

# A challenge to unified accounts of non-canonical belief reports\*

Ido Benbaji-Elhadad

Massachusetts Institute of Technology

## 1. Introduction

This paper is concerned with two observations about belief attribution. The first is Fauconnier’s (1984) observation that we can felicitously ascribe to  $x$  belief in a proposition  $p$ , even when our pre-theoretical intuitions suggest that  $x$  does not *really* believe  $p$ . For example, in scenario (1a), Matt does not believe Alex is on flight AF62, so thus does not believe she landed. Yet (1b) still seems appropriate, perhaps because had Matt known what Sue knows about Alex’s whereabouts, he would have believed Alex landed. Following Blumberg and Holguín (2018), I will call such belief reports *ultra-liberal belief reports*, or ULBs, in short.

The second observation, due to Szabó (2010), is that when  $x$  believes *de re*, of  $n$  individuals, that they bear a property  $P$ , we can attribute to  $x$  the belief that “ $n$  things are  $P$ ” even if  $x$  claims not to know how many relevant *de re* beliefs she has. In scenario (2a), for instance, Alex did not count her allegations, so (2b) is not used to express a *de dicto* belief of Alex that eleven thieves live near her. It also does not convey *de re* that there are eleven *actual* thieves who Alex thinks live near her, since the reporter in (2b) is skeptical of Alex’s claims. Rather, (2b) is used to express what (2c) does, effectively “summing up” in a single report eleven distinct *de re* beliefs that Alex expressed about people depicted in the photos; I thus follow Szabó in calling the intended reading of (2b) its *summative reading*, or SR.

- (1) a. **Scenario:** Alex is on flight AF62 to Boston. Sue and Carol wonder whether she landed. Their pal Matt, who knows of Alex but not of her travels, works at the airport and reports to Sue that he thinks AF62 landed. Sue says to Carol:  
b. Matt believes Alex landed. (adapted from Percus 2020)
- (2) a. **Scenario:** Under police questioning, Alex is shown several photos, and asked, for each, whether it depicts a thief and if so, where the alleged thief lives. Later, Alex notes that she does not think there are any thieves beyond those

---

\*Thanks to Omri Doron, Kai von Fintel, Danny Fox, Keely New, Viola Schmitt, Jad Wehbe and to the reviewers and participants of NELS56 at New York University for valuable feedback. Errors are my own.

she pointed to, and that she did not count her allegations. An officer who did count tallied seventeen alleged thieves, eleven of whom Alex claimed live near her. While skeptical of Alex’s claims, the officer can report to her colleagues:

- b. Alex believes that eleven thieves live near her.
  - c. Eleven individuals are s.t. Alex believes them to be thieves who live near her.
- (adapted from Szabó 2010)

For lack of a better, theory-neutral term, I group (1-2) under the umbrella term *non-canonical belief reports*, intended to indicate that both challenge canonical theories of belief attribution, to be discussed below. However, grouping (1-2) together fails to reflect the fact that they have been introduced in isolation of each other, as evidence for distinct theoretical claims (also discussed further below). Only recently had it been argued that SRs just are ULBs (Benbaji 2021, Mayr and Schmitt 2024, 2026; henceforth M&S). A particular argument for this position, implicit in M&S 2026 and made explicit below, goes as follows: ULBs and SRs seem to be subject to the same distributional constraint. Thus, assuming that the restricted distribution of ULBs results from a feature of the mechanism that derives them, economy considerations demand that we use that mechanism to derive SRs as well. Thus, whatever renders (1b) felicitous in (1a) also renders (2b) felicitous in (2a).

Whether or not this conclusion is justified has non-trivial consequences. In Benbaji 2021, I point out that a successful unification of ULBs and SRs may cast doubt on the theoretical claims that have been motivated by the existence of SRs. In particular, for Szabó (2010), the SR in (2) is used to motivate a non-standard scoping mechanism, which allows a determiner (e.g., “two”) and its NP complement (e.g., “thieves”) to scope in different positions relative to an intensional operator (e.g., “believe”). However, the mechanisms posited in the literature to derive ULBs are all orthogonal to the scopal properties of the DPs in a belief report. Thus, if SRs can be successfully reduced to ULBs, then the existence of SRs would not inform our understanding of DP-scope. So there are reasons beyond theoretical economy to explore further whether unification is empirically justified.

