Counterfactual attitude contents and the semantics of plurals in belief contexts[1]

Nina HASLINGER — *Georg-August-Universität Göttingen*
Viola SCHMITT — *Humboldt-Universität Berlin*

Abstract. Indefinites in the complement clause of *believe* may participate in cumulative relations with a plural attitude subject, even under an intensional reading (Schmitt 2020; see also Pasternak 2018). Schmitt (2020) provides a compositional semantics for such constructions, where indefinites introduce pluralities of partial individual concepts. Crucially, however, such pluralities can only be formed if, intuitively, the concepts are 'distinct enough' (Haslinger and Schmitt to appear). We here implement this 'distinctness constraint' compositionally, combining Schmitt's (2020) ideas with a new proposal concerning the semantics of plural indefinites: We argue that the latter introduce a distinctness requirement that appeals to the attitude subjects' counterfactual beliefs – a requirement that is visible in cumulative belief sentences, but is trivialized (and thus unnoticeable) in extensional contexts. To compose our new DP-denotations with other material in the sentence, we use a version of Yalcin's (2007) 'domain semantics'.

Keywords: attitudes, cumulativity, Plural Projection, counterfactuals, intentional identity

1. Introduction

This paper develops a descriptively adequate, fully compositional analysis of **cumulative belief attributions**: In examples like (1a) or (1b), a plural expression in the complement of *believe* seems to participate in a cumulative relation with a plural attitude subject.

(1) a. *Abe and Bert believe that two monsters were roaming the castle.*
b. *Abe und Bert glauben, dass zwei Monster im Schloss waren.*
Abe and Bert believe that two monsters in.the castle were
'Abe and Bert believe that two monsters were in the castle.' German

For many (but not all) speakers, these sentences have a reading, brought out by the ZOMBIE VS. GRIFFIN scenario in (2), that is puzzling in two respects (see Pasternak 2018; Schmitt 2020, Haslinger and Schmitt to appear). First, it does not attribute the belief that there were *two* monsters in the castle to either Abe or Bert, so the 'distributive' paraphrase in (3) is inadequate. Second, the relevant reading is intensional: It does not entail the existence of monsters, and substituting a coextensional expression for *monster* will not necessarily preserve its truth value.

(2) ZOMBIE VS. GRIFFIN: Abe and Bert spent the night at Roy's castle. Abe believes in zombies, Bert in griffins. At 1am, Abe thinks there is a zombie in his room. At 2am, Bert thinks there is a griffin on his bed. They don't discuss it among themselves, but each of them tells Roy about his belief. In fact, there are no monsters at the castle. Roy says: "Well, I know that people find it spooky here, but ..." (1a), (1b) **%true**

(3) $$\mathbf{believe}_{w_0}(\lambda w.\exists x[\mathbf{monsters}_w(x) \wedge |x| = 2 \wedge \mathbf{in\text{-}castle}_w(x)])(\mathbf{a}) \wedge$$
$$\mathbf{believe}_{w_0}(\lambda w.\exists x[\mathbf{monsters}_w(x) \wedge |x| = 2 \wedge \mathbf{in\text{-}castle}_w(x)])(\mathbf{b})$$

[1]We thank Maria Barouni, Stergios Chatzikyriakidis, Rory Harder, Winnie Lechner, Magdalena Kaufmann, Rick Nouwen, Rob Pasternak, Orin Percus, Magdalena Roszkowski, Frank Sode, Giorgos Spathas, Clemens Steiner-Mayr, Tim Stowell, Peter Sutton, Marcin Wągiel, Thomas Weskott, Henk Zeevat, Ede Zimmermann, Sarah Zobel and audiences at SuB25, SinFonIJA 12, *Semantics in Athens*, the University of Oslo and the University of Cologne. VS's work was supported by the Austrian Science Fund (FWF), project P 32939, 'The typology of cumulativity'.

 In: P. G. Grosz, L. Martí, H. Pearson, Y. Sudo, and S. Zobel (eds.)
Proceedings of Sinn und Bedeutung 25, pp. 394–411. University College London and Queen Mary University of London.

The first puzzling property of (1a)/(1b) under the relevant reading, non-distributivity, can be accounted for in different ways. Pasternak's (2018) analysis of such examples makes use of a new lexical entry for *believe*, which encodes a notion of 'collective belief' when it combines with a plural subject. In this paper, we will adopt Schmitt's (2020) contrasting view that cases like (1a) and (1b) do not involve a special plural-sensitive meaning for *believe*, but are instances of cumulativity.[2] To see what this means, recall the standard truth conditions for a simple cumulative sentence without intensional operators, exemplified in (4) (see e.g. Scha 1981).

(4) *Abe and Bert read two books.*
'There are two distinct books x, y such that each of Abe and Bert read at least one of x and y and each of x and y was read by at least one of Abe and Bert.'

In (4), the subject plurality $\mathbf{a}+\mathbf{b}$ stands in a cumulative relation to some actual plurality of two books. So how could we give an analogous paraphrase for (1a)/(1b) that involves cumulativity, but does not entail the existence of a plurality of two monsters? We will approach this question by appealing to a pre-theoretical intuition taken from the philosophical literature on intentional identity (e.g. Edelberg 1986, 1992; cf. Geach 1967), namely that (1a)/(1b) involve beliefs about **intentional objects** or **belief objects** distinct from ordinary individuals. The crucial property of belief objects, for our purposes, is that i) they may 'pick out' different individuals in different worlds and ii) they do not have to 'pick out' an individual in every world. We can then give the following incomplete paraphrase of the truth conditions of (1a)/(1b):

(5) 'There are two distinct belief objects f and g corresponding to monsters such that each of Abe and Bert believes at least one of f and g is roaming the castle, and each of f and g is such that at least one of Abe and Bert believes it is roaming the castle.'

Thus, in the ZOMBIE VS. GRIFFIN scenario, Abe's attitude would involve a belief object that 'picks out' a zombie in Abe's room in any world where it is defined, and Bert's attitude would involve a belief object 'picking out' a griffin in Bert's bed in any world where it is defined. But to make sense of (5) and derive predictions from it, we must spell out what belief objects are and under which conditions two belief objects f and g count as 'distinct' for the purposes of plural semantics. That this is not trivial is illustrated by the alternative scenario in (6):

(6) UNSPECIFIED MONSTERS: [Roy's castle, no monsters exist ...] At 1am, Abe thinks there is a monster in his room, but cannot specify in more detail what the monster is like. At 2am, Bert thinks there is a monster in his room, but cannot specify what kind of monster it is. [They don't discuss it ...] (1a), (1b) **not true**

The contrast between (2) and (6) shows that Abe's belief that a monster was at the castle and Bert's belief that a monster was at the castle do not always add up to a cumulative belief with the content that two monsters were at the castle. To license the numeral *two*, Abe's belief object and Bert's belief object must meet a certain **distinctness condition**. In Haslinger and Schmitt (to appear), we argue that this contrast, and others like it, can be captured within a relatively conservative possible-worlds framework: If belief objects are modeled as partial individual concepts, the distinctness condition can be characterized in terms of the attitude subjects' counterfactual beliefs. But we still lack a compositional analysis of sentences as in (1) that incorporates this notion of distinctness (Schmitt (2020) and Pasternak (2018) do not discuss

[2]See Schmitt (2020) for potential ways of distinguishing between these two approaches.

the problem of distinctness in detail). Here we will fill this gap by combining the counterfactual distinctness condition with Schmitt's (2020) compositional semantics for sentences like those in (1): In our system, the distinctness condition is part of the basic semantics of plural DPs, regardless of whether they occur in a belief context. Our paper thus follows a more general research program that tries to analyze the non-trivial notions of individuation or distinctness underlying the semantics of ordinary DPs[3], and takes semantic dependencies between DPs in intensional contexts (Geach, 1967) to reveal otherwise unnoticed properties of DP semantics.

In Section 2, we present the counterfactual distinctness condition from Haslinger and Schmitt (to appear) and provide a paraphrase of (1a)/(1b) that incorporates this condition. Section 3 introduces the 'Plural Projection' approach to semantic composition in cumulative sentences (Haslinger and Schmitt, 2018; Schmitt, 2019, 2020) and gives an analysis of cumulative belief within this framework. Section 4 provides an analysis of counterfactual belief attributions based on Lewis's (1973) ordering semantics and Yalcin's (2007) 'domain semantics'. In Section 5, we put these pieces together to derive the truth conditions of (1a)/(1b). Section 6 briefly explores the consequences of our proposal for the semantics of DPs in extensional contexts.

