

Revisionist Reporting*

Kyle Blumberg
Harvey Lederman[†]

June 16, 2020

Abstract

Several theorists have observed that attitude reports have what we call “revisionist” uses (Pryor, 2004; Hawthorne & Manley, 2012; Recanati, 2012; Blumberg & Holguín, 2018). For example, even if Pete has never met Ann and has no idea that she exists, Jane can still say to Jim ‘Pete believes Ann can learn to play tennis in ten lessons’ if Pete believes all six-year-olds can learn to play tennis in ten lessons and it is part of Jane and Jim’s background knowledge that Ann is a six-year-old. Jane’s assertion seems acceptable because the claim she reports Pete as believing (that Ann can learn to play tennis in ten lessons) is entailed by Pete’s beliefs if they are revised in light of Jane and Jim’s background knowledge. We provide a semantic theory of revisionist reports based on this idea. We observe that the admissible “revisions” are *limited* in a striking way. Jane cannot say ‘Pete thinks Ann is a six-year-old and can play tennis in ten lessons’ in the same context that she can say ‘Pete believes Ann can learn to play tennis ten lessons’, even though this too follows from Jane and Jim’s background knowledge together with what Pete believes. Our theory predicts the infelicity of these latter reports. It also has the resources to predict the truth of “exported” attitude reports (Quine, 1956; Sosa, 1970; Kripke, 2011) and casts the relationship between these reports and “singular thought” in a new light. We conclude by discussing how revisionist reports make trouble for a simplistic view of the connection between

*Forthcoming in *Philosophical Studies*. Please cite published version.

[†]The authors contributed equally. This paper took its final shape in response to a series of thoughtful, inspiring comments from a number of people. Most notably: Cian Dorr and Jeremy Goodman convinced us to reconsider the arguments of Holguín (2018) that motivate using questions in the manner discussed in §3.2; Simon Goldstein, Daniel Rothschild, and Bernhard Salow persuaded us to use the official account of §3.3, as opposed to the disjunctive one now discussed in §4; and Matt Mandelkern pointed us to a new style of formulation of an older non-redundancy condition, altering the direction of §5. Dorr, Goldstein, Goodman, Mandelkern and Rothschild also gave extremely helpful comments on the rest of the paper, as did Ben Holguín, Ezra Keshet, Henry Schiller, Tomasz Zyglewicz and an anonymous referee. We thank them all.

the relations expressed by attitude verbs in natural language and the relations of most interest to philosophers of mind and cognitive science.

1 Introduction

Consider the following case:

Tennis: Ann is a six-year-old girl whom Pete, an expert in tennis pedagogy, has never met and whose existence he is unaware of. Pete believes that every six-year-old can learn to play tennis in ten lessons. Jane, Ann’s aunt, is aware of Pete’s views on the matter. Jane wants to encourage Ann’s father, Jim, to sign Ann up for tennis lessons, so in conversation with Jim she asserts the following:

- (1) Pete believes Ann can learn to play tennis in ten lessons.¹

(1) is acceptable in context, and bears none of the hallmarks of insincere or non-literal speech. Thus, we have reason to think that it is true. But if it is true, then this is quite surprising. After all, Pete has no idea who Ann is. So, on the one hand, good semantic methodology suggests that (1) is not just acceptable, but true. But, on the other, reflection on facts about Pete’s doxastic state—what, as one might overemphatically put it, Pete *really* believes—suggests the report is false.²

The first aim of this paper is to develop a semantics for attitude reports which resolves this tension, by allowing that the reports are true, while also making some sense of the idea that Pete does not really believe that Ann can learn to play tennis in ten lessons. We will call the relevant readings of reports like (1) “revisionist attitude reports”. In our view, these ascriptions illustrate that speakers can truly report ascribees as having attitudes which are “revised” in light of information that is shared among the speaker’s interlocutors (but which may not be available to the ascribee). For instance, Jane and Jim know that Ann is a six-year-old, although Pete does not. Pete

¹This example is taken from (Blumberg & Holguín, 2018). Pryor (2004), Hawthorne & Manley (2012), and Recanati (2012) discuss similar cases—see §2. Blumberg & Holguín (2018) call these ascriptions “ultra-liberal de re reports”. However, as we will show, the phenomenon of which (1) is an instance is exhibited by a range of attitude reports, not just those that are in some sense “de re”.

²As Blumberg & Holguín (2018) point out, reports such as (1) also pose a problem for so-called “descriptivist” approaches to attitude reports, e.g. (Kaplan, 1968), (Lewis, 1979), (Aloni, 2005), (Maier, 2009). Roughly, these accounts maintain that (1) indirectly ascribes Pete a belief involving a definite description that denotes Ann. However, Pete needn’t have any such belief.

does however believe that all six year olds can play tennis in ten lessons. It follows from the conjunction of the claims that Ann is a six year old and that all six year olds can learn to play tennis in ten lessons that Ann can play tennis in ten lessons. The truth of (1) illustrates that Jane can “revise” Pete’s beliefs with Jane and Jim’s background information, and report this consequence of the revised belief as something Pete himself believes.

Revisionist reporting is visible in a wide range of constructions. Consider:

Running: I tell Mary that some of my friends have entered the race. Mary thinks that all of my friends are fit and strong, so she says ‘A friend of yours will win the race’. Mary has no idea which of my friends entered. Now suppose that we’re at the start looking at the racers lined up. Of my friends, only Bill, Ted and Ann have entered the race, and they’re standing to my left, mixed in with some other racers. I say (2):

- (2) Mary thinks that a person to my left will win the race.³

(2) is acceptable in context. But neither of the standard resolutions of the ambiguity that arises when descriptions (‘a person to my left’) interact with intensional operators can be used to explain this fact. On a ‘narrow-scope’ reading of the indefinite, the sentence could be paraphrased as ‘Mary thinks there is a person to my left who will win the race’, and requires that Mary thinks that there are some people to the left of the speaker in the lineup. But this reading won’t be true in the scenario above: given our setup, it’s consistent with what Mary believes that no one at all is to the left of the speaker in the lineup. On a ‘wide-scope’ reading of the indefinite, by contrast, the sentence could be paraphrased as ‘There is a person to my left who Mary thinks will win the race’. But this reading won’t be true in the scenario above either: its truth requires that Mary have beliefs about which friends of the speaker’s have entered the race, and she has no such beliefs.⁴ But in context, the revisionist report (2) seems to be true. On the theory we will develop below, the report is true because the speaker can

³This case is similar to that of *Foyle’s investigation* in (Schwager, 2009). Note though that the theory Schwager develops in response to her example does not smoothly handle ours; the problem for her theory arises from the fact that in our example (unlike Schwager’s) the predicate used in the prejacent (‘person to my left’) has an extension at the actual world which is not a superset of the extension of the predicate (‘one of my friends’) most saliently used to describe what Mary “really” believes.