This short note explores just that. I show that the distributional restrictions on ULBs differ non-trivially from those regulating the availability of SRs, and that if SRs were ULBs we would not expect to observe them in the contexts in which we in fact do. The discussion is organized as follows: Section 2 presents the challenges that non-canonical belief reports posit for canonical theories of belief attribution. Section 3 reviews the mechanisms proposed to defuse these challenges, and Section 4 motivates the unification of non-canonical belief reports under such mechanisms. Section 5 then offers an empirical challenge for unification. Section 6 concludes by weighing the theoretical implications of the challenge, particularly for Szabó’s claims regarding SRs and DP scope.

## 2. What is non-canonical about non-canonical belief reports?

Most contemporary work on belief takes as its starting point the idealization, due to Hintikka (1962), that belief involves universal quantification over worlds (3). On this view, attributing to individual  $x$  in world  $w$  belief in  $p$  amounts to saying that  $p$  holds in all worlds compatible with the what  $x$  believes in  $w$ ; i.e., in  $x$ ’s *doxastic state* in  $w$ , or  $\text{DOX}(x, w)$ .

$$(3) \quad \llbracket \text{believe} \rrbracket = \lambda p_{\langle s, t \rangle} \lambda x_e \lambda w_s. \forall w' \in \text{DOX}(x, w) : p(w') = 1$$

The two phenomena observed above challenge the Hintikkaan idealization in different ways. ULBs are instances where belief can be felicitously attributed even when the predicate in (3) does not hold between  $x$  (aka the *belief holder*) and  $p$  (aka the *prejacent proposition*) in  $w$  (aka the *evaluation world*); in (1), it simply is not the case that throughout Matt’s doxastic state, Alex landed (for all Matt knows, Alex might not even be traveling).

The case of SRs is more complicated. As Szabó notes, given its closure properties, Hintikkaan belief in (3) actually validates (2b) in the scenario in (2a). First, (3) is closed under conjunction:  $\llbracket \text{believe} \rrbracket(p, x, w)$  and  $\llbracket \text{believe} \rrbracket(q, x, w)$  entail  $\llbracket \text{believe} \rrbracket(p \wedge q, x, w)$ , and given that in scenario (2a) Alex has eleven beliefs, of distinct individuals, that they are her neighborly thieves, she also believes the conjunction of these beliefs. Second, belief is closed under entailment:  $\llbracket \text{believe} \rrbracket(p, x, w) \rightarrow \llbracket \text{believe} \rrbracket(q, x, w)$  whenever  $p \rightarrow q$ , and the conjunction of Alex’s individual beliefs entails the existence of eleven neighborly thieves.

SRs are non-canonical not because the Hintikkaan idealization does not validate them (it does), but because it validates them only in virtue of some of its properties widely considered quite problematic. The closure properties exhibited by (3) render it unrealistic as an account of belief; our beliefs do not seem to be deductively closed. And while (3) is often a useful simplification, “if we don’t want to take the blame when the simplifying assumption of logical omniscience leads to unacceptable predictions we should not take credit when it accidentally delivers the right result” as is arguably the case with SRs (Szabó 2010: 33). In other words, SRs are non-canonical because once we reject (3) in favor of an account of belief sans closure properties, (2b) is no longer automatically validated in scenario (2a). If  $x$  can believe  $p$  and  $q$  without believing  $p \wedge q$ , or believe  $p$  while failing to believe  $q$  even when  $p \rightarrow q$ , having eleven distinct *de re* beliefs, of eleven individuals, that they are thieves would no longer entails believing there are eleven thieves.