2. Distinctness as a counterfactual notion

Let us first recapitulate the distinctness condition from Haslinger and Schmitt (to appear) and some of its motivation. Our starting point is a naive analysis of 'belief objects' as partial individual concepts (pICs), partial functions from possible worlds to truth values. On this view, *two monsters* in (1a) ranges over pluralities of two **monster**-concepts in the following sense:

(7) For any predicate $P \in D_{\langle s,et\rangle}$, a ***P*-concept** is a partial function $f \in D_{\langle s,e\rangle}$ such that for any world $w \in \mathrm{dom}(f)$, $P(w)(f(w)) = 1$.

Analyzing belief objects as functions suggests a straightforward notion of distinctness: Two pICs are distinct iff their domains are distinct or they yield distinct values for at least one world. Yet, the UNSPECIFIED MONSTERS scenario in (6) shows that this cannot be the linguistically adequate notion of distinctness. Intuitively, the belief objects in this scenario are not distinct enough to license the numeral *two*, but we could immediately find two individual concepts f and g that satisfy the notion of distinctness just described, for instance those in (8).[4]

(8) $f = [\lambda w :$ there was a monster in A's room in w.the monster that was in A's room in $w]$,
$g = [\lambda w :$ there was a monster in B's room in w.the monster that was in B's room in $w]$

Our first attempt at defining distinctness is thus insufficient, but there seems to be an obvious way to improve on it: Maybe what is wrong with f and g in (8) is that in *some* worlds, they return the same monster. So why don't we adopt (9), which disallows overlap between the pICs f and g, following existing work on 'conceptual covers' (Aloni, 2001; Schwager, 2007)?

(9) Two pICs f and g count as distinct iff for any $w \in \mathrm{dom}(f) \cap \mathrm{dom}(g)$, $f(w) \neq g(w)$.

(9) correctly predicts (1a)/(1b) to be fine in our ZOMBIE VS. GRIFFIN scenario: The pICs in (10) won't yield the same value in any world.[5] It also correctly predicts that (1a)/(1b) are bad

---

[3]See Aloni 2001; Condoravdi et al. 2001; Schwager 2007 for more arguments that individual concepts play a role in regular DP semantics, and Wągiel 2018 for related work on individuation in extensional contexts.
[4]From now on, we shorten our representation of individual concepts, using the ι-operator as a notational device.
[5]If it is impossible to be both a zombie and a griffin. Our final condition (15) no longer requires this assumption.

in our UNSPECIFIED MONSTERS scenario, as there will be worlds where Abe and Bert were visited by the same monster.

(10) $f = [\lambda w.\iota x.\textbf{zombie-in-A's-room}_w(x)], g = [\lambda w.\iota x.\textbf{griffin-in-B's-room}_w(x)]$

But (9) turns out to be inadequate in two respects. First, any two individual concepts f and g whose domains do not overlap are classified as distinct, as (9) is trivially satisfied. Hence, if the attitude subjects have disjoint sets of doxastic alternatives, we should always find two concepts that count as distinct for grammatical purposes. This prediction is incorrect: (1a)/(1b) are not true in the scenario in (11), but since Abe's and Bert's belief states are disjoint, (9) would lead us to expect that we should be able to find two non-overlapping pICs, e.g., those in (12).

(11) 3 MONSTERS VS. 4 MONSTERS: [Roy's castle, no monsters ...] Abe heard a sound and believes it was a monster. He read somewhere that there are exactly three monsters in the area. Bert heard also heard a sound and believes it was a monster, but he was told there are exactly 4 monsters in the area. (1a), (1b) **not true**

(12) $f = [\lambda w.\iota x.\textbf{monster}_w(x) \wedge \textbf{in-the-castle}_w(x) \wedge |\textbf{monster}_w| = 3]$,
$g = [\lambda w.\iota x.\textbf{monster}_w(x) \wedge \textbf{in-the-castle}_w(x) \wedge |\textbf{monster}_w| = 4]$

Second, since the no-overlap condition in (9) involves unrestricted quantification over possible worlds, it falsely predicts that plurals can range only over pluralities of pICs that *necessarily* yield distinct values. It seems that the relevant property should not be a total lack of overlap between the two pICs, but a lack of overlap relative to the subjects' belief states. As an illustration, consider the following German example from Haslinger and Schmitt (to appear)[6]:

(13) SHAPE-SHIFTING: [Roy's castle, no monsters ...] Abe and Bert believe in ghosts and think ghosts cannot change their appearance. At 1am, Abe thinks he saw a tall, red-haired ghost. At 2am, Bert thinks he saw a short, black-haired ghost. They tell Roy about their beliefs. Roy isn't sure whether ghosts exist, but he is convinced that *if* ghosts exist, they can shape-shift. Roy says:
Abe und Bert glauben, dass zwei Geister im Schloss waren. *Aber selbst wenn sie wirklich jeder einem Geist begegnet sind, war es wahrscheinlich ein und derselbe.*
'Abe and Bert think there were two ghosts at the castle. But even if they really each encountered a ghost, it was probably one and the same.'

To us, Roy's utterance in (13) is felicitous and correctly reports Abe's and Bert's beliefs. This poses a problem for the unrestricted quantification in (9): Given the possibility of shape-shifting ghosts, there may be worlds in which Abe's tall, red-haired ghost and Bert's short, black-haired ghost are one and the same. In fact, we must allow such worlds in our model to account for the second half of Roy's utterance. So the most natural pICs in (13) do not meet the condition in (9). Accordingly, none of the obvious notions of distinctness for pICs are empirically adequate.

The reason why (13) is still acceptable seems to be that Abe's and Bert's beliefs rule out the existence of shape-shifting monsters: It follows from each of their belief states that a tall, red-haired ghost must be distinct from a short, black-haired ghost. (13) thus suggests a notion of

[6]Thanks to Magdalena Kaufmann, Sarah Zobel and a reviewer for the suggestion that we consider examples where the attitude subjects and the speaker disagree on the relevant individuation criteria. While our analysis of cumulative belief will not appeal to the *speaker's* individuation criteria, we do not want to rule out the possibility that plural semantics could be sensitive both to subject-relative and to speaker-relative individuation criteria.

distinctness that is relativized to the attitude subjects' individuation criteria. If we can attribute to Abe and Bert the plausible belief that a zombie must be distinct from a griffin, the acceptability of (1a)/(1b) in the ZOMBIE VS. GRIFFIN scenario would still be accounted for. In the UNSPECIFIED MONSTERS scenario, on the other hand, the pICs in (8) would count as distinct only if Abe and Bert both believe that a monster appearing in Abe's room must be distinct from a monster appearing in Bert's room (possibly at a different time). As the scenario does not provide any reason to ascribe this odd belief to them, (1a)/(1b) are still not predicted to be true.

Accordingly, to develop a more plausible notion of distinctness for pICs, we have to consider the subjects' own individuation criteria, which are part of their belief states. But to implement this idea, we must go beyond the common practice of representing a belief state as a mere set of possible worlds (cf. Hintikka 1969). This is because, informally speaking, whether or not a subject takes two pICs to be distinct doesn't seem to depend on whether the subject takes these pICs to actually be instantiated. To see the point, consider (14), a variant of our ZOMBIE VS. GRIFFIN scenario. In (14), neither subject believes in both kinds of monsters; yet, they may still have individuation criteria that tell them a zombie must be distinct from a griffin.

(14) ZOMBIE VS. GRIFFIN (INCOMPATIBLE): [Roy's castle, no monsters ...] Abe believes in zombies but believes that griffins don't exist. Bert believes in griffins but believes that zombies don't exist. At 1am, Abe thinks there is a zombie in his room. At 2am, Bert thinks there is a griffin on his bed. [No discussion ...] (1a), (1b) **%true**

Crucially, this variation does not reverse the judgments – the sentences are still true (for speakers who accept cumulative belief sentences to begin with). This suggests that, informally speaking, a subject can 'compare' two pICs in terms of distinctness even if they don't believe both pICs to be instantiated. To make sense of this intuition, we submit that judgments of distinctness can be thought of as involving (potentially) *counterfactual* beliefs. In Haslinger and Schmitt (to appear), we therefore propose the distinctness constraint stated informally in (15):

(15) Relative to two attitude subjects **a** and **b**, two individual concepts f, g count as distinct iff **a** and **b** both believe the counterfactual *If both f and g were instantiated, their values would be distinct* and the domains of f and g overlap.

Given (15), we can now give a first informal rendering of the truth-conditions associated with the cumulative reading of (1a)/(1b), again following Haslinger and Schmitt (to appear):

(16) $[\![(1a)]\!]^w = [\![(1b)]\!]^w = 1$ iff there is a plurality $f + g$ of two **monster**-concepts such that
   a. Abe and Bert each believe at least one of the propositions $[\lambda w.\mathbf{roam\text{-}the\text{-}castle}_w(f(w))]$ and $[\lambda w.\mathbf{roam\text{-}the\text{-}castle}_w(g(w))]$
   b. and each of these two propositions is believed by at least one of Abe and Bert
   c. and Abe believes: If both f and g were instantiated, their values would be distinct
   d. and Bert believes: If both f and g were instantiated, their values would be distinct

(16) raises two questions that we will now address in turn: First, it seems to be in conflict with the expectation that the truth conditions of (1a)/(1b) somehow involve the proposition [$\lambda w.$ two monsters are roaming the castle in w]. Rather, (16) requires Abe and Bert to stand in a cumulative belief relation to what could be described as 'parts' of this proposition. So our first question is: How do we obtain such propositional 'parts', how are they connected to the different **monster**-concepts and how can the notion of cumulative belief be spelled out?