⁴We’ve labeled these readings in a way intended to be familiar to philosophers, as though they arise from a scope-ambiguity, but this should be seen as just a rough description: a great deal of recent work in linguistics has focused on developing non-scopal accounts of the relevant ambiguity, so that they would be better (more neutrally) called ‘transparent’ and ‘opaque’ readings, respectively. Recent, sophisticated theories of the interaction of determiners with intensional operators also make available a third reading

report Mary as believing what is entailed by her beliefs revised in light of further information, in this case, in light of who is standing to the speaker's left. The example illustrates that revisionist reporting is admissible not just when moving from a general belief (that all six year olds can learn to play tennis in ten lessons) to a specific one (that Ann can learn to play tennis ten lessons), but also with other entailment patterns. In the next section we will also show how revisionist reporting arises with a range of attitude verbs.

After outlining a basic semantics for revisionist reports, we highlight further features of the data, and consider connections between revisionist reports and other, more familiar phenomena. First, we observe that revisionist reports are *limited* in a striking way. On our account, the acceptability of (1) and (2) can be explained by the fact that the prejacent in each report is entailed by the conjunction of the conversational participants' background information with the ascriber's beliefs. But a speaker cannot truly report *any* proposition of this kind. Although (1) is perfectly acceptable in the context of *Tennis*, (3) is not (as indicated by the '#' preceding the example):

- (3) # Pete believes Ann is a six-year-old.

But the proposition that Ann is a six-year-old is entailed by the conversational participants' background information, and so *a fortiori* is entailed by the conjunction of their background information and Pete's beliefs. We suggest the infelicity of, e.g. (3) can be explained by the fact that revisionist reports carry a *non-redundancy condition*: that the subject's mental state cannot be redundant in deriving the ascribed attitude, or, more exactly, that the proposition in which belief is ascribed cannot already be entailed by the information used for revision.

Second, we show how our approach can be used to account for a phenomenon which has been discussed a great deal in the philosophical literature on attitude ascriptions, that of "exportation" (e.g. Quine (1956); Sosa (1970); Kripke (2011)). The cases we will be interested in under this heading may be roughly described as those in which a subject who has a merely descriptive belief can be reported as having a specific one. For instance, consider the following case inspired by Sosa (1970):

Shorty: On our army base we overhear the general say 'I think the shortest man should charge first'. A while later we are in

for (2) (Keshet, 2008, 2010; Schwarz, 2012). Roughly, on this reading the report is true just in case Mary thinks that some person who is *actually* to my left will win. This third reading would also be false in the scenario above: given our setup, Mary only has a general belief that some friend or other will win, not the more specific belief that Bill, Ted, or Ann will win.

conversation with some other members of our platoon. We all know that Shorty is the shortest soldier on base. I say (4):

- (4) The general thinks Shorty should charge first.

An earlier discussion of the connection between exportation and revisionist reports (Blumberg & Holguín (2018)) showed that existing theories of exportation cannot account for revisionist reports. In this paper we take the opposite tack: we show that an adequate account of revisionist reporting—in particular, the theory we develop—can account for examples of exportation like (4).

The paper is structured as follows. §2 introduces further examples of revisionist attitude reports. Our account of revisionist reports is presented in §3. In §4 we discuss an issue concerning the logic of ‘believe’, while in §5 we present a non-redundancy condition on revisionist reporting. §6 shows how our proposal gives a natural treatment of exportation, and §7 distinguishes the phenomenon of revisionist reporting from various phenomena which have been called “*de re* reports”. Finally, §8 concludes by discussing the connections between the relations expressed by attitude verbs in context and the underlying mental relations to content of most interest to philosophers of mind and cognitive science.

2 Scope of the phenomenon

Before presenting our account of revisionist reporting, in this section we give examples to illustrate that revisionist reporting can arise with a range of attitude verbs, and that the reported material can be obtained by employing a variety of inference rules.

To begin with, here is a case involving ‘doubt’:

Doubting Tom: Tom thinks that running a minute mile is virtually impossible. Amy’s friend Ben is training to do this. Tom has never met Ben. Amy and Carol are both friends of Ben and of Tom. Amy says to Carol:

- (5) Tom doubts Ben will be able to break the minute barrier, but I really hope he does.

Our idea will be that (5) is licensed because Amy and Carol can revise the reported belief in light of their shared background information that Ben is trying to run a minute mile.

Here is a similar case inspired by (Pryor, 2004) involving ‘want’:

Music: A music professor says to a dark, crowded auditorium: ‘I want all the doctors here tonight to clap on the third beat...’ My sister Dorothy, who is a doctor, and her husband were listening to their iPods. The husband unplugs his earphones and asks, ‘What’s going on?’ I say:

- (6) He wants Dorothy to clap on the third beat.

Another class of cases that may involve revisionist reporting (but have not been analyzed as such) are discussed by Jerzak (forthcoming). For instance:

Wine Time: You’ve been invited to a party and tasked with bringing the wine. You know that the other guests have reasonably strong preferences, but unfortunately you know basically nothing about wine. At the grocery store you have a choice between a Pinot Noir (from California) and a Malbec (from Argentina), which you take to be equally likely to be the optimal choice for the other guests. If an onlooker (somehow) knows of my situation and of the other guests’ actual preferences (suppose it’s for the Pinot Noir), she may think to herself:

- (7) That person wants to buy the Pinot Noir.

In short, the intuitive gloss on revisionist reports—that speakers can report subjects as believing, doubting or wanting things which are entailed by their beliefs or desires when revised in light of the speaker’s background knowledge—appears to apply in a broad array of cases. The project in most of the remainder of the paper is to develop a unified account of these data.^{5,6}

⁵Although Jerzak (forthcoming) develops an account of desire reports in order to handle (7), his theory cannot be adapted to other attitude reports, e.g. belief ascriptions.