Now, it is not immediately obvious whether one’s failure to count the objects about which one holds a certain belief should be modeled as a failure of conjunction or entailment closure; e.g., whether in (2a) Alex fails to believe the conjunction of her distinct *de re* beliefs, of eleven individuals, that they are thieves, or fails to believe that there are eleven thieves, despite believing the conjunction? It seems to me that entailment closure is the culprit here. To see this, consider another, somewhat simpler context involving a toddler who has not yet learned to count.<sup>1</sup> She can complain that her friend took her stuffed animal, her toy car, and her baby doll—indicating that she *believes* that her friend took her stuffed animal, and that he took her toy car, and that he took her baby doll—without believing that her friend took *three* of her toys. Nevertheless, an adult observer can ascribe to her this latter belief (perhaps, say, as an explanation for why she is upset). Thus, we need an account of belief without entailment closure; i.e., an account that fits the schema in (4).

- $$(4) \quad \begin{array}{ll} \text{a.} & \llbracket \text{believe} \rrbracket(p, x, w) = \mathcal{B}(p, x, w) \\ \text{b.} & \exists x, p, q, w : p \rightarrow q \wedge \mathcal{B}(p, x, w) \wedge \neg \mathcal{B}(q, x, w) \end{array}$$

<sup>1</sup>I thank Danny Fox for suggesting this to me in personal communication.

There is a long and involved discussion in the philosophical literature regarding the correct way to rid our account of belief of its various closure properties.<sup>2</sup> As it turns out, closure under entailment is perhaps “stickier” than other closure properties, in that it is harder to get rid of (cf. Yalcin 2018). I cannot wade into that discussion here, and will have to stipulate that the correct account of belief—whatever it turns out to be—meets the schema in (4). Note that I will also take for granted that under an account without entailment closure, both ULBs and SRs remain problematic; i.e., that such an account still falsifies (1b) and (2b) in scenarios (1a) and (2a), respectively.

### 3. Accounting for non-canonical belief reports

In what follows, I will be concerned with the claim that SRs like (2) are subsumed under the mechanism that derives ULBs like (1). Let us start, then, by reviewing the mechanisms proposed in the literature: On the *replacement* approach (Percus 2020, Benbaji-Elhadad 2023, M&S 2024, 2026), (1b) ascribes belief to Matt, but not belief in the overt prejacent in the sentence; rather, it ascribes to him belief in some proposition that stands in a particular relation to the prejacent. The nature of the relation is up for debate, but (5a) is a useful approximation, demanding that the proposition, together with something the reporter believes, entail the prejacent.<sup>3</sup> Thus, (1b) is true in scenario (1a), because Matt believes that flight AF62 landed, which in turn, together with Sue’s knowledge that Alex is on flight AF62, entails that Alex did. On the *revision* approach, (1b) ascribes belief in the overt prejacent, but not to Matt; rather, it ascribes that belief to someone like Matt, but whose doxastic state has been revised with some of the attitude reporter’s beliefs (Blumberg and Lederman 2021). In (5b) this is spelled-out as demanding that (counterfactually) updating the attitude holder’s doxastic state with a proposition believed by the reporter, would result in them believing the prejacent. The truth of (1b) given (1a) is explained here by the fact that, had Matt known—like Sue—that Alex is on flight AF62, he would have believed that she landed.

- (5)  $x$  can truthfully ascribe to  $y$  a belief  $p$  in world  $w$  iff *either*  $\llbracket \text{believe} \rrbracket(p, y, w)$  *or*:
- a. (On the *replacement* approach:) There is a proposition  $q$  that together with what  $x$  believes entails  $p$ , s.t.  $\llbracket \text{believe} \rrbracket(q, y, w) = 1$
  - b. (On the *revision* approach:) There is a proposition  $q$  known to  $x$ , s.t., had  $\llbracket \text{believe} \rrbracket(q, y, w)$  held,  $\llbracket \text{believe} \rrbracket(p, y, w)$  would have held too.

Some terminology: When the report that  $y$  believes  $p$  in  $w$  is true in virtue of the truth of  $\llbracket \text{believe} \rrbracket(p, y, w)$ , I will say that  $y$  *strictly-speaking-believes*  $p$ , or:  $S\text{-believe}(y, p)$ , omitting mentions of the evaluation world going forward, for readability. If the report is true in virtue of one of the conditions in (5a)-(5b), I will say that  $y$  *R-believes*  $p$  or  $R\text{-believe}(y, p)$ , where  $R$  stands for revision or replacement, depending on one’s preferred theory.