This issue will be addressed in Section 3, where we summarize Schmitt's (2020) account of cumulative belief within the 'Plural Projection' theory of cumulativity, which lets us derive propositional pluralities from pluralities of pICs. The second notion in need of clarification in (16) is counterfactual belief, which seems to be beyond the scope of the standard possible-worlds treatment of belief (Hintikka, 1969). This will be the topic of Section 4, where we give a simple analysis of counterfactual belief based on subject-dependent ordered sets of worlds.

3. Cumulative belief in the Plural Projection framework

This section lays some groundwork for our analysis of cumulative belief attributions involving 'belief objects' or, on our analysis, pICs. We will first give some less complex examples of cumulative belief that do not involve pICs and then use these examples to introduce the Plural Projection analysis of cumulativity developed in Haslinger and Schmitt (2018); Schmitt (2019, 2020). Readers familiar with any of these works can skip Section 3.2.

3.1. Simpler cases of cumulative belief

Let us first look at other cases of cumulative belief. Recall that cumulativity is a hallmark of plurality: The weak truth conditions of (4) above are tied to the presence of two semantically plural expressions. Now consider the German sentence (17a), which is true in scenario (17b). While it is not obvious under which conditions someone counts as having two beliefs about the World Cup rather than one[7], it seems clear that neither Abe nor Bert individually satisfies these conditions in scenario (17b). (17) thus appears to involve a cumulative reading of the plurals *Abe und Bert* and *zwei Sachen* 'two things'. This reading is easily accounted for if *zwei Sachen* is taken to range over pluralities of belief contents. But this entails that the notion of plurality, usually restricted to primitives like individuals or events, should be extended to propositions.

(17) a. *Was das WM-Finale betrifft, glauben Abe und Bert zwei Sachen.*
what the World.Cup-final concerns believe Abe and Bert two things
'Concerning the World Cup final, Abe and Bert believe two things.'

b. SCENARIO Soccer experts Abe and Bert are asked who will be in the World Cup final. Abe believes Germany will be in the final. Bert believes Brazil will be in the final. They have no other relevant beliefs. (17a) **true**

The same point can be made without relying on semantically exceptional nouns like *Sache* 'thing'. As discussed by Schmitt (2019, 2020), *and*-conjunctions of any category permit cumulative readings, including sentential conjunctions in the object position of *believe*. This is illustrated by (18b), which is true in scenario (18a).

(18) a. SCENARIO: A liberal newspaper is hiring two editors. Carl and Dean have the best CVs, but the HR people, Abe and Bert, have looked them up on social media. Abe thinks Carl is a Trump supporter and a 9/11 truther. Bert thinks Dean believes in lizard people. They tell their boss about their findings. She tells her colleague:

b. *I really don't think we should hire Carl and Dean. Abe and Bert believe* [[$_p$ ***that they are conspiracy theorists***] *and* [$_q$ ***(that) Carl is a Trump supporter***]].

The sentential conjunction in (18b) has a cumulative reading: Neither Abe nor Bert believes

[7]There are no obvious criteria that determine what counts as 'one' or 'two' beliefs (see e.g. Quine 1960: §44). Haslinger (2020) argues that this is because the quantificational domain of such DPs is context-dependent, more specifically QUD-dependent. This is why we explicitly introduce a discourse topic and set up a QUD in (17).

both of the conjuncts p and q. But as scenario (18a) also forces a cumulative reading of the pronoun *they* in the first conjunct relative to *Abe and Bert*, the truth conditions do not even require each of Abe and Bert to believe at least one of p and q. In other words, the plural expression *they* in (18b) is embedded in the sentential conjunction *p and q*, which itself receives a plural interpretation, and both expressions participate in a cumulative reading relative to a higher plural. Schmitt (2019) argues that such configurations are beyond the scope of analyses where cumulativity is a property of relation-denoting expressions (e.g. Beck and Sauerland 2000). What we need instead to account for (18b) is a theory that 1) builds cumulativity into the compositional mechanism, rather than the meanings of predicates, and 2) admits pluralities of higher-type objects like propositions. Within such a theory, the complement clause in (18b) can be analyzed as denoting a plurality of three propositions: p_1 = **that Carl is a conspiracy theorist**, p_2 = **that Dean is a conspiracy theorist** and q = **that Carl is a Trump supporter**. The truth conditions of (18b) can then be expressed via a cumulative relation between the individual plurality $\mathbf{a}+\mathbf{b}$ and this propositional plurality: Each of **a** and **b** must believe at least one of p_1, p_2 and q, and each of p_1, p_2 and q must be believed by at least one of **a** and **b**.

In summary, both the use of pluralities of higher-type objects (like propositions) and the notion of cumulative belief can be motivated by examples that do not involve puzzles about identity or distinctness conditions for belief objects. We will thus start by developing an analysis of these simpler cumulative belief sentences, and then combine it with an analysis of indefinite plurals that uses pICs and incorporates the counterfactual distinctness constraint in (16).

3.2. Plural Projection

We will now sketch a stripped-down version of the compositional semantics for cumulativity proposed by Schmitt (2019, 2020), using (18b) as our motivating example. Schmitt (2019, 2020) starts off by generalizing the notion of plurality across semantic categories. Every semantic domain – including those for complex types – is enriched with a set of pluralities that stand in a one-to-one correspondence with nonempty sets of the usual domain elements (19).

(19) a. For any type a we have a set A_a of 'atomic' meanings. The full domain D_a is a set isomorphic to, but disjoint from, the set of nonempty subsets of A_a, with a sum operation $+_a$ corresponding to the union of nonempty subsets of A_a.
b. For a functional type $\langle a,b\rangle$, $A_{\langle a,b\rangle}$ is the set of all partial functions from D_a to D_b.

So besides plural individuals like $\mathbf{Abe}+\mathbf{Bert}$ in D_e, we have pluralities of propositions in $D_{\langle s,t\rangle}$: For instance, as $D_{\langle s,t\rangle}$ contains the propositions p_1, p_2 and q from example (18b), it also contains their sum p_1+p_2+q, from which the three atomic parts p_1, p_2 and q can be recovered. A further example directly relevant to our purposes is that, given (19), $D_{\langle s,e\rangle}$ contains arbitrary pluralities of pICs. Hence $D_{\langle s,e\rangle}$ contains the pICs f and g from (8), but also their sum $f+g$.

The second nonstandard aspect of this system is that the compositional mechanism relies on alternative sets containing pluralities, so-called **plural sets**. This feature is essentially an adaptation of Alternative Semantics approaches to wh-questions (Hamblin, 1973) and indefinites (Kratzer and Shimoyama, 2002). For example, we argued that the plural indefinite *zwei Sachen* 'two things' in (17a) ranges over contextually provided pluralities of propositions. Given the 'atomic' propositions p, q and r, it thus denotes the set of pluralities from $D_{\langle s,t\rangle}$ in (20):

(20) $[\![\textit{zwei Sachen}]\!] = \{p+q, q+r, p+r\}$

For non-plural expressions, such as proper names or simple predicates, the plural sets will be singletons (21a). The semantic contribution of *and* is to combine two plural sets by forming all sums of an element of the first set and an element of the second set, as illustrated in (21b). Thus, *and* expresses a plurality-forming operation for conjuncts of any category.

(21) a. $[\![\textit{Carl}]\!] = \{\mathbf{c}\}$; $[\![\textit{Dean}]\!] = \{\mathbf{d}\}$; $[\![\textit{conspiracy theorist}]\!] = \{\mathbf{C}\}$
b. $[\![\textit{Carl and Dean}]\!] = \{\mathbf{c}+\mathbf{d}\}$

A full compositional analysis of (18b) requires a general way of extending composition rules, such as functional application, to these plural sets in a way that encodes cumulativity. Here, we give a rough sketch of how this is done and refer the reader to Haslinger and Schmitt (2018); Schmitt (2020) for the formal details. Let's assume we want to combine two plural sets S_a and S_b of types a and b, respectively, and that we already have a composition rule that combines 'atomic' (non-plural) denotations of type a with those of type b. We then introduce the notion of a **compositional cover** of a plurality in S_a and a plurality in S_b, which is defined as follows:

(22) A **compositional cover** of two pluralities x and y is a relation $C \subseteq \{x' \mid x' \text{ is an atomic part of } x\} \times \{y' \mid y' \text{ is an atomic part of } y\}$ in which every atomic part of x and every atomic part of y appears at least once.