⁶We will focus primarily on propositional attitude verbs, and will not consider in detail whether revisionist reports can also arise with so-called ‘object-directed’ attitude verbs, like ‘admire’, or ‘detest’ (Montague, 2007; Grzankowski, 2012). Suppose that Evan has no idea who Francesca is, but he admires anyone who can play Bach’s *Inventions*. If George knows that Francesca can play the *Inventions*, then it seems he can say ‘Evan admires Francesca’. Such reports might seem to be ‘revisionist’ in our sense, but it’s controversial whether they are. For instance, Drucker (2019) develops a theory of some object-directed attitudes on which they are ‘thin’, so that Evan doesn’t need to know anything at all about Francesca to count as admiring her. On this view, there would not be any need to say that George’s report is revisionist. But even if this is a revisionist report, our theory would extend straightforwardly to cover it. Just as there is an intuitive notion of beliefs or desires being revised in light of new information, there is also an intuitive notion of the objects of a person’s admiration being revised in light of new information, very roughly, to the objects a person would admire were they to know thus and such facts (see below §3.4). So George’s report might be true because of facts about Evan’s dispositions to revise whom he admires—roughly, the fact that Evan is disposed to admire Francesca if he learns that she can play the *Inventions*.

3 The account

In this section, we build our account of revisionist reports by focusing primarily on belief reports. In §3.1 we present a first proposal, which treats revision simply as strengthening. In §3.2 we suggest that reports are revised not by a proposition, but by the true answer to a contextually salient question. In §3.3 we show that strengthening is not flexible enough, and turn to a more general account of revision.

3.1 Strengthening

Our hypothesis about why (1) (‘Pete believes Ann can learn to play tennis in ten lessons’) is acceptable is that speakers can amend the contents they report others as believing by revising the ascriber’s “real” beliefs in light of background information of their own. This idea may strike one as somewhat bizarre on first encounter: why would a linguistic community have a practice of describing others’ minds in such an apparently inaccurate way? As we develop our theory, it will be helpful to have in the background a proposed answer this question. Partly, this proposed answer will help to make sense of how a community could have a practice of this kind; partly it will serve as a narrative within which the components of our theory can be more readily understood. Our basic idea is that revisionist reporting is used as a shorthand to communicate relevant aspects of the ascriber’s dispositions, very roughly, what the ascriber *would* believe if they were a part of our conversation, and knew what we know. In revisionist reports, we leave out the qualification that this is what the subject is merely disposed to believe, and report them as if it is what they *do* believe.⁷

As this description presages, we will develop a semantic account of revisionist reports. Our theory is designed to predict that revisionist reports like (1) can be true. But we will not argue directly for the claim that they can be. We see it simply as a natural null hypothesis. Still, we think our theory would be of interest even if it turned out that these reports are false but acceptable for some kind of pragmatic reason. The acceptability of the reports exhibits a good deal of systematicity, which a fully pragmatic theory must still explain; such a theory would presumably draw on resources similar to those we develop here.⁸

⁷If a question concerning Pete’s belief state is posed in the subjunctive mood, e.g. if someone asks ‘Would Pete believe Ann can play if he knew she was six?’, the response ‘He thinks she can play’ sounds quite bad. This fact might lend credence to the idea that revisionist reports are a type of short-hand. By explicitly using the subjunctive, one draws attention to the fact that Pete’s doxastic state doesn’t actually settle the question, and once this is made salient, answering in the indicative is no longer appropriate.

⁸An anonymous reviewer suggested that reports such as (1) should be considered false,

In presenting our theory, we will employ a possible worlds approach to semantic content, where propositions are modeled as sets of possible worlds, namely the set of worlds where the propositions are true. This approach is controversial (Soames, 1987), but we will use it here for convenience and concreteness. We do not believe that the account of revisionist reporting we will provide can only be cashed out by appealing to possible worlds.

Our semantics for belief reports will take inspiration from Hintikka (1962). On Hintikka’s account, ‘believe’ is given a quantificational semantics involving a lexically-determined accessibility relation. More specifically, ‘believe’ denotes a relation that holds between an agent S and a proposition p just in case every world compatible with what S believes is one in which p is true, i.e. a p -world. On this approach, ‘Bill believes that Ann left’ is true just in case every world compatible with what Bill believes is one where Ann left. We will build on the Hintikka-style account by letting $B_{S,w}$ be the set of worlds compatible with what S believes* in world w , where ‘believe*’ expresses a primitive doxastic relation. One might think heuristically of believe* as the relation theories of belief in cognitive science and the philosophy of mind aim to characterize. As we have already intimated, on our theory believe* will turn out to be distinct from the relation expressed by ‘believe’ in many contexts in natural language.

Our first-run account of revisionist belief reports employs a flatfooted account of “revision” as “strengthening”. The idea is that when we report others’ beliefs we can strengthen them with some of our background information:⁹

Account 1 (to be revised)

$\lceil S$ believes $P \rceil$ is true relative to w and proposition q , just in case:

$$(B_{S,w} \cap q) \subseteq p.$$

On this account, every proposition entailed by $B_{S,w}$ can be reported, since for any q , $(B_{S,w} \cap q) \subseteq B_{S,w}$. Note also that if q is the tautologous proposition, i.e. the set of all worlds, then $(B_{S,w} \cap q) = B_{S,w}$. Thus, the relation expressed by ‘believe’ extends the primitive doxastic relation.¹⁰ On Account

since if someone responded to Jane with ‘Really? But Pete doesn’t even think Ann exists’, Jane would retract her statement. But Sosa’s cases of exportation, mentioned above in §1, are widely agreed to be literally true in context, but they are subject to precisely the same charge: if someone responded to a speaker’s assertion of (4) with ‘Really? But the general doesn’t think Shorty is the shortest man on the base’, the speaker would be inclined to retract my statement. So, if we take the “really” challenge as a diagnostic, then either a fairly broad swathe of acceptable reports should be taken to be false, or we should take examples such as (1) to be true in context. We have opted for the latter approach here.

⁹We let ‘ S ’ range over the names of agents and let ‘ S ’ range over the corresponding agents denoted by ‘ S ’. Similarly, we let ‘ P ’ range over the logical forms of proposition-denoting strings and let ‘ p ’ range over the corresponding propositions denoted by ‘ P ’.

¹⁰We use ‘extend’ to mean ‘weakly extend’ throughout the paper: on this usage, a

1, “revisionist reports” are those that report propositions to which the subject is related by the extended relation but not by the primitive doxastic relation. What exactly this extended relation looks like depends on the relevant value of q . To illustrate, suppose that q is the proposition that Ann is a six-year-old in *Tennis*. Then this account predicts that (1) should be acceptable, since the proposition that Ann can learn to play tennis in ten lessons is entailed by $B_{\text{Pete},w_{\text{@}}} \cap q$.¹¹

One might think that a further constraint on q is justified, namely that it needs to be part of the information that is mutually taken for granted by the discourse participants, i.e. that it must be part of the *common ground* (Stalnaker, 1978, 1984). This constraint does seem to be supported by *Tennis*, and the other cases considered in §2. However, there are other examples which show that such an amendment would be problematic:

Tennis 2: Ann is a six-year-old girl whom Pete, an expert in tennis pedagogy, has never met and whose existence he is unaware of. Pete believes that every six-year-old can learn to play tennis in ten lessons. Jane and Jim wrongly believe that Ann is five, and that Pete believes that every five-year-old can learn to play tennis in ten lessons. Jane says (1) (‘Pete believes Ann can learn to play tennis in ten lessons’).