<sup>2</sup>Cf. Fagin et al. 1995 for an instructive overview of this discussion in its early stages.

<sup>3</sup>See Benbaji-Elhadad 2023 for reasons why entailment given  $x$ ’s beliefs is an over-simplification.

Both (5a) and (5b) extend to SRs. In scenario (2a), Alex undergoes two rounds of questioning. First, she identifies photos that she thinks depict thieves, and then she identifies which of the photos she identified as thieves depict her neighbors. Thus, she *S*-believes that every photo she identified in both rounds depicts a thief who lives near her. The police officer uttering (2b) counted Alex’s allegations, concluding that she identified eleven photos in both rounds. Thus, (2b) is rendered true in scenario (2a) on the revision view (5b), because had Alex *S*-believed what the officer does about the number of her allegations, she would have *S*-believed eleven thieves live in her building. The replacement view in (5a) verifies (2b) in scenario (2a), because there is a proposition that Alex *S*-believes—namely, that every photo selected in the two rounds depicts a thief who lives near her—that together with what the officer believes entails the prejacent proposition of the report in (2b).

That the mechanism designed to derive ULBs extends to SRs is already *ipso facto* an argument for unifying the two phenomena under this mechanism. Why should our semantics contain a particular mechanism to derive SRs, when the one it posits to derive ULBs already does that? Next, I will outline a stronger argument for unification, which rests on the observation that ULBs and SRs seem subject to the same distributional constraint.

#### 4. The argument for unification

M&S (2024, 2026) observe that there are instances in which  $x$  does not *S*-believe  $p$ , in which attributing to  $x$  belief in  $p$  seems unacceptable despite the fact that  $x$  does in fact *R*-believe  $p$ . In particular, they suggest, the *question-under-discussion* (henceforth: QuD) plays a role in regulating when *R*-belief may be appealed to in evaluating the truth of a belief report. The QuD is an implicit question discourse participants would like to resolve, posited as a parameter of interpretation (cf. Roberts 1996/2012), in addition to the context set  $c$  (6a). It is defined in (6b) as a partition of  $c$  into cells, representing the “task” discourse participants are faced with at a given stage in discourse; namely, the task of agreeing on where the actual world resides, among these cells.

- (6) a.  $c = \{w : w \text{ is compatible with the shared-knowledge of discourse participants}\}$   
 b. QuD  $Q$  is a partition of  $c$ , i.e. a set of non-empty elements  $\{q_1, \dots, q_n\}$  s.t.  
 $\bigcup Q = c \wedge \forall w \in c \exists q \in Q : [w \in q \wedge \neg \exists q' \neq q : w \in q']$

M&S observe that  $R\text{-believe}(x, p)$  must be *indistinguishable* from  $S\text{-believe}(x, p)$  for the purpose of answering the QuD, if *R*-belief is to be appealed to in verifying a belief report; i.e., that both must entail the same answer to the QuD. This is formalized in (7), though in a way that departs significantly from M&S’s formulation, for ease of exposition.

- (7) **M&S’s QuD constraint:** A statement ascribing to  $x$  belief in proposition  $p$ , uttered in world  $w$  and relative to context  $c$  and QuD  $Q$ , can be true in virtue of  $x$ ’s *R*-belief that  $p$  only if:  $S\text{-believe}(x, p)$  and  $R\text{-believe}(x, p)$  resolve  $Q$  in the same way; i.e.,  $\exists Q' \subseteq Q : S\text{-believe}(x, p) \subseteq \bigcap Q' \wedge R\text{-believe}(x, p) \subseteq \bigcap Q'$ .