For instance, if S_a is the singleton set $\{\mathbf{C}\}$ from (21a) and S_b the singleton set $\{\mathbf{c}+\mathbf{d}\}$ from (21b), there is only one compositional cover of a plurality from S_a and a plurality from S_b:[8]

(23) $\{\langle \mathbf{C},\mathbf{c}\rangle, \langle \mathbf{C},\mathbf{d}\rangle\}$

The two sets S_a and S_b are then combined as follows: For any compositional cover of a plurality in S_a and a plurality in S_b, we apply the regular composition rule for types a and b to each pair in the cover and form a plurality from the resulting values. The pluralities corresponding to the different covers are then collected into a new plural set. Consider the output of this rule for the first conjunct of the embedded clause in (18b): Let us assume that, in the given context, *they* denotes the plural set in (21b). The only compositional cover of this set and the set $\{\mathbf{C}\}$ for the predicate is the one in (23). Assuming for simplicity that $\mathbf{C}$ is of type $\langle e, \langle s,t\rangle\rangle$, the next step is to perform regular functional application for each pair in the cover and sum up the results. We obtain the plurality $\mathbf{C}(\mathbf{c}) + \mathbf{C}(\mathbf{d})$, which consists of two 'atomic' propositions:

(24) $[\![\textit{they are conspiracy theorists}]\!] = \{\mathbf{C}(\mathbf{c}) + \mathbf{C}(\mathbf{d})\}$

The second embedded conjunct, *Carl is a Trump supporter*, does not contain any plurals and thus denotes a singleton set with an atomic proposition, $\{\mathbf{T}(\mathbf{c})\}$. Now we combine the plural sets for the two conjuncts via $[\![\textit{and}]\!]$. Recall that $[\![\textit{and}]\!]$ takes two plural sets by forming all sums of an element of the first set and an element of the second set. So we end up with the denotation we wanted for the clausal conjunction, a set containing a plurality of three propositions:

(25) $[\![\textit{they are conspiracy theorists and Carl is a Trump supporter}]\!] = \{\mathbf{C}(\mathbf{c}) + \mathbf{C}(\mathbf{d}) + \mathbf{T}(\mathbf{c})\}$

The next step is to compose (25) with the matrix verb *believe*. For the moment, we assume that *believe* has its standard Hintikka-style meaning, lifted to a singleton plural set:

(26) $[\![\textit{believe}]\!] = \{\lambda p_{\langle s,t\rangle}.\lambda x.\lambda w.\forall w' \in \text{DOX}_w(x).p(w')\}$

[8]A relation like $\{\langle \mathbf{C},\mathbf{c}\rangle\}$ would not count as a compositional cover, as it does not use all the atomic parts of $\mathbf{c}+\mathbf{d}$.

Applying our composition principle to (25) and (26), there will again be only one compositional cover, which matches the standard denotation of *believe* with each of the three atomic propositions in the plurality in (25). We thus end up with the following plural set:

(27) $[\![$*believe that they are conspiracy theorists and that Carl is a Trump supporter*$]\!]$
$= \{[\lambda x.\lambda w.\forall w' \in \textsc{dox}_w(x).\mathbf{C}(\mathbf{c})(w')] + [\lambda x.\lambda w.\forall w' \in \textsc{dox}_w(x).\mathbf{C}(\mathbf{d})(w')]$
$+ [\lambda x.\lambda w.\forall w' \in \textsc{dox}_w(x).\mathbf{T}(\mathbf{c})(w')]\}$

This set contains a sum of three properties of individuals: the property of believing Carl is a conspiracy theorist, that of believing Dean is a conspiracy theorist and that of believing Carl is a Trump supporter. The last step is to combine this with the plural set $\{\mathbf{a}+\mathbf{b}\}$ for *Abe and Bert*. Here, the notion of a compositional cover is non-trivial: As both sets contain proper pluralities, we can match Abe and Bert with the propositions in (27) in multiple ways. (28) gives some of the compositional covers. For each compositional cover of the plurality in (27) and the plurality $\mathbf{a}+\mathbf{b}$, we now perform functional application for each pair in the cover and sum up the resulting propositions. Thus, the covers in (28a) and (28b) yield the sums of propositions in (29).

(28) a. $\{\langle[\lambda x.\lambda w.\forall w' \in \textsc{dox}_w(x).\mathbf{C}(\mathbf{c})(w')], \mathbf{a}\rangle, \langle[\lambda x.\lambda w.\forall w' \in \textsc{dox}_w(x).\mathbf{C}(\mathbf{d})(w')], \mathbf{b}\rangle,$
$\langle[\lambda x.\lambda w.\forall w' \in \textsc{dox}_w(x).\mathbf{T}(\mathbf{c})(w')], \mathbf{a}\rangle\}$
b. $\{\langle[\lambda x.\lambda w.\forall w' \in \textsc{dox}_w(x).\mathbf{C}(\mathbf{c})(w')], \mathbf{b}\rangle, \langle[\lambda x.\lambda w.\forall w' \in \textsc{dox}_w(x).\mathbf{C}(\mathbf{d})(w')], \mathbf{b}\rangle,$
$\langle[\lambda x.\lambda w.\forall w' \in \textsc{dox}_w(x).\mathbf{T}(\mathbf{c})(w')], \mathbf{a}\rangle\}$

(29) a. $[\lambda w.\forall w' \in \textsc{dox}_w(\mathbf{a}).\mathbf{C}(\mathbf{c})(w')] + [\lambda w.\forall w' \in \textsc{dox}_w(\mathbf{b}).\mathbf{C}(\mathbf{d})(w')]$
$+ [\lambda w.\forall w' \in \textsc{dox}_w(\mathbf{a}).\mathbf{T}(\mathbf{c})(w')]$
b. $[\lambda w.\forall w' \in \textsc{dox}_w(\mathbf{b}).\mathbf{C}(\mathbf{c})(w')] + [\lambda w.\forall w' \in \textsc{dox}_w(\mathbf{b}).\mathbf{C}(\mathbf{d})(w')]$
$+ [\lambda w.\forall w' \in \textsc{dox}_w(\mathbf{a}).\mathbf{T}(\mathbf{c})(w')]$

There are many other compositional covers besides those in (28). Each cover, and each propositional plurality obtained from a cover, corresponds to a particular cumulative scenario for (18b); e.g., (29a) characterizes scenario (18a), while (29b) corresponds to a scenario where Bert believes Carl and Dean are conspiracy theorists and Abe believes Carl is a Trump supporter. The final step in our derivation of (18b) and other plural sentences is to define how a truth value is assigned to a plural set of propositions in a world: (30) uses existential quantification over the set, and universal quantification over the atomic parts of each plurality in it:

(30) A plural set $S \subseteq D_{\langle s,t\rangle}$ is **true** in a world w iff there is at least one propositional plurality in S all atomic parts of which are true in w.

Our next goal is to incorporate plural indefinites like *two monsters* in (1a) into this framework in a way that accounts for the distinctness restriction described in Section 2. Before we do so, let us highlight an important property of the present analysis of cumulative belief sentences: It does not rely on a lexical treatment of *believe* specifically designed for plural subjects. Rather, a highly general approach to cumulativity that effectively builds it into the composition rules is combined with the standard Hintikka (1969) denotation of *believe*. As a consequence, our distinctness constraint from Section 2 cannot be specific to cumulative belief sentences either. Instead, we will argue that this constraint is present in the semantics of all plural indefinites.

4. Encoding distinctness: A counterfactual component for plurals

As shown in Section 3, the Plural Projection system takes plural indefinites to introduce sets of pluralities. In Section 2 we further assumed that plural indefinites range over pICs rather than actual individuals in order to account for their behavior in cumulative belief sentences. Combining these two ideas, we obtain the naive semantics for plural indefinites in (31).

(31) $[\![$*two monsters*$]\!] = \{f+g \mid f,g \in A_{\langle s,e\rangle} \wedge f,g$ are $[\![$*monster*$]\!]$-concepts $\wedge f \neq g\}$

(31) obviously does not capture the distinctness constraint motivated in Section 2. For example, it fails to exclude the sum of the pICs f and g in (8), which are intuitively not 'distinct enough' to license the numeral *two*. So it seems we must somehow supplement (31) with our account of intentional distinctness, repeated in (32), to overcome this problem.

(32) Relative to two attitude subjects **a** and **b**, two individual concepts f,g count as distinct iff **a** and **b** both believe the counterfactual *If both f and g were instantiated, their values would be distinct* and the domains of f and g overlap.