(1) is true in *Tennis 2*, even though the proposition that Ann is six isn’t part of the common ground. So, although one might be tempted to constrain q by tying it to the common ground, there are reasons against doing so.

3.2 Strengthening by an answer

Although Account 1 captures an important part of our ultimate theory, it needs some refinement. To see this, consider the following case:

Health: Fred is in the hospital and Jim and Jane are concerned about his health. They ask a nurse for information. The nurse replies: ‘The doctor thinks anyone with Fred’s blood type shouldn’t eat gluten. Sorry, that’s all I know. Excuse me, I have

binary relation R extends a binary relation R' iff for every x and every y if $xR'y$, then xRy .

¹¹There is a problem with Account 1 that we won’t be able to resolve fully, namely, that it is difficult to capture Pete’s lack of a conception of Ann in the possible worlds setting. To have something concrete to work with, we will be assuming informally that all this lack of conception comes to is that Ann exists in some but not all of Pete’s belief worlds. But there are reasons to be dissatisfied with this approach. One possible solution appeals to the sort of question-sensitivity briefly discussed in fn. 14.

another call...’ Jim and Jane don’t know what Fred’s blood type is. Jane says:

- (8) The doctor thinks Fred shouldn’t be eating gluten.

(8) is true in *Health*. However, there doesn’t appear to be an appropriate proposition that can play the role of q here. One might suspect that q could simply be the proposition specifying what Fred’s blood type actually is (even though Jim and Jane don’t know it). But this proposal doesn’t get the modal profile of Jane’s assertion right. For instance, suppose that Fred’s blood type is AB at $w_{@}$, and that q is the proposition that Fred’s blood type is AB. Then although Account 1 can explain why (8) is true, it can’t explain why Jane’s utterance would *still* have been true at a world w' where Fred’s blood type had been O, and the doctor at w' thinks that anybody with blood type O should not eat gluten. In short: we want the value of q to be able to somehow vary from world to world. Perhaps this point can be brought out most clearly by considering the behavior of revisionist reports under modals.

Health 2: Fred is in the hospital and Jim and Jane are concerned about his health. They ask a nurse for information. The nurse replies: ‘All I know is that the doctor thinks that anybody with blood type AB should not eat gluten, and that anybody with blood type O should eat gluten. Excuse me, I have another call...’ Jim and Jane know that Fred’s blood type is either AB or O. Jane says:

- (9) The doctor might think Fred shouldn’t eat gluten, it depends what his blood type is.

(9) is perfectly felicitous in *Health 2*. However, Account 1 can’t explain this. For suppose that q is the proposition that Fred’s blood type is O. Then Jane’s utterance would be unacceptable, since she knows that the doctor thinks that anybody with blood type O should eat gluten. So, every world w' compatible with Jane’s knowledge is such that $(B_{\text{Doc},w'} \cap q) \subseteq \{w \mid \text{Fred should eat gluten at } w\}$. Without getting too far into the semantics of epistemic modals, it is plausible that the might-claim in (9) should then be false. On the other hand, suppose that q is the proposition that Fred’s blood type is AB. Jane knows that the doctor thinks that anybody with blood type AB shouldn’t eat gluten. So, every world w' compatible with Jane’s knowledge is such that $(B_{\text{Doc},w'} \cap q) \subseteq \{w \mid \text{Fred shouldn’t eat gluten at } w\}$. But then Jane would be able to say ‘I know the doctor thinks Fred shouldn’t eat gluten’, which she clearly can’t.

To capture (9), we intuitively want to allow the value of q to vary across the worlds compatible with Jane’s knowledge: at worlds compatible with her knowledge where Fred’s blood type is O, we want q to be the proposition that Fred’s blood type is O; and at worlds compatible with Jane’s knowledge where Fred’s blood type is AB, we want q to be the proposition that Fred’s blood type is AB.¹²

In order to handle revisionist readings under modals, we propose that the relevant restriction mechanism should be given by a *question* rather than a proposition. To make this idea precise, we need a semantics for questions. A fairly straightforward account will suffice for our purposes.¹³ We will identify questions with sets of propositions; intuitively: the set of *complete answers* to that question. Complete answers are assumed to be incompatible and exhaustive, so that the set of complete answers corresponds to a partition of the set of possible worlds. To illustrate, (10b) is the denotation of the logical form in (10a) (suppose that the domain only consists of Ann and Mary); and (11b) is the denotation of the logical form in (11a):

- (10) a. Who can learn tennis?
 b. $\{\{w \mid \text{Only Ann can learn at } w\}, \{w \mid \text{Only Mary can learn at } w\},$
 $\{w \mid \text{Both Ann and Mary can learn tennis at } w\},$
 $\{w \mid \text{Nobody can learn tennis at } w\}\}$
- (11) a. Can Ann learn tennis?
 b. $\{\{w \mid \text{Ann can learn at } w\}, \{w \mid \text{Ann can't learn at } w\}\}$

With questions on board, our account is now as follows:

Account 2 (to be revised)

$\ulcorner S$ believes $P \urcorner$ is true relative to w and question Q , just in case:

$$(B_{S,w} \cap Q(w)) \subseteq p.$$

Here we are still supposing that revision be understood as strengthening. On this account, revisionist reports arise when the subject’s belief-set at a world is strengthened with the true answer to the question at that world. Let us suppose that the relevant question, Q , in *Tennis* is *Is Ann a six-*

¹²The argument from revisionist reports embedded under modals is inspired by Dorr & Hawthorne (2013) who posit a question-sensitive restriction mechanism for embedded epistemic modals, as well as Holguín (2018) who considers similar arguments involving attitude ascriptions embedded under probability operators (these arguments are also taken up in (Blumberg & Holguín, 2019, p.27, fn.35)).