The constraint is satisfied in (1), since the question Sue and Carol wish to resolve is *whether Alex landed*, and given that Matt is an authority on flight-schedules (i.e., given that it is part of the common ground that him *S*-believing a flight landed is sufficient indication that the flight did, in fact, land), his *S*-belief and *R*-belief that Alex landed both resolve the question affirmatively. Thus, (1b) can be true in virtue of Matt's *R*-belief in scenario (1a).<sup>4</sup>

However, (8) illustrates a violation of the constraint. The QuD in (8a) is a question about Matt's *S*-beliefs—e.g., “*is the surprise spoiled?*” or “*does Matt suspect Alex landed?*”—which is not resolved in the same way by Matt *S*-believing and *R*-believing that Alex landed. In fact, while it is resolved, affirmatively, by the former, the latter leaves it unresolved: that Matt *would have S*-believed that Alex landed had he known that she is on AF62 (on the revision view), or that something else that he believes—presumably, that AF62 landed—entails that Alex landed given Sue's beliefs (on the replacement view), simply does not bear on the QuD. Thus, (8b) cannot be true if Matt does not *S*-believe that Alex landed, irrespective of him *S*-believing that AF62 landed, and thus, *R*-believing that Alex did. This is in line with our intuitions about (8).

- (8) a. **Scenario:** Alex is on AF62. Her visit is a surprise organized by Sue and Carol for the birthday of airport employee Matt. Matt reports to Sue that he believes AF62 landed. Independently, Carol is worried that the surprise has been spoiled. She asks Sue whether that is the case. Sue answers:  
b. #Matt thinks that Alex landed.

Prima facie, M&S's constraint in (7) also seems to regulate the availability of SRs, supporting the claim that ULBs and SRs form a natural class. This can be shown by examining whether, e.g., (2b) has a SR even when the QuD has to do with Alex's *S*-beliefs regarding the *number* of individuals she identified as thieves. Such a scenario is provided in (9a). There, the police officers engaged in conversation are presumably interested in a question like “*how many individuals does Alex think she alleged are thieves who live near her?*” or “*did Alex count her allegations correctly?*” Alex's *S*-belief that she identified eleven individuals would resolve these questions, while her *R*-belief that she did so simply does not bear on them. Indeed, the report in (9b) seems infelicitous in such a scenario, as predicted if (7) constrains the distribution of ULBs, and SRs just are a special case of ULBs.

- (9) a. **Scenario:** Alex is a witness in a murder trial, having seen the suspect hand the victim a certain number of pills claimed to be the cause of death. Whether

---

<sup>4</sup>There is a critical complication here that I set aside. *Q* is defined as a partition of *c*, so every cell in *Q* is a subset of *c*. But *R*-belief can be felicitously appealed to in, e.g., (1) even when the respective *S*-belief is false in *c*; i.e., even if it is presupposed that Matt *S*-believes Alex is *not* on AF62, and thus did not land. In that case no subset of *c*—and therefore no cell in *Q*—is entailed by Matt's *S*-belief that Alex landed. As far as I can tell, this wrinkle applies to M&S system as well, and addressing it in full is beyond my immediate goals. For the current purposes, as an *ad hoc* solution, we can take *Q* in (7) to be a partition not of *c*, but rather of the set of worlds most similar to *c* in which there are some worlds where *S*-believe(*x*, *p*) holds. Note that *Q* cannot be a partition of logical space, as *some* contextual information must be retained for the condition in (7) to be met; e.g., Matt's *S*-belief and *R*-belief that Mary landed only settle the question of whether she did if Matt is presumed an authority on flight schedules, which is—of course—not a logical truth.

### *A challenge to unified accounts of non-canonical belief reports*

Alex counted the pills correctly may decide the trial, and the police want to test her reliability. An officer who knows Alex to be paranoid about thieves, subjects her to the kind of questioning described in scenario (2a). If Alex ends up counting her allegations correctly, that would attest to her strength as a witness in this case. A group of officers present in Alex's interview count her allegation, tallying eleven of them. Alex later reports to one of the officers that she did not count her allegations. That officer reports to her colleagues:

- b. Alex believes that eleven thieves live near her.

However, to motivate the unification thesis, it is not enough to show that when M&S's constraint on ULBs in (7) is violated, SRs are unavailable. One must also show that when SRs are available, the constraint is *not* violated. This, I will argue, is the sticking point in the argument, as SRs can be available despite clear violations of the QuD constraint.