4.1. A problem with compositionality

It is, however, unclear how to add (32) to our treatment of indefinites without running into a compositionality problem, which has its roots in the fact that (32) appeals to a counterfactual belief. Following the intuition behind the 'ordering semantics' of Lewis (1973), we take this counterfactual to involve reference to a set of 'closest' worlds where both f and g are instantiated. What does 'closest worlds' mean in a belief context? For one, these 'closest' worlds will not necessarily be among the attitude subjects' doxastic alternatives: As the ZOMBIE VS. GRIFFIN (INCOMPATIBLE) scenario in (14) shows, cumulative belief sentences do not require both subjects to believe that all the relevant pICs are instantiated. Crucially, however, the SHAPE-SHIFTING scenario in (13) suggests that the subjects' beliefs do play a role in determining which worlds count as 'closest'. Thus, to determine the right set of worlds to evaluate the consequent of the counterfactual, we need access to each subject's belief state as a whole. This 'access to the entire belief state' is at odds with the standard possible-worlds treatment of *believe* (Hintikka, 1969), in which the complement is evaluated separately in each of the attitude subject's doxastic alternatives, so that the semantics of expressions within the complement cannot make use of global properties of the subject's belief state. So the problem is that operators within the complement of *believe* are standardly assumed to be unable to access the subject's full doxastic state, yet this seems to be exactly what we need to make sense of (32).

4.2. Tackling the compositionality problem

This type of 'lookahead' problem is not specific to cumulative belief attributions.[9] For example, Yalcin (2007) argues that to account for the behavior of epistemic modals like *might* in embedded contexts, we must relativize the truth value of a sentence not just to a world, but also to a set of worlds – a 'domain' – representing an epistemic state. More specifically, he relativizes the extensions of natural language expressions to complex indices, which are pairs consisting of a world and a domain. In addition to quantifying over the world parameter, certain embedding operators shift the value of the domain parameter. Modal expressions embedded under such operators then have access to global properties of the relevant domain.

This independently motivated notion of complex indices will help solve the compositionality problem raised by the distinctness constraint (32). But Yalcin's view of epistemic states as sets

[9]See Schulz 2007 for a detailed investigation of the role of 'global' operations on an epistemic state in the semantics of counterfactuals. We cannot incorporate her insights here for reasons of space.

of worlds is insufficient to account for the (potentially) counterfactual nature of (32) – we must consider worlds that might not be among either subject's doxastic alternatives. We thus combine his insight with an ordering semantics for counterfactuals, in the sense of Lewis 1973, and model indices as pairs $i = \langle w_i, \preceq_i \rangle$, with w_i a world and $\preceq_i$ a partial ordering among worlds. Unlike Lewis, however, we take the ordering $\preceq_i$ to represent an epistemic state, rather than the facts and generalizations holding in a particular world. The belief state of an individual x in a world w can thus be represented as an ordering $\preceq_{w,x}$. As x's belief state does not identify a unique world that x locates herself in, we will not assume that this ordering has a unique minimal element. Rather, each of x's doxastic alternatives will be a minimal element of $\preceq_{w,x}$. The ordering $\preceq_{w,x}$ therefore encodes both x's set of doxastic alternatives (in w) and the information which of the worlds ruled out by x's beliefs are 'closer' to being doxastic alternatives than others. This 'closeness' relation, which could be conceived of as encoding a set of law-like generalizations x believes to hold (cf. Schulz 2007), underlies x's counterfactual beliefs.

Before we return to plural indefinites, we will illustrate the formal details of this approach using overt counterfactuals embedded under *believe* as in (33).

(33) *Abe believes that [$_p$ if there had been a monster], [$_q$ there would have been a noise].*

We start with the counterfactual. (34) encodes the idea that counterfactuals quantify over a set of worlds determined exclusively by the ordering $\preceq_i$ of the evaluation index i; the world component of i is ignored. The counterfactual selects the minimal elements relative to $\preceq_i$ among the worlds making the antecedent true, and requires the consequent to be true at all indices consisting of one of these worlds and $\preceq_i$. These do not have to be minimal elements of $\preceq_i$ in the global sense, which is crucial for the account of counterfactual belief attributions.

(34) $[\![\textit{if } p, q]\!] =$
$$\lambda i.\forall w[[w \in \mathrm{dom}(\preceq_i) \wedge [\![p]\!](\langle w, \preceq_i \rangle) \wedge \neg\exists w'[w' \prec_i w \wedge [\![p]\!](\langle w', \preceq_i \rangle)]] \rightarrow [\![q]\!](\langle w, \preceq_i \rangle)]$$

Next, we revise the semantics of *believe*. The lexical entry in (35) quantifies over all worlds that are minimal within the ordering $\preceq_{w_i,x}$ that represents the subject x's belief state in the world w_i provided by the evaluation index i. In the case of a simple complement without modal expressions, as in (36), it replicates the results of the Hintikka semantics, as the semantics of such complements does not depend on the ordering component of an index.[10]

(35) $[\![\textit{believe}]\!] =$
$$\lambda i.\lambda p_{\langle s,t \rangle}.\lambda x_e.\forall w'[w' \in \mathrm{dom}(\preceq_{w_i,x}) \wedge \neg\exists w''[w'' \prec_{w_i,x} w'] \rightarrow p(\langle w', \preceq_{w_i,x} \rangle)]$$

(36) $[\![\textit{Abe believes there was a monster}]\!]$
$$= \lambda i.\forall w'[w' \in \mathrm{dom}(\preceq_{w_i,\mathbf{a}}) \wedge \neg\exists w''[w'' \prec_{w_i,\mathbf{a}} w'] \rightarrow [\![\textit{there was a monster}]\!](\langle w', \preceq_{w_i,\mathbf{a}} \rangle)]$$
$$= \lambda i.\forall w'[w' \in \mathrm{dom}(\preceq_{w_i,\mathbf{a}}) \wedge \neg\exists w''[w'' \prec_{w_i,\mathbf{a}} w'] \rightarrow \exists y.\mathbf{monster}(w')(y)]$$

The crucial difference from the Hintikka semantics, however, is that the belief content p is evaluated at complex indices consisting of one such world and the entire ordering $\preceq_{w_i,x}$: Attitude verbs do not only shift the world component, but also the ordering component of an index. This makes it possible for material in the complement of attitude verbs to access global properties of the belief state encoded in $\preceq_{w_i,x}$. This is exactly what happens in the case of counterfactuals

[10]We assume that the meanings of expressions of lexical categories like *monster* never depend on the ordering component of an index, so we use metalanguage predicates like **monster** with a simple world argument.

under attitudes like (33), as shown in (37): The second step of (37) states that the counterfactual must be true at all indices consisting of the ordering that encodes Abe's belief state, $\preceq_{w_i,\mathbf{a}}$, and a minimal world in that ordering. But crucially our semantics for counterfactuals in (34) only cares about the ordering component of an index; the quantification over Abe's belief worlds is thus vacuous. As shown in the third and fourth step of (37), the counterfactual considers the minimal worlds relative to $\preceq_{w_i,\mathbf{a}}$ at which there was a monster, and requires there to be a noise in all of these worlds. These minimal worlds where there was a monster do not have to be minimal elements of $\preceq_{w_i,\mathbf{a}}$ in the global sense, so they do not have to be among Abe's belief worlds (in fact they won't be, if we factor in the presupposition of the counterfactual). But Abe's beliefs still matter for selecting the right set of worlds to quantify over.

(37) $[\![\textit{(33)}]\!] = \lambda i.\forall w'[w' \in \mathrm{dom}(\preceq_{w_i,\mathbf{a}}) \wedge \neg\exists w''[w'' \prec_{w_i,\mathbf{a}} w']$
$\rightarrow [\![\textit{if there had been a monster, there would have been a noise}]\!](\langle w', \preceq_{w_i,\mathbf{a}}\rangle)]$
$= \lambda i.\forall w[w \in \mathrm{dom}(\preceq_{w_i,\mathbf{a}}) \wedge [\![\textit{there was a monster}]\!](\langle w, \preceq_{w_i,\mathbf{a}}\rangle) \wedge \neg\exists w'[w' \prec_{w_i,\mathbf{a}} w \wedge$
$[\![\textit{there was a monster}]\!](\langle w', \preceq_{w_i,\mathbf{a}}\rangle)] \rightarrow [\![\textit{there was a noise}]\!](\langle w, \preceq_{w_i,\mathbf{a}}\rangle)]$
$= \lambda i.\forall w[w \in \mathrm{dom}(\preceq_{w_i,\mathbf{a}}) \wedge \exists y.\mathbf{monster}(w)(y) \wedge \neg\exists w'[w' \prec_{w_i,\mathbf{a}} w \wedge \exists y.\mathbf{monster}(w')(y)] \rightarrow$
$\exists z.\mathbf{noise}(w)(z)]$

4.3. Back to distinctness

Given these assumptions, we return to the question of how the distinctness constraint can be incorporated into the meaning of plural indefinites. Our explication of distinctness in (32) uses the attitude subjects' belief states to select the 'closest' worlds for each subject at which both of the relevant pICs are instantiated, even if these are not among the subjects' belief worlds. At the complex indices corresponding to these worlds, the values of the pICs must be distinct. In our formalism, this means that for each attitude subject x, we must consider the doxastic ordering $\preceq_{w,x}$ in the evaluation world w, and check if the relevant concepts have distinct values at indices corresponding to the minimal worlds in $\preceq_{w,x}$ at which they are both instantiated. Now one of the semantic effects we ascribed to *believe* in (35) is to evaluate the complement clause – and any material within it – at indices with the ordering component shifted to $\preceq_{w,x}$. This means we can now build this condition into the meaning of the indefinite without facing a compositionality problem: If the indefinite is evaluated at such complex indices, the ordering component of the index gives it access to the subject's entire belief state.