¹³The literature on questions is large and growing. Classic works include (Hamblin, 1973) and (Groenendijk & Stokhof, 1984), while (Ciardelli *et al.*, 2016) presents a more recent approach.

year-old?.¹⁴ Then Account 2 predicts that (1) should be true, since the true answer to Q at $w_{@}$ is that Ann is six. But the account also handles (9) so long as we assume that Q is the question *What is Fred's blood type?*: at worlds in which Fred's blood-type is AB, the doctor's belief set will be intersected with the proposition that Fred's blood type is AB; and at worlds in which Fred's blood type is O, the doctor's belief set will be intersected with the proposition that Fred's blood type is O. Thus, at some (but not all) of the worlds w compatible with Jane's knowledge, the revisionist report 'The doctor thinks Fred shouldn't eat gluten' is true at w .¹⁵

Note that if Q is the tautologous question, i.e. $Q(w)$ yields the tautologous proposition for every world w , then $(B_{S,w} \cap Q(w)) = B_{S,w}$. Thus, just as on Account 1, on Account 2 in every context the relation expressed by 'believe' in that context extends the primitive doxastic relation.

We will not here try to resolve how exactly Q gets determined in context. This is obviously a significant issue, and an important lacuna in our account. One would ideally like an explanation of why some questions and not others are available for revisionist reporting. It is an important task for future work to pin down the ways in which this parameter gets determined; without a predictive theory of how the parameter is determined, our account may seem to overgenerate, allowing too many reports to be true. But giving such a theory goes beyond the scope of the present discussion.¹⁶ Our central, more

¹⁴Recently, a number of authors have suggested that an interesting notion of belief is limited by an agent's conceptual repertoire, and that this conceptual repertoire is well represented by a question (Fritz & Lederman (2015), Yalcin (2016) (cf. Yalcin (2011)), Elliott (2019), cf. Kets (2012)). The basic motivation for these approaches is to eliminate certain problematic features of the standard possible-worlds approach to modeling belief. For instance, Yalcin (2016) motivates his account by considering examples like the following (originally from Stalnaker): it may be true that William the Conqueror believed that there wouldn't be war, while being false that he believed that there wouldn't be nuclear war. On the standard approach from Hintikka, the latter would follow from the former, so Yalcin recommends modifying the standard approach.

Although this use of questions may be relevant to some of the examples presented above, e.g. capturing the claim that Pete has no conception of Ann in *Tennis*, it cannot account for the phenomenon that has been the focus of this paper.

¹⁵On a popular approach, questions are represented by functions from worlds to sets of worlds, whose images partition logical space. But it is worth observing that we do not actually require this partitional structure. For instance, in (9) all that matters is that worlds where Fred has blood type AB yield the set of worlds where Fred has blood type AB, and worlds where Fred has blood type O yield the set of worlds where Fred has blood type O. It is of no concern whether worlds in which Fred has a different blood type, e.g. A, yield the set of worlds where Fred has that blood type. Nevertheless, for convenience we will continue to appeal to objects which partition logical space. But it should be kept in mind that (not necessarily partitional) functions from worlds to sets of worlds would also work instead.

¹⁶One might wonder whether Q is identical to the "question under discussion" (QUD), as developed by Roberts (2004, 2012). But we doubt it is. Roberts assumes that the overall strategy of a discourse, namely eliminating non-actual possibilities, is mediated by

modest, aim here is to outline the general shape that an account of revisionist reports would need to take in order to capture the relevant patterns; we leave the important project of developing a predictive theory of how the parameter is determined to future work.¹⁷

3.3 Beyond strengthening: revision

The theory presented so far is simple and captures a broad range of data. But, as we show in this section, there are some examples it doesn't handle well, owing to its overly simplistic identification of revision with strengthening. These examples motivate a more flexible account of revisionist reports.

Consider the following scenario:

Tennis 3: Pete thinks that all six-year-olds can learn to play tennis in ten lessons, but happens to think that Jonny is the only six-year-old. However, if he thought that there were other

raising and answering a set of questions: the QUDs. However, even in a context where both Jane and Jim know that Ann can learn tennis in ten lessons, (1) is acceptable. In this case the question whether Ann can learn to play tennis is *not* a QUD—Jane and Jim *already* take this question to be settled, since they are already convinced that Ann will be able to learn to play.

Several theorists have argued that certain constructions in natural language are question-sensitive, where this question cannot be identified with the QUD, e.g. definite plurals (Malamud, 2012; Križ, 2015). So our proposal is not without precedent in this respect.

¹⁷Account 2 closely resembles the “constraint” account of knowledge reports developed by Holguín (2018). As noted in fn.12, our arguments for question-sensitivity are inspired by similar arguments from Holguín. He presents the following striking data:

Treasure Hunt: Peggy and Pete are each given a map to the buried treasure. Before setting off, they are informed that exactly one of their maps is accurate.

- (12) [We are told that it was Peggy who got the good map, so I say:]
Peggy knows where the treasure is buried.

Holguín maintains that the ascription is licensed because we constrain Peggy's knowledge that if her map is good, then the treasure is buried in location X in light of our knowledge that her map is good. Revisionist reports suggest that Holguín's basic ideas about 'knowledge'-ascriptions have broad application to other reports as well. As we will see, developing this generalization requires some significant modifications of his account. But his insights about knowledge-ascriptions can be seen as underlying our final proposal.

Account 2 is in some respects similar to the “restriction” account proposed by Blumberg & Holguín (2019). On the face of it, this appears to be largely accidental, since there are differences in the underlying phenomena. Blumberg & Holguín are concerned with embedded attitude reports such as ‘If Bill is on a plane to Cuba, then I'm surprised that he left without saying goodbye’. As Blumberg & Holguín note, such constructions only seem to be acceptable in the first person. By contrast, revisionist reports are a distinctly third-person phenomenon, and can occur unembedded.

six-year-olds, he'd think they could learn too. I want to encourage Ann's father, Jim, to sign Ann up for tennis lessons, so in conversation with Jim I say:

- (1) Pete believes Ann can learn to play tennis in ten lessons.
- (13) # Pete believes Ann is an extraterrestrial sent to Earth in order to play tennis.

(1) is acceptable here, but (13) isn't. Our account so far cannot produce this contrast, however, because what the subject believe*s is incompatible with the true answer to what we have been assuming is the contextually salient question. Let us suppose, as above, that Q is the question *Is Ann a six-year-old?*. Then $B_{\text{Pete}, w_{\text{@}}} \cap Q(w_{\text{@}}) = \emptyset$. This means that although our account predicts that (1) should be acceptable, it also predicts that (13) should be acceptable, since the proposition expressed by the complement is a superset of \emptyset .