## **5. The argument against unification**

Szabó (2010) observed that (2b) has a SR in scenario (2a). Now, whether or not this is expected on a theory that unifies SRs and ULBs depends on the QuD relative to which (2b) is uttered, which is not entirely transparent in (2a). It seems to me that the most natural QuD, given that interlocutors in that scenario are officers discussing the results of Alex's police interview, is "*how many individuals did Alex accuse of being her neighboring thieves?*"

This QuD does not distinguish *R*-belief from *S*-belief, but only on certain conditions. In particular, that Alex *S*-believes there to be eleven neighbor-thieves, and that she *R*-believes so, would resolve the QuD in the same way only on the assumption that Alex is in general an accurate counter who does not *S*-believe there to be any thieves near her beyond those she identified. These assumptions are necessary because for the *S*-belief to bear on the QuD in the first place, it would have to have been formed solely on the basis of the allegations Alex made under questioning, and on her having accurately counted these allegations.<sup>5</sup> If that *S*-belief could, in principle, have been formed on the basis of (e.g.,) Alex's independent intuitions about the number of thieves living near her, it would simply fail to bear on the QuD, which inquires about the results of Alex's interview.

The argument against unification begins with the observation that even when these conditions are not met, (2b) seems to have a SR. For instance, consider (10).

- (10) a. **Scenario:** Several people are under questioning, each presented with photos and asked, as in (2a), to first identify thieves then identify neighbors among their alleged thieves. All officers involved know what the task consists of, but only one tallies allegations. Upon leaving, each interviewee reports to all officers whether they counted their allegations; among those who did not count, some may have had prior beliefs about the number of thieves living near

---

<sup>5</sup>One would also have to assume that the photos Alex was shown depict distinct individuals, and that Alex is aware of this. Otherwise, her *S*-belief that eleven thieves live near her might be formed on the basis of her failing to notice that she identified several photos of the same individual under questioning. I set aside this issue, first raised by Ben-Yami (2014), for the remainder of this paper.

them, and among those who did, some are known to be unreliable counters who may have formed a belief that there is a certain number of thieves near them, distinct from the number of people they actually identified. The officer who kept track of the results, reports them to her colleagues by attributing, to each interviewee, a belief that  $n$  thieves live near them, where  $n$  is the number of photos identified by that interviewee as depicting neighborly thieves.

- b. Alex believes that eleven thieves live near her.

Given scenario (10a), the QuD relative to which (10b) is uttered seems to be an inquiry about the results of Alex's questioning. However, none of those conditions that render  $R$ -belief indistinguishable from  $S$ -belief modulo *that* QuD hold here. First, Alex is not guaranteed to be an accurate counter, as it is established that at least some of the interviewees miscounted. Second, Alex is not guaranteed not to have a prior belief, one which is not based on the task she was asked to perform, about the number of thieves who live near her. Therefore, her  $S$ -belief that eleven thieves live near her would not bear on the QuD, whereas her  $R$ -belief that this is so would resolve it. So (10b) has a SR even when (7) is violated. If it is indeed the constraint in (7) that regulates the availability of ULBs, then the unification thesis undergenerates attested instances of SRs.<sup>6</sup>

---

<sup>6</sup>As an exploratory aside, note that (10) may be indicative of a deeper problem for the unification thesis: Even if one manages to weaken the QuD constraint just enough so that it does not rule out (10b), while making sure it still rules out infelicitous examples like (8b), an independent constraint on ULBs is needed that may itself already rule out (10b). To see why, observe that in scenario (10a), it is common ground among the police officers engaged in discourse that each interviewee  $S$ -believes that all photos they identified in both rounds of questioning depict thieves who live in their building. So the only "new information" contributed by the belief report in (10b) is the knowledge of the reporting officer that Alex identified eleven such photos. Alex's own belief state, it seems, contributes no new information, and is thus *redundant* given (10a).