To encode this idea, we first formalize our notion of distinctness relative to an ordering $\preceq$ in the way just described (38): Two pICs f and g are distinct iff at each of the indices corresponding to a minimal world in $\preceq$ where both f and g are defined, they have distinct values.

(38) For any partial ordering $\preceq$ on worlds and $f, g \in D_{\langle s,e\rangle}$: $\textsc{distinct}_{\preceq}(f,g)$ holds iff
$\forall w[[w \in \mathrm{dom}(\preceq) \wedge \langle w, \preceq\rangle \in \mathrm{dom}(f) \cap \mathrm{dom}(g) \wedge \neg\exists w'[w' \prec w \wedge \langle w', \preceq\rangle \in \mathrm{dom}(f) \cap$
$\mathrm{dom}(g)]] \rightarrow f(\langle w, \preceq\rangle) \neq g(\langle w, \preceq\rangle)]$

The idea is to add such a distinctness condition to each part of each plurality of pICs introduced by the indefinite. The simplest implementation would be to make the distinctness condition a presupposition, but then we would have to explore how presupposition failures interact with the Plural Projection system, which is orthogonal to our purposes. We would also face empirical issues: If Abe does not believe two monsters exist, a (non-cumulative) belief ascription with *two monsters* in the belief context should come out as false, not as a presupposition failure.

We thus take a more complex route, encoding the condition in the *assertive* meaning component of each atomic part of each plurality in $[\![$*two monsters*$]\!]$. As this is impossible if these parts are of type $\langle s, e\rangle$, we type-lift them as in (39) to combine with a property argument P of type $\langle s, \langle e, t\rangle\rangle$. Take a plurality of two pICs $f + g$: The lifted version of f is a function that takes a property argument P and maps it to true only if a) f is instantiated at the evaluation index i, b) f and g are distinct at i (in the sense of (38)) and c) the value of f at i satisfies P at i. The lifted version of g is analogous, so it also encodes the condition that f and g count as distinct at i. Each plurality $f + g$ of pICs is thus mapped to a plurality of generalized quantifiers, each of which mimics the behavior of one of the pICs, but also requires f and g to count as distinct relative to the given ordering (which in a belief context encodes a subject's belief state).

(39) $[\![\textit{two monsters}]\!] = \{(\lambda i.\lambda P_{\langle s,et\rangle}.i \in \mathrm{dom}(f) \wedge \textsc{distinct}_{\preceq_i}(f,g) \wedge P(i)(f(i)))$
$+(\lambda i.\lambda P_{\langle s,et\rangle}.i \in \mathrm{dom}(g) \wedge \textsc{distinct}_{\preceq_i}(f,g) \wedge P(i)(g(i)))$
$\mid f, g \in A_{\langle s,e\rangle} \wedge f, g$ are $[\![\textit{monster}]\!]$-concepts $\wedge f \neq g\}$

Our new DP-denotation follows the general 'format' of plural indefinites within the Plural Projection system but also encodes distinctness: If one of the quantifiers in a plurality in the set (39) is evaluated at an index i with an ordering that does not meet the distinctness constraint for f and g, the proposition derived by applying it to a property P will be false at i, no matter what P is. While we cannot give the internal composition for plural indefinites here, it seems plausible to assume that our distinctness condition is contributed by the plural morpheme. This would mean that the contribution of nominal plural morphology includes not only the extensional operation of closing a predicate under sum, but also a non-trivial intensional condition.

5. Deriving the truth conditions of cumulative belief sentences

We now have the ingredients needed to analyze our sentences in (1a) (= (40)) and (1b): The semantics for plural indefinites in (39) with the counterfactual distinctness condition, and the compositional mechanism for cumulative sentences within the Plural Projection system. So: How do we derive a plural set of propositions as the denotation for the complement clause via our new DP-meaning, and how does this plural set interact with our new semantics for *believe*?

(40) *Abe and Bert believe that two monsters were roaming the castle.*

We first consider the embedded clause of (40). Since the VP does not contain any plurals, it denotes a plural set with a single atomic element, given in (41). The DP denotes the set in (39) above, which contains all quantifier pluralities derived by type-lifting two partial **monster**-concepts f and g to take a property argument and conjoining each of the resulting quantifiers with a distinctness condition. Given an index i, the distinctness condition states that in the closest worlds according to the ordering $\preceq_i$ where both f and g yield values, these values are different. For each plurality of **monster**-concepts f and g, we obtain only one compositional cover when matching the corresponding quantifier plurality with the VP-denotation: $\{\langle \mathbf{roam}, (\lambda i.\lambda P_{\langle s,et\rangle}.i \in \mathrm{dom}(f) \wedge \textsc{distinct}_{\preceq_i}(f,g) \wedge P(i)(f(i)))\rangle, \langle \mathbf{roam}, (\lambda i.\lambda P_{\langle s,et\rangle}.i \in \mathrm{dom}(g) \wedge \textsc{distinct}_{\preceq_i}(f,g) \wedge P(i)(g(i)))\rangle\}$. Our composition principle from Section 3 then requires us to combine the two components of each pair in the cover in a type-driven manner. In the present case, the composition rule we need to apply for each pair is intensional functional application. After applying this rule for each functor-argument pair in the cover, we sum up the results for each cover and collect the resulting sums in one big plural set. This gives us the plural set displayed at the top node in (41): Each plurality of propositions in this set has atomic parts that reflect the parts of one of our original pluralities of **monster**-concepts. Note further that each part of a propositional plurality encodes the distinctness constraint for the corresponding pICs.

(41)

$$\begin{array}{c}\{(\lambda i.i \in \mathrm{dom}(f) \wedge \textsc{distinct}_{\preceq_i}(f,g) \wedge \mathbf{roam}(i)(f(i)))\\ +(\lambda i.i \in \mathrm{dom}(g) \wedge \textsc{distinct}_{\preceq_i}(f,g) \wedge \mathbf{roam}(i)(g(i)))\\ \mid f,g \in A_{\langle s,e\rangle} \wedge f,g \text{ are } [\![\mathit{monster}]\!]\text{-concepts } \wedge f \neq g\}\end{array}$$

$$\begin{array}{c}\{(\lambda i.\lambda P_{\langle s,et\rangle}.i \in \mathrm{dom}(f) \wedge \textsc{distinct}_{\preceq_i}(f,g) \wedge P(i)(f(i)))\\ +(\lambda i.\lambda P_{\langle s,et\rangle}.i \in \mathrm{dom}(g) \wedge \textsc{distinct}_{\preceq_i}(f,g) \wedge P(i)(g(i)))\\ \mid f,g \in A_{\langle s,e\rangle} \wedge f,g \text{ are } [\![\mathit{monster}]\!]\text{-concepts } \wedge f \neq g\}\end{array} \qquad \{\mathbf{roam}\}$$

two monsters — *roaming the castle*

Next we compose this plural set with $[\![\mathit{believe}]\!]$. Recall that our revised semantics for *believe* in Section 4 led to the function **believe** defined in (35) which, unlike the standard analysis, quantifies over complex indices consisting of the subject's ordering relation and one of their belief worlds. Given the basic assumptions of the Plural Projection framework, *believe* denotes a singleton set containing the function in (35). So for each propositional plurality $p+q$ in the plural set in (41), we obtain exactly one compositional cover, $\{\langle \mathbf{believe}, p\rangle, \langle \mathbf{believe}, q\rangle\}$. Following our composition principle for plural sets, we once more perform intensional functional application for each pair in the cover, sum up the values per cover, and collect the results in a plural set. We end up with the plural set of predicates at the top node in (42). The atomic parts of the predicate pluralities in this set reflect the atomic parts of the initial pluralities of **monster**-concepts $f+g$. The predicate corresponding to the pIC f imposes the following conditions on its subject x: Each index $\langle w', \preceq_{w_i,x}\rangle$ corresponding to the subject's belief state $\preceq_{w_i,x}$ and one of her belief worlds w' is such that f is instantiated at that index, and the value of f at that index is roaming the castle. Analogous conditions have to hold for the predicate corresponding to g. Importantly, in addition, both atomic parts of the predicate plurality still carry the distinctness condition – and this condition is evaluated relative to the ordering component of the indices **believe** quantifies over, i.e. relative to the attitude subject's belief state.[11]