We suggest that this problem can be handled by appealing to *revised belief states*.¹⁸ The literature on belief revision is enormous, but we will illustrate our general point with a few fairly straightforward ideas (Grove, 1988; Hansson, 2017). We will assume that each person S in w has an associated *plausibility ordering* $\succeq_{S,w}$, which satisfies the axioms of a partial order. Informally, $w' \succeq_{S,w} w''$ just in case it is at least as plausible for S (in w) that all of the events in w' occur than that all of the events in w'' occur. Note that given $w' \succeq_{S,w} w''$, we no longer need to take $B_{S,w}$ to be a primitive notion; instead we can define it to be the set of top-ranked worlds given by $w' \succeq_{S,w} w''$. Then let $\text{Rev}_{S,w}(A)$ denote the revision of $B_{S,w}$ by A . $\text{Rev}_{S,w}(A)$ is a subset of A that meets three conditions: (i) if $A = \emptyset$, $\text{Rev}_{S,w}(A) = \emptyset$; if $A \neq \emptyset$ then (ii.a) $\text{Rev}_{S,w}(A) = A \cap B_{S,w}$, if $A \cap B_{S,w} \neq \emptyset$; and (ii.b) if $A \cap B_{S,w} = \emptyset$, then $\text{Rev}_{S,w}(A)$ is the set of $w' \in A$ such that for every $w'' \in A$ if $w'' \succeq_{S,w} w'$ then $w' \succeq_{S,w} w''$.¹⁹ With revised belief states on board, the account then looks as follows:

Revisionist Belief

$\ulcorner S$ believes $P \urcorner$ is true relative to w and Q just in case:

$$\text{Rev}_{S,w}(Q(w)) \subseteq p.$$

This account predicts that (1) should be acceptable in *Tennis 3* (given the appropriate value for Q). Moreover, it does *not* predict that (13) should

¹⁸One might try to account for why (13) is bad by simply stipulating that enrichment is not allowed when $B_{S,w} \cap Q(w) = \emptyset$. However, this would predict that (1) should be unacceptable in *Tennis 3*, which it isn't.

¹⁹There are other ways of spelling out condition (ii.b). For instance, by appealing to similarity spheres in the style of Lewis's (1973) account of counterfactual conditionals. See (Hansson, 2017) for further discussion.

be acceptable, since Pete’s belief set revised by the proposition that Ann is a six-year-old does not entail that Ann is an extraterrestrial. The account handles all of the data we considered in §3.2 in the same way as earlier versions of it, since when $B_{S,w} \cap Q(w) \neq \emptyset$, $\text{Rev}_{S,w}(Q(w)) = B_{S,w} \cap Q(w)$. In those cases, then, the new entry collapses into the old one. However, Revisionist Belief does have an important new feature by comparison to our earlier accounts. On the earlier accounts, the relation expressed by ‘believe’ extended the primitive doxastic relation. But this doesn’t necessarily hold on Revisionist Belief: S can believe* a proposition p without standing in the relation expressed by ‘believe’ to p . Here we just want to mention this point—we will discuss it in more detail in §4.

In Revisionist Belief, the notion of revision is officially to be understood as a black box. We have introduced a concrete theory of belief-revision described simply in order to illustrate how our theory could be spelled out in detail to make predictions in specific cases. But there are reasons to reject this concrete account of revision. For instance, there are reasons to think that, in general, belief-revision is non-monotonic, i.e. it does not even collapse into strengthening in cases where the proposition by which we revise the ascriber’s beliefs is consistent with what they believe*.²⁰ Such cases might generate counterexamples to the concrete formal theory we have presented above, but they would not be straightforward counterexamples to our official proposal. Provided these examples are still examples of revision, they are consistent with our official theory.

3.4 Other attitudes

The load-bearing notion in our account is that of an attitude being revised in light of some information. In the case of belief, we described one prominent concrete implementation of belief-revision and sketched how our account can take on board this concrete implementation. But much less has been done on revision of other attitudes, so there are fewer ready-made formal tools one

²⁰Jeremy Goodman highlighted the importance of this point to us, using the following case inspired by (Dorr *et al.*, 2014):

Coins: We’re going to conduct an experiment involving flipping a fair coin 100 times. Knowing this, Pete utters ‘I think the coin will land heads in the first ten flips’. Then Pete goes away and doesn’t watch the result of the experiment. We watch and see the first nine flips land tails. Then I say:

- (14) Pete thinks the coin will land heads after 10-19 total flips.

(14) has a true reading in *Coins*. However, this isn’t predicted by our account. Assuming that Q is the question *Did the first nine flips land tails?*, all we’re allowed to report is ‘Pete thinks the tenth flip will land heads’. Intuitively, we want to be able to revise Pete’s beliefs *even though Pete’s beliefs are compatible with the results of the experiment so far*.

might look to in giving a more worked out theory of, e.g. revisionist desire reports. Still, it is clear that the question of how a person’s desires would be revised in light of new information is in as good standing as the question of how their beliefs would be revised in light of new information. So the basic outline of our theory can still be used for these other attitudes. In this section, we will flesh out this idea a little by briefly sketching how one might give an account of some other revisionist attitude reports (in particular, of doubt and desire), by piggy-backing on an account of belief-revision.

Beginning with doubt reports, recall that the sort of ascription we would like to capture is (5) (‘Tom doubts Ben will be able to break the minute barrier, but I really hope he does’). For simplicity, we take ‘S doubts P’ to be synonymous with ‘S believes not-P’. Given our account of revisionist belief reports, this gives us the following:

Revisionist Doubt

‘S doubts P’ is true relative to w and Q just in case:

$$\text{Rev}_{S,w}(Q(w)) \subseteq \neg p$$

It can be checked that this account handles (5).

For desire reports, we will build on so-called “ideal worlds” accounts (von Fintel, 1999; Crnič, 2011). On this approach, ‘want’ is treated similarly to ‘believe’: that is, as a function such that ‘Bill wants Ann to leave’ is true just in case every world that conforms to what Bill desires—every world in Bill’s *desire set*—is one where Ann leaves. Most theorists who endorse an ideal worlds account put constraints on which worlds can appear in a subject’s desire set. It is generally assumed that a subject’s desires generate a preference ordering over possible worlds: for any subject S : $w' >_{S,w} w''$ iff w' is more desirable to S than w'' in w . $>_{S,w}$ is a strict partial order. The idea is that the subject’s desire set is constrained by their beliefs: the subject’s desire set is comprised of all and only their top-ranked *belief worlds*, as ordered by $>_{S,w}$. We will use something a bit more general. Given a set of worlds A , let $>^{S,w}(A)$ denote the top-ranked worlds in A as ordered by $>_{S,w}$. Then our account of revisionist desire reports is the following:

Revisionist Desire

‘S wants P’ is true relative to w and Q just in case:

$$>^{S,w}(\text{Rev}_{S,w}(Q(w))) \subseteq p$$

On this account, revisionist desire reports essentially involve first updating the subject’s beliefs with $Q(w)$, and then reporting the result of calculating the subject’s preferences over this restricted set. It can be shown that this account also handles, e.g. (6) (‘He wants you to clap on the third beat’).²¹

²¹One might have wondered whether revision is needed at all in the case of desire, or

In summary, the account of revisionist belief reports that we have developed generalizes smoothly to some other attitudes where revision can be characterized in terms of belief-revision. For those attitudes where revision can't be characterized in this way, we must appeal to a more abstract notion of revision. Presumably such a notion can be made sense of for just about any attitude. But it is beyond the scope of the present paper to consider formal implementations of revision for these other attitudes.²²

4 Logic of belief

In this section, we return to an issue briefly mentioned in §3.3, namely that it is consistent with Revisionist Belief that S believe* a proposition p and yet not stand in the relation expressed by 'believe' in a given context to p .

