi \wedge \chi$, then

$$\text{Pos}(\varphi) = \{\varphi\} \cup \text{Pos}(\psi) \cup \text{Pos}(\chi);$$

6. if $\varphi \doteq \neg(\psi \wedge \chi)$, then

$$\text{Pos}(\varphi) = \{\varphi\} \cup \text{Pos}(\neg\psi) \cup \text{Pos}(\neg\chi).$$

With this notion in place, $(\wedge D)$ can be seen as a particular instance of the following constraint:

Decomposition for all $\varphi \in \mathcal{L}$, $v \in W$, and assignments β :

$$\text{if } \psi \in \text{Pos}(\varphi), \text{ then } \text{Dox}_{B_I(v, \varphi, \beta)} \subseteq \text{Dox}_{B_I(v, \psi, \beta)}.$$

With Decomposition in place belief frame will validate a number of interesting inferences:

Proposition 8. *Let F be a belief frames that respects **Decomposition**, then*

- $F \Vdash B\neg\neg\varphi \Rightarrow B\varphi$;
- $F \Vdash B(\varphi \wedge \psi) \Rightarrow B\varphi \wedge B\psi$;
- if Dox_α is dense for all $\alpha \in \text{WB}_F$, then

$$F \Vdash \text{BB}\varphi \Rightarrow B\varphi.$$

- if Dox_α has the First Introspection Property for all $\alpha \in \text{WB}_F$, then

$$F \Vdash B\neg B\varphi \Rightarrow B\neg\varphi.$$

Let us introduce the Introspection property alluded to:

Definition 9 (Introspection Properties). *Let $\alpha \in \text{WB}_F$ and $w \in W$. Then the First Introspection Property is the following condition:*

$$\text{for all } v \in W : \text{if } \text{Dox}_\alpha(w, v), \text{ then for some } u \in W : \\ \text{Dox}_\alpha(w, u) \text{ s.t. } z = v \text{ for all } z \in W \text{ with } \text{Dox}_\alpha(u, z);$$

whereas the Second Introspection Property is the condition:

$$\text{for all } v \in W : \text{if } \text{Dox}_\alpha(w, v), \text{ then for some } u \in W : \text{Dox}_\alpha(w, u) \& \text{Dox}_\alpha(v, u).$$

Decomposition analyses the positive built up of a formula χ and, in some sense, takes this structure to be transparent: the ways of believing produced by χ “encompasses” the ways of believing of its subformulas. What about the converse direction, that is, the direction where we start from formulas and move to more complex formulas that have been constructed out of the initial formulas? We are led to the notion of *positive closure*.

Definition 10 (Positive Closure). *Let $X \subseteq \mathcal{L}$. Then $X^{\text{cl}} \subseteq \mathcal{L}$ is defined by the following induction:*

1. if $\varphi \in X$, then $\varphi \in X^{\text{cl}}$;

2. if $\varphi \in X^{\text{cl}}$, then $\neg\neg\varphi \in X^{\text{cl}}$;
3. if $\varphi, \psi \in X^{\text{cl}}$, then $\varphi \wedge \psi \in X^{\text{cl}}$;
4. if $\neg\varphi, \neg\psi \in X^{\text{cl}}$, then $\neg(\varphi \wedge \psi) \in X^{\text{cl}}$;
5. nothing else is in X^{cl} .

The first constraint Weak Closure is essentially the counterpart to Decomposition:

Weak Closure For all $X \subseteq \mathcal{L}$, $v \in W$, and assignments β :

$$\bigcap_{\psi \in X} \text{Dox}_{B_I(v, \psi, \beta)} \subseteq \bigcap_{\psi \in X^{\text{cl}}} \text{Dox}_{B_I(v, \psi, \beta)}.$$

Most importantly, Weak Closure gives us Double Negation Introduction and Or-Introduction in belief contexts.

Proposition 11. *Let F be a belief frame that respects Weak Closure, then*

- $F \Vdash B\varphi \Rightarrow B\neg\neg\varphi$;
- $F \Vdash B\neg\varphi \wedge B\neg\psi \Rightarrow B\neg(\varphi \wedge \psi)$;
- if Dox_α is transitive for all $\alpha \in \text{WB}_F$, then

$$F \Vdash B\varphi \Rightarrow BB\varphi.$$

- if Dox_α is has the Second Introspection Property for all $\alpha \in \text{WB}_F$, then

$$F \Vdash B\neg\varphi \Rightarrow B\neg B\varphi.$$

Weak closure is not sufficient for Conjunction Introduction in belief context. A stronger closure condition is required:

Strong Closure For all $X \subseteq \mathcal{L}$, $\chi \in X^{\text{cl}}$, $v \in W$, and assignments β :

$$\text{if } \bigcup_{\psi \in X} \text{Pos}(\psi) = \text{Pos}(\chi), \text{ then } \bigcup_{\psi \in X} \text{Dox}_{B_I(v, \psi, \beta)} \subseteq \text{Dox}_{B_I(v, \chi, \beta)}.$$

Lemma 12. Strong Closure implies Weak Closure.

Proposition 13. *Let F be a belief frame that respects Strong Closure and Decomposition, then*

- $F \Vdash B\varphi \Leftrightarrow B\neg\neg\varphi$;
- $F \Vdash B\varphi \wedge B\psi \Leftrightarrow B(\varphi \wedge \psi)$;
- $F \Vdash B\neg\varphi \wedge B\neg\psi \Rightarrow B\neg(\varphi \wedge \psi)$;

- if Dox_α is dense and transitive for all $\alpha \in \text{WB}_F$, then

$$F \Vdash \text{B}\varphi \leftrightarrow \text{BB}\varphi.$$

- if Dox_α has both the First and the Second Introspection Property for all $\alpha \in \text{WB}_F$, then

$$F \Vdash \text{B}\neg\varphi \leftrightarrow \text{B}\neg\text{B}\varphi.$$

Even if Decomposition and Strong Closure are adopted Soames's argument cannot be carried out as the step from 1 to 2 will not go through. For this step to go through we need to impose a constraint that ways of believing are substitution invariant:

Substitution Let F be a belief frame, c_1, c_2 individual constants and I an interpretation. Then

$$\text{if } I(c_1) = I(c_2), \text{ then } \mathcal{B}_I(v, \varphi, \beta) = \mathcal{B}_I(v, \varphi[c_1/c_2], \beta).$$

If Substitution is imposed we can derive 2 of Soames's argument, that is,

$$a \text{ bel } \|\varphi(h)\|$$

but we still cannot derive 3, i.e.,

$$a \text{ bel } \|\varphi(h) \wedge \exists x\varphi\|.$$

Finally, as we have seen 3 to 4 requires Decomposition.

This does not mean that imposing all these constraints is without costs, but it shows that Soames's argument is blocked not because of an inherent predictive weakness of the semantics, but for principled reasons. For example, Strong Closure and Substitution taken together rule out the possibility that

- The ancients believed that 'Hesperus' refers to Hesperus.
- The ancients believed that 'Hesperus' did not refer to Phosphorus.

are both true at a world w and context v , that is, if we disallow paraconsistent worlds.

An option available to theorists that on the one hand are attracted by the flexibility of the semantics, but, on the other hand, would also like to have the possibility to employ suitable doxastic logics in transparent cases is to introduce an awareness operator (cf. Fagin et al., 1995). Such an operator could be used to move back and forth between contextualized and non-contextualized belief ascriptions.⁵²

⁵²See Stern (2021) for an implementation of the idea in a related semantics, although Stern works with an awareness predicate.