(42)

$$\begin{array}{c}\{(\lambda i.\lambda x_e.\forall w'[w' \text{ is } \preceq_{w_i,x}\text{-minimal}\\ \rightarrow \langle w', \preceq_{w_i,x}\rangle \in \mathrm{dom}(f) \wedge \textsc{distinct}_{\preceq_{w_i,x}}(f,g) \wedge \mathbf{roam}(w')(f(\langle w', \preceq_{w_i,x}\rangle))])\\ +(\lambda i.\lambda x_e.\forall w'[w' \text{ is } \preceq_{w_i,x}\text{-minimal}\\ \rightarrow \langle w', \preceq_{w_i,x}\rangle \in \mathrm{dom}(g) \wedge \textsc{distinct}_{\preceq_{w_i,x}}(f,g) \wedge \mathbf{roam}(w')(g(\langle w', \preceq_{w_i,x}\rangle))])\\ \mid f,g \in A_{\langle s,e\rangle} \wedge f,g \text{ are } [\![\mathit{monster}]\!]\text{-concepts } \wedge f \neq g\}\end{array}$$

$$\{\lambda i.\lambda p_{\langle s,t\rangle}.\lambda x_e.\forall w'[w' \text{ is } \preceq_{w_i,x}\text{-minimal} \rightarrow p(\langle w', \preceq_{w_i,x}\rangle)]\} \qquad (41)$$

believe

In the final step, the set of predicate pluralities in (42) is combined with the subject (43). This

[11]For simplicity, we will use the sloppy rendering 'w' is $\preceq_{w_i,x}$-minimal' for '$w' \in \mathrm{dom}(\preceq_{w_i,x}) \wedge \neg\exists w''[w'' \prec_{w_i,x} w']$'.

means that each predicate plurality in (42) is 'matched' with the plurality $\mathbf{a}+\mathbf{b}$ contained in $[\![$*Abe and Bert*$]\!]$ by taking all the compositional covers, performing functional application for each pair in the cover, summing up the results and collecting them in a plural set of propositions. In (43), only two examples of such pluralities are indicated: The first corresponds to a scenario in which Abe's belief is about a **monster**-concept f and Bert's belief about a **monster**-concept g; the second captures a scenario in which Bert's belief is about f' and Abe's belief about g'.

(43)

$$\begin{gathered}
\{(\lambda i.\forall w'[w' \text{ is } \preceq_{w_i,\mathbf{a}} \text{-minimal} \\
\rightarrow \langle w', \preceq_{w_i,\mathbf{a}}\rangle \in \mathrm{dom}(f) \wedge \textsc{distinct}_{\preceq_{w_i,\mathbf{a}}}(f,g) \wedge \mathbf{roam}(w')(f(\langle w', \preceq_{w_i,\mathbf{a}}\rangle))]) \\
+(\lambda i.\forall w'[w' \text{ is } \preceq_{w_i,\mathbf{b}} \text{-minimal} \\
\rightarrow \langle w', \preceq_{w_i,\mathbf{b}}\rangle \in \mathrm{dom}(g) \wedge \textsc{distinct}_{\preceq_{w_i,\mathbf{b}}}(f,g) \wedge \mathbf{roam}(w')(g(\langle w', \preceq_{w_i,\mathbf{b}}\rangle))]), \\
(\lambda i.\forall w'[w' \text{ is } \preceq_{w_i,\mathbf{b}} \text{-minimal} \\
\rightarrow \langle w', \preceq_{w_i,\mathbf{b}}\rangle \in \mathrm{dom}(f') \wedge \textsc{distinct}_{\preceq_{w_i,\mathbf{b}}}(f',g') \wedge \mathbf{roam}(w')(f'(\langle w', \preceq_{w_i,\mathbf{b}}\rangle))]) \\
+(\lambda i.\forall w'[w' \text{ is } \preceq_{w_i,\mathbf{a}} \text{-minimal} \\
\rightarrow \langle w', \preceq_{w_i,\mathbf{a}}\rangle \in \mathrm{dom}(g') \wedge \textsc{distinct}_{\preceq_{w_i,\mathbf{a}}}(f',g') \wedge \mathbf{roam}(w')(g'(\langle w', \preceq_{w_i,\mathbf{a}}\rangle))]),\ldots\}
\end{gathered}$$

$\{\mathbf{a}+\mathbf{b}\}$ (42)

Abe and Bert

Let us now return to the puzzle we started this paper with – the question how to derive the fact that (1a) may be judged true in the ZOMBIE VS. GRIFFIN scenario, but not in the UNSPECIFIED MONSTERS scenario. According to our truth-definition in (30) above, (43) counts as true iff it contains a plurality consisting exclusively of true propositions.

Take the pICs $f = [\lambda w.\iota x.\textbf{zombie-in-A's-room}_w(x)]$ and $g = [\lambda w.\iota x.\textbf{griffin-in-B's-room}_w(x)]$ from the ZOMBIE VS. GRIFFIN scenario. The sum of these pICs corresponds to a predicate plurality $P_f + P_g$ in (42), where both P_f and P_g come with the entailment that at all the minimal worlds relative to the subject's ordering at which f and g are both defined (where there is both a zombie in Abe's room and a griffin in Bert's room) – the values of f and g are distinct. One of the covers used in deriving the set (43) will match Abe with P_f and Bert with P_g, giving rise to a plurality of propositions that each entail this distinctness condition. The proposition $p_{\mathbf{a},f}$ derived by applying P_f to Abe will thus entail that at the closest worlds relative to Abe's belief state at which there is both a zombie in Abe's room and a griffin in Bert's room, the zombie and the griffin are distinct individuals. The proposition $p_{\mathbf{b},g}$ derived from P_g and Bert comes with a similar entailment about Bert's belief state. We can safely assume that both these distinctness conditions are met. Further, $p_{\mathbf{a},f}$ asserts that at each w among Abe's belief worlds, the unique zombie in Abe's room in w is roaming the castle, and $p_{\mathbf{b},g}$ asserts that at each w among Bert's belief worlds, the unique griffin on Bert's bed in w is roaming the castle. Since $p_{\mathbf{a},f}$ and $p_{\mathbf{b},g}$ are thus both true in the scenario, we correctly predict the sentences (1a) and (1b) to be true.

In the UNSPECIFIED MONSTERS scenario, the pICs $f' = [\lambda w.\iota x.\textbf{monster-in-A's-room}_w(x)]$ and $g' = [\lambda w.\iota x.\textbf{monster-in-B's-room}_w(x)]$ also give rise to a plurality of two propositions $p_{\mathbf{a},f'} + p_{\mathbf{b},g'}$ in the way just described. Since the scenario is such that Abe believes there was a monster in his room and Bert believes there was a monster in his room, why aren't the propositions in this plurality true? The reason is the distinctness condition: $p_{\mathbf{a},f'}$ will entail that at each of the closest worlds relative to Abe's belief state at which there was a monster in Abe's room and a monster in Bert's room, these two monsters are distinct. An analogous condition must hold for Bert. If Abe and Bert consider it plausible that a monster can appear to different individuals at different times, these conditions are not met. So even the most plausible choice of **monster**-concepts will not give us a plurality of propositions that are all true in the scenario.

What distinguishes our account from Schmitt (2020) (apart from the broader empirical scope) is the idea to encode an intensional notion of distinctness in these propositional pluralities. We still view cumulative belief as a relation between a plural individual and a plurality of propositions, but treat the distinctness constraint as a condition entailed separately by *each* of the propositions in the plurality. Thus we can compositionally derive cumulative truth conditions while the distinctness constraint is required to hold in a 'distributive' manner for each subject.

To conclude our discussion, let us point out an independent reason for type-lifting the partial individual concepts to quantifiers. Consider our examples (1a)/(1b) in the following scenario:

(44) NO MONSTERS: Abe and Bert both believe that monsters do not exist, but make standard assumptions about how they could be distinguished if they existed (in particular, one and the same monster may appear in different places at different times).