But ULBs generally seem infelicitous when the attitude holder's belief state is redundant in this way. For instance, consider the variation on Percus's original scenario below. In scenario (ia), Sue cannot attribute to Matt the belief that Alex is flying business on AF62 (ib), despite the fact that Matt  $R$ -believes that she is: His  $S$ -belief that Alex always flies business, together with Sue's knowledge that Alex is on AF62, entails the prejacent in (ib), so the replacement view (5a) predicts (ib) to be felicitous. And revising Matt's belief state with Sue's knowledge that Alex is on AF62 would result in him  $S$ -believing that Alex is flying business on AF62, so the revision view (5b) predicts the same thing.

- (i) a. **Scenario:** Alex is on an incoming flight to Boston, and Carol wants to know which. Sue knows the flight is AF62. Their pal Matt, who knows Alex but does not know she is traveling, works at the airport and they expect him to run into Alex when she lands. Independently, they both know that Matt falsely believes that when Alex does travel, she flies business. Sue to Carol:
- b. #Matt believes Alex is in the business class of flight AF62.

As things stand, then, both theories of ULBs must be further constrained so as to prevent one from appealing to  $R$ -belief in cases like (i), and it seems to me quite difficult to formulate a constraint that rules out (ib) in scenario (ia) without ruling out (10b) in scenario (10a), given that the examples are construed in the same way. To see this, consider how the infelicity of (ib) might be explained on the replacement view. That view allows  $x$  to attribute to  $y$  belief in proposition  $p$  even when  $y$  does not  $S$ -believe  $p$ , as long as  $y$   $S$ -believes some proposition  $q$ , and  $q$  together with some proposition  $x$  takes to be true, entails  $p$ . The relevant  $q$  that Matt  $S$ -believes in scenario (ia) is the proposition that Alex always flies business, and the fact that Matt  $S$ -believes this proposition is common ground among Sue and Carol, so (ib) contributes no new



## 6. Theoretical implications

This paper examines the thesis that two kinds of belief reports which both challenge canonical theories of belief attribution, can be unified under a single theory. As I note at the outset, the two phenomena involved—*ultra-liberal belief reports* (ULBs) and *summative readings* (SRs)—have been introduced in isolation of each other, and used to motivate different theoretical claims. On the one hand, ULBs have been used to motivate a general claim about the way we evaluate the truth of belief reports; namely, the claim that we must posit a replacement (5a) or revision (5b) mechanism that can verify belief reports in virtue of *R*-belief, at least in certain scenarios. SRs, on the other hand, have been appealed to more specifically in a debate about the scope of DPs.

Szabó (2010) introduces the SR in (2) as evidence that DPs can take “split scope” relative to an attitude predicate. In particular, it is argued that the truth of (2b) in scenario (2a) is due to the determiner of the DP “*eleven thieves*” effectively splitting from its NP complement, taking quantificational scope above the belief predicate, while leaving its restrictor behind in that predicate’s scope to be interpreted opaquely. In other words, Szabó argues that (2b) is mapped to a structure along the lines of (11), deriving the proposition that there are eleven individuals of whom Alex believes that they are thieves who live in her building.<sup>7</sup>

$$(11) \quad [[\exists \text{ eleven}] [\lambda_i [\text{Alex believes } [[x_i \text{ thieves}] \text{ live in her building}]]]]$$

Such an unorthodox scoping mechanism has been argued to be too powerful, allowing DPs to take split scope even when this would give rise to unattested readings,<sup>8</sup> and the unification of SRs and ULBs has been cited as a way to avoid it (cf. Benbaji 2021, Mayr and Schmitt 2026). This is because the mechanisms proposed to derive ULBs in (5) operate independently of the scopal properties of the DPs in the prejacent of the belief reports involved. Thus, if SRs can be derived via such mechanisms, their existence would not inform our understanding of DP scope, contra the conclusion pursued by Szabó.