To illustrate this point, consider the following scenario:

Tennis 4: Pete thinks that all and only six-year-olds can learn to play tennis in ten lessons, and happens to think, incorrectly, that Ann is five years old. However, if he thought that Ann was six, he'd think that she could learn tennis in ten lessons as well. I want to encourage Ann's father, Jim, to sign Ann up for tennis lessons, so in conversation with Jim I say:

whether strengthening on its own would have sufficed. The following case suggests that revision is needed:

Maltese: Obi loves small dogs, and cares especially about the treatment of the ones in his town of Smallville; Obi wants every Maltese in Smallville to be fed liver regularly. As it happens, Obi incorrectly believes that he is the only Smallvillian to own a Maltese, In conversation with Peng, another Smallvillian with a Maltese, I say (15):

(15) Obi wants you to feed your Maltese liver regularly.

Even though the proposition that Peng has a Maltese is incompatible with Obi's beliefs, (15) is acceptable because the top-ranked worlds in Obi's belief set revised by the proposition that Peng has a Maltese are ones in which Peng feeds Tess liver regularly.

²²Some attitudes that can't be given a Hintikka-style analysis also exhibit revisionist readings. For example, if we overhear Pete say 'I wonder how many six-year-olds will attend tennis camp this summer', then I can truly assert (16):

(16) Pete wonders whether Ann will attend tennis camp.

It is plausible that such cases can be handled by taking on board the semantics for 'wonder' developed by Ciardelli & Roelofsen (2015). The basic idea is to intersect the subject's "inquisitive state" pointwise with $Q(w_{\text{@}})$, i.e. the proposition that Ann is a six-year-old. It can be shown that the resulting set of information states settles the question denoted by the complement in (16), i.e. (16) should be true.

(1) Pete believes Ann can learn to play tennis in ten lessons.

(1) is acceptable in the context of *Tennis 4*. Supposing that Q is *Is Ann a six-year-old?*, this is exactly what our account predicts: Pete’s belief set revised by the information that Ann is six entails that she can play tennis in ten lessons. However, it might be argued that the following report should *also* be true in this context:

(17) Pete believes Ann cannot learn to play tennis in ten lessons.

Pete plausibly believe*s that Ann cannot learn to play tennis, so one might think that (17) should be true in every context. But our account predicts that (17) is *false* in the context specified above: Pete’s belief set revised by the information that Ann is six does not entail that she cannot play tennis in ten lessons.

Revisionist Belief can of course accommodate *some* true readings of (17). For instance, this sentence will be true if uttered in a context where the salient Q is taken to be the trivial question, whose answer is the tautologous proposition. Finding such a true reading is not the problem. The problem is that Revisionist Belief allows that there are also some contexts (e.g. the context in which we have been assuming (1) is uttered) in which (17) is false, and this conflicts with a putative desideratum for an account of revisionist reports: that in every context, speakers be able to truthfully report every proposition entailed by the subject’s belief set, i.e. what the subject believe*s.

One could avoid this problem by amending Revisionist Belief as follows:

Disjunctive Revisionist Belief

$\ulcorner S$ believes $P \urcorner$ is true relative to w and Q just in case:

- (A) $B_{S,w} \subseteq p$, or
- (B) $\text{Rev}_{S,w}(Q(w)) \subseteq p$.

Clause (A) guarantees that in every context the relation expressed by ‘believe’ in that context extends the belief* relation. For instance, since every world in Pete’s belief set in *Tennis 4* is one where Ann can’t learn tennis in ten lessons, (17) is predicted to be true on Disjunctive Revisionist Belief.

But this new account brings with it a new problem. A different desideratum on an account of belief-reports is that in every context, every instance of the following principle is true:

CONJUNCTION: If S believes that P , and S believes that Q , then S believes that P and Q .

Revisionist Belief satisfies this desideratum. But Disjunctive Revisionist Belief does not. On the latter account, in the context described above (1) (‘Pete believes Ann can learn to play tennis in ten lessons’) is true (via clause (B)), and (17) (‘Pete believes Ann cannot learn to play tennis in ten lessons’) is true (via clause (A)), but ‘Pete believes Ann can learn to play tennis in ten lessons and Ann cannot learn to play tennis in ten lessons’ is false (since neither Pete’s belief set nor his revised belief set entails the proposition expressed by the complement).

The choice between Disjunctive Revisionist Belief and Revisionist Belief turns on a subtle question about which of these two general constraints is more important for a semantic theory. Should the theory respect the intuitive idea that, in every context, what a subject believe*s is in the extension of the relation expressed by ‘believe’ in that context? Or should it rather ensure that in every context, every instance of CONJUNCTION is true, guaranteeing an interesting ‘logic of “belief”’? The general questions raised by how to trade these two desiderata against one another are difficult to resolve conclusively. But our own inclination—which we won’t argue for here—is to take logical principles like CONJUNCTION to play a more important role in constraining semantic theory than judgments about what ought to be in the extension of the relation expressed by ‘believe’ in a given context. Thus, although we think Disjunctive Revisionist Belief is worth exploring further, for the remainder of the paper we will use Revisionist Belief as our official account. In the next section, we will continue to explore the consequences of adopting this ‘logic-first’ approach to revisionist reports.

5 Non-redundancy

In this section, we discuss a striking feature of revisionist reports that the account developed in §3 does not explain. We will present a pragmatic theory of this phenomenon, and argue that it is to be expected given what we take to be the overall use and purpose of revisionist reporting.

We have suggested that the acceptability of, e.g. (1) (‘Pete believes Ann can learn to play tennis in ten lessons’) can be explained by the fact that the prejacent in the report is entailed by facts about how the ascriber’s “real” beliefs are revised in light of information the conversational participants have. However, a speaker cannot truly report *any* proposition of this kind. Consider the following report, uttered in *Tennis* (repeated from §1):

(18) # Pete believes Ann is a six-year-old.

This result cannot be predicted by our account to this point. If the salient question is *Is Ann six-years old?*, then of course Pete’s beliefs revised by the

true answer to this question entails that Ann is a six year-old. The problem is not that our theory predicts no false reading of this sentence; in many contexts it will be false. But we have been working under the assumption that *Tennis* suggests a context in which the proposition that Ann is a six year-old is entailed by the true answer to the salient question. This assumption strongly suggests at least that there would be salient readings of (18) on which it is acceptable. But no such readings seem available.

One response to this observation would be to revise our semantic theory to predict that, after all, (18) is false in a relevant context in which (1) is true. An alternative response—which we favor, for reasons we will discuss shortly—is to add a further condition on the felicity of assertions, which predicts that while (18) is true, it is infelicitous. We propose that revisionist reports carry a pragmatic *non-redundancy condition*, namely, that the ascriber’s mental state cannot be redundant in deriving the ascribed attitude, or, more formally, that the proposition in which belief is ascribed cannot already be entailed by the information used for revision. This idea promises to explain why (18) is unacceptable: the proposition in which belief is ascribed is already entailed by the true answer to the relevant question. However, we need to take some care in giving a precise formulation of the condition. The issue is that, like (18), (19a) and (19b) are equally unacceptable in the context of *Tennis*:

- (19) a. # Pete believes Ann is a six-year-old and can learn to play tennis in ten lessons.
- b. # Pete believes Ann can learn to play tennis in ten lessons and is no older than seven.

In each of these cases, the proposition expressed by the entire complement clause is not entailed by the information used for revision. However, in each case the complement clause is a conjunction and one of its conjuncts *is* entailed by the information used for revision. We take this to show that the non-redundancy condition doesn’t just apply at the level of the *entire* complement clause, but rather relative to *parts* of the complement clause as well. Both (19a) and (19b) are unacceptable because one of the conjuncts of the complement clause in each case is entailed by the information used for revision. To make this precise, we will appeal to the notion of a *context set*, denoted C : a set of worlds representing the information that is mutually taken for granted by conversational participants (Stalnaker, 1978, 1984). We say that a sentence P is entailed in context by a proposition $Q(w)$ just in case the semantic value of the sentence in context is entailed by the proposition. Our proposal is then the following:

Non-redundancy condition: The report $\lceil S \text{ believes } P \rceil$ can be felicitously asserted only if (i) at every world w in $C+S$ *believes*

P : $B_{S,w} \subseteq \llbracket P \rrbracket$, or (ii) no part P' of P is entailed in context by $Q(w)$, where $Q(w)$ is the information used for revision.

This necessary condition on the acceptability of reports can be satisfied in two ways. If at every world w in the context set updated by the report, $B_{S,w} \subseteq \llbracket P \rrbracket$, then the non-redundancy constraint need not apply.²³ This part of the necessary condition ensures that we do not prohibit reporting propositions that the subject believes*, even if these propositions are entailed by $Q(w)$. The second way of satisfying the necessary condition has bite only when this first one is not satisfied, i.e. when what is reported is not entailed by the subject's belief*s, and we are engaged in revisionist reporting.²⁴

This non-redundancy condition fits neatly into the just-so story we have been offering about the function of revisionist reports. We have suggested that revisionist reporting is used as a kind of shorthand to communicate relevant aspects of the ascriber's dispositions, very roughly, what the ascriber *would* believe if they were a part of our conversation, and knew what we know. But if the aim of the shorthand is to allow speakers to quickly report facts about the ascriber's dispositions, the rules of the shorthand should not allow them to make reports which are uninformative about those dispositions.

²³ $C+S$ believes P denotes the set of worlds w in C such that $\lceil S$ believes $P \rceil$ is true in w . We use $C+S$ believes P here rather than just C since we want to allow speakers to be able to add information regarding what S believes* to the context set without triggering infelicity.

²⁴As it stands, our non-redundancy condition only applies to *sentential* parts of P , so it can't account for the infelicity of reports such as (20):

(20) Pete thinks Ann is a six-year-old who can play tennis.

However, such cases can be captured if we adopt a more complex account of non-redundancy that appeals to a notion of *semantic equivalence relative to $Q(w)$* (Fox, 2007; Schlenker, 2008; Mayr & Romoli, 2016). Let us say that P and Q are equivalent relative to $Q(w)$ if P and Q have the same truth-value at every $w \in Q(w)$ (Singh, 2011). We can then replace condition (ii) in the non-redundancy condition with the following:

(ii)' For no simplification P' of P are P' and P equivalent relative to $Q(w)$, where P' is a simplification of P if P' can be derived from P by replacing nodes in P with their subconstituents.

'Ann can play tennis' is a simplification of 'Ann is a six-year-old who can play tennis'. Moreover, relative to any context which carries the information that Ann is six, they will be equivalent. Thus, the revised non-redundancy condition predicts that (20) should be unacceptable, as required.

One would ideally like our non-redundancy condition to be derivable from more general pragmatic principles. One possibility is that the question used for revision, Q , is syntactically realized, and appears as a constituent in the logical form of the complement of belief reports. Using mainstream accounts of non-redundancy, e.g. those cited above, it is plausible that a condition extensionally equivalent to (ii)' can be recovered. However, we must leave a more detailed investigation of this proposal for future work. Thanks to Matt Mandelkern for helpful discussion on these points.

Without the non-redundancy condition, we would expect ‘Everyone on earth believes Ann is six’ to be acceptable in the context of *Tennis*; this use of the shorthand would be entirely uninformative about people’s dispositions to believe.

We have framed non-redundancy as a constraint on felicity rather than truth. We have opted for this pragmatic version of the condition, as opposed to a semantic one, for reasons connected to the logic of belief. It is plausible that in every context, every instance of the following principle is true:

LIMITED CLOSURE: If S believes that A is F and knows that anything which is F is G , then S believes that A is G .

However, we can spell out a variant on our original case which would give a counterexample to this principle if the non-redundancy condition were taken to be a constraint on truth, as opposed to felicity. Suppose that Pete believes that the only creatures in the universe which can learn to play tennis by taking lessons are mammals, so that (21) is true:

- (21) Pete believes that anything which can learn to play tennis in ten lessons is a mammal.

Suppose, moreover, that this sentence is true in a context in which (1) is true. The problem is that the following sentence is unacceptable in describing the scenario of *Tennis*:

- (22) # Pete believes Ann is a mammal.

The prejacent of this sentence is redundant (