In this scenario, (1a)/(1b) are intuitively false. But if we had based our composition directly on pluralities of pICs, rather than their type-lifted correlates which are total functions, we would have predicted a presupposition failure. The reason is unrelated to the distinctness condition and concerns our general decision to use partial, rather than total, individual concepts. In a system where the indefinite denotes a plural set of pICs, this set would combine with the VP denotation via extensional functional application to yield the denotation in (45) for the embedded clause. The problem is that each atomic part of each plurality in this set will *presuppose* the existence of a monster. Consider now the presupposition-filtering behavior of *believe*: Within a Hintikka-style attitude semantics, it is usually assumed that the presupposition of the complement of *believe* must be met in each of the subject's doxastic alternatives. If so, each part of each propositional plurality in the set (45) would give rise to a presupposition failure in scenario (44) once this set is combined with $[\![believe]\!]$ and the plural subject, because the presupposition that at least one monster exists would not be met in either subject's belief worlds.

(45) $\{(\lambda i : i \in \mathrm{dom}(f) \wedge \textsc{distinct}_{\preceq_i}(f,g).\mathbf{roam}(i)(f(i))) + (\lambda i : i \in \mathrm{dom}(g) \wedge \textsc{distinct}_{\preceq_i}(f,g).\mathbf{roam}(i)(g(i))) \mid f,g \in A_{\langle s,e\rangle} \wedge f,g \text{ are } [\![monster]\!]\text{-concepts} \wedge f \neq g\}$

6. Conclusion

We have given a compositional semantics for cumulative belief sentences like (1a), repeated in (46), focussing on the intuition that the use of a numeral-modified plural indefinite in the embedded clause requires the attitude subjects to have sufficiently distinct 'belief objects'. More specifically, we tried to implement the distinctness constraint described in Haslinger and Schmitt (to appear), which appeals to counterfactual beliefs of the attitude subjects.

(46) *Abe and Bert believe that two monsters were roaming the castle.*

Like Schmitt (2020), we used the Plural Projection account of cumulativity, on which a plurality of individual concepts introduced by *two monsters* 'projects' to a plurality of propositions that is cumulatively believed by the attitude subjects. However, we argued that the DP-denotation itself must be much more complex than assumed in previous work: We had to encode the counterfactual distinctness constraint by lifting the relevant pICs to quantifiers and supplementing each such quantifier with a distinctness requirement – arguably introduced by the plural morpheme – which is in turn inherited by each part of the corresponding propositional plurality.

This should not be an idiosyncratic property of plural DPs in cumulative belief contexts, so we are in fact making a rather radical claim about the semantics of plural DPs in general: They always involve an intensional distinctness condition. Haslinger and Schmitt (to appear) present other constructions that are sensitive to such conditions, but the question still arises: Why don't we usually see any effects of these conditions? Put differently: How do such conditions surface with a semantically singular attitude subject, as in (47a), or in extensional contexts as in (47b)?

(47) a. *Abe believes that two monsters are roaming the castle.*
b. *Two monsters are roaming the castle.*

Consider first (47a), which has a semantically singular subject and thus differs from our original (1a)/(1b) in two crucial respects. First, going back to (42), the denotation for the matrix VP in (47a), it should be clear that our composition principle yields a plural set of propositions for (47a). Each propositional plurality in this set 'matches' Abe with all of the parts of the corresponding predicate plurality, i.e. for any predicate plurality $\mathbf{believe\ p} + \mathbf{believe\ q}$, we have only one cover $\{\langle \mathbf{Abe, believe\ p} \rangle, \langle \mathbf{Abe, believe\ q} \rangle\}$. So for (47a) to be true, there must be a propositional plurality (corresponding to some plurality $f+g$ of **monster**-concepts) such that Abe believes *all* of its parts. The second difference is that since the subject is semantically singular, we no longer evaluate the distinctness constraint relative to a belief state for which the pICs in question might not be instantiated. Our complex DP-semantics is thus 'trivialized', in the sense that its requirements reduce to finding some plurality of **monster**-concepts $f+g$ such that both f and g are instantiated in all of Abe's belief worlds and f and g are distinct relative to his belief state. The latter condition makes reference to the minimal worlds in the ordering $\preceq_{w_i,\mathbf{a}}$ in which both f and g are instantiated – but as we just saw, these are just Abe's belief worlds in the standard sense. The predictions of our theory for this example are therefore indistinguishable from those of more traditional analyses of plural indefinites.

Our condition is similarly 'trivialized' in the non-embedded case in (47b) – the denotation of which is (41). Recall that our distinctness condition appeals to an ordering component of the index that is shifted by the attitude verb (so that it then represents the belief state of the subject). Clearly, any expression that is sensitive to this ordering component – our DP-meanings or also the counterfactuals from Section 4 – should be sensitive to it also in non-embedded contexts. The question thus is what determines this ordering in non-embedded cases (see Yalcin 2007 for the same issue in a domain semantics without orderings). The obvious answer seems to be that this ordering is determined by a contextually provided body of information. We could explicate this for example via the notion of the common ground (see again Yalcin 2007) and generalize the standard conception of a context set to an ordered set. (47b) would then resemble (47a) in that the distinctness condition is relativized to a single ordering. Further, the truth conditions of (47b) would require the existence of a plurality $f+g$ such that both f and g are instantiated in each world compatible with the common ground (since these will be the minimal elements of the 'non-shifted' ordering) and also require f and g to be distinct in all these worlds (as these will be the minimal worlds where both concepts are instantiated). So, regardless of how exactly the relevant ordering is determined in non-embedded contexts, the distinctness condition is trivialized in the sense that it only looks at the minimal worlds in the ordering, which we may identify with the worlds in the context set, following Yalcin (2007). We then predict exactly the same truth-conditions as a traditional extensional analysis of plural indefinites.

References

Aloni, M. (2001). *Quantification under Conceptual Covers*. Ph. D. thesis, University of Amsterdam.

Beck, S. and U. Sauerland (2000). Cumulation is needed: A reply to Winter (2000). *Natural Language Semantics 8*(4), 349–371.

Condoravdi, C., D. Crouch, and M. van den Berg (2001). Counting concepts. In *Proceedings of the Thirteenth Amsterdam Colloquium*, Amsterdam, pp. 67–72.

Edelberg, W. (1986). A new puzzle about intentional identity. *Journal of Philosophical Logic 15*, 1–25.

Edelberg, W. (1992). Intentional identity and the attitudes. *Linguistics and Philosophy 15*(6), 561–596.

Geach, P. (1967). Intentional identity. *Journal of Philosophy 64*, 627–632.

Hamblin, C. L. (1973). Questions in Montague English. *Foundations of Language 10*(1), 41–53.

Haslinger, N. (2020). Individuating beliefs: DP objects of attitude verbs and their domains of quantification. In M. Franke, N. Kompa, M. Liu, J. L. Mueller, and J. Schwab (Eds.), *Proceedings of Sinn und Bedeutung 24*, pp. 305–322. Universität Konstanz.

Haslinger, N. and V. Schmitt (2018). Scope-related cumulativity asymmetries and cumulative composition. In S. Maspong, B. Stefánsdóttir, K. Blake, and F. Davis (Eds.), *Proceedings of SALT 28*, pp. 197–216.

Haslinger, N. and V. Schmitt (to appear). Distinguishing belief objects. To appear in M. Dočekal and M. Wągiel (Eds.), *Formal Approaches to Number in Slavic and Beyond*.

Hintikka, J. (1969). Semantics for propositional attitudes. In *Models for Modalities. Selected Essays*, pp. 87–111. Dordrecht: Reidel.

Kratzer, A. and J. Shimoyama (2002). Indeterminate pronouns: The view from Japanese. In Y. Otsu (Ed.), *The Proceedings of the Third Tokyo Conference on Psycholinguistics*, Tokyo, pp. 1–25. Hituzi Syobo.

Lewis, D. (1973). *Counterfactuals*. Harvard University Press.

Pasternak, R. (2018). Thinking alone and thinking together. In S. Maspong, B. Stefánsdóttir, K. Blake, and D. Forrest (Eds.), *Proceedings of SALT 28*, pp. 546–565.

Quine, W. V. O. (1960). *Word and Object*. Cambridge, MA/London: MIT Press.

Scha, R. (1981). Distributive, collective and cumulative quantification. In J. Groenendijk, T. Janssen, and M. Stokhof (Eds.), *Formal Methods in the Study of Language*, Volume 2, pp. 483–512. Amsterdam: Mathematisch Centrum.

Schmitt, V. (2019). Pluralities across categories and plural projection. *Semantics and Pragmatics 12*(17), 1–49.

Schmitt, V. (2020). Cumulation across attitudes and plural projection. *Journal of Semantics 37*(4), 557–609.

Schulz, K. (2007). *Minimal models in Semantics and Pragmatics*. Amsterdam: ILLC.

Schwager, M. (2007). Bodyguards under cover: the status of individual concepts. In *Proceedings of SALT 17*, pp. 246–263.

Wągiel, M. (2018). *Subatomic Quantification*. Ph. D. thesis, Masaryk University, Brno.

Yalcin, S. (2007). Epistemic modals. *Mind 116*, 983–1026.