This paper illustrates that unification still faces an empirical challenge: ULBs seem to be sensitive to the QuD in a way that SRs are not. In particular, if the mechanisms that

---

information beyond Sue’s knowledge that Alex is on flight AF62. We can thus reformulate the replacement approach as allowing  $x$  to attribute to  $y$  belief in proposition  $p$  given a context  $c$ , as long as  $y$  *S*-believes some proposition  $q$  s.t.  $q$  together with some proposition  $x$  takes to be true entails  $p$ , **and**  $c$  does not already entail that  $y$  *S*-believes  $q$  (i.e., it is not common ground that  $y$  *S*-believes  $q$ ). But, of course, the scenario in (10a) is one in which it is common ground that Alex *S*-believes the proposition  $q$ , which together with the reporting officer’s knowledge entails the prejacent in (10b); namely, the proposition that all photos Alex identified in both rounds of questioning depict neighborly-thieves.

I am of course not ruling out the possibility that *some* formulation of a constraint on ULBs may rule out (ib) but not (10b) in their respective scenarios, but merely pointing out that this is an open challenge to unification on top of the challenge presented above. Note that as far as I know, Blumberg and Lederman (2021) were the first to note the need for a non-redundancy constraint on ULBs, but their examples differ non-trivially from (i) and it is beyond the scope of this paper to explore this difference.

<sup>7</sup>I abstract over the technical details of Szabó’s proposals for brevity.

<sup>8</sup>See discussion in Santorio 2013, Keshet and Schwarz 2019, Benbaji 2021, Benbaji-Elhadad 2026.

derive ULBs are subject to a QuD constraint akin to (7), as argued by Mayr and Schmitt (2024, 2026), these mechanisms are not expressive enough to generate all instances of SRs. Thus, without overcoming this empirical challenge, one cannot yet rule out the claim that SRs really are indicative of some non-standard scoping mechanism as Szabó argues.

The unification thesis is appealing: Not only does it let us avoid the overgeneration that comes with adopting a DP-splitting operation for SRs, but it also accounts for those parallels between ULBs and SRs reviewed above. The goal of this paper is to point out that there is still no precise generalization of when the two phenomena pattern together and when they come apart, which seems to be a precondition for a successful unified account.

## References

- Ben-Yami, Hanoch. 2014. Bare quantifiers? *Pacific Philosophical Quarterly* 95:175–188.
- Benbaji, Ido. 2021. Restricting the fourth reading. *SALT* 31:221–240.
- Benbaji-Elhadad, Ido. 2023. Preliminaries for a substitution theory of de re. *SALT* 33:358–378.
- Benbaji-Elhadad, Ido. 2026. Reconstruction, bare quantification, and the fourth reading. Manuscript, MIT.
- Blumberg, Kyle, and Ben Holguín. 2018. Ultra-liberal attitude reports. *Philosophical Studies* 175:2043–2062.
- Blumberg, Kyle, and Harvey Lederman. 2021. Revisionist reporting. *Philosophical Studies* 178:755–783.
- Fagin, Ronald, Joseph Y Halpern, Yoram Moses, and Moshe Vardi. 1995. *Reasoning about knowledge*. MIT press.
- Fauconnier, Gilles. 1984. *Mental spaces*. Cambridge, MA: MIT Press.
- Hintikka, Jaakko. 1962. *Knowledge and belief*. Ithaca, NY: Cornell University Press.
- Keshet, Ezra, and Florian Schwarz. 2019. De re/de dicto. In *The oxford handbook of reference*, ed. by Jeanette Gundel and Barbara Abbott, 167–202. Oxford: Oxford University Press.
- Mayr, Clemens, and Viola Schmitt. 2024. Non-de dicto construals as a unified phenomenon. In *SALT* 34.
- Mayr, Clemens, and Viola Schmitt. 2026. Generalized non-de dicto: Underspecification and discourse stability. Unpublished Manuscript.
- Percus, Orin. 2020. Index-dependence and embedding. *The Wiley Blackwell Companion to Semantics*. 1–31.
- Roberts, Craige. 1996/2012. Information structure: Towards an integrated formal theory of pragmatics. *Semantics and pragmatics* 5:6–1.
- Santorio, Paolo. 2013. Descriptions as variables. *Philosophical Studies* 164:41–59.
- Szabó, Zoltán Gendler. 2010. Specific, yet opaque. *Amsterdam Colloquium* 17:32–41.
- Yalcin, Seth. 2018. Belief as question-sensitive. *Philosophy and Phenomenological Research* 97:23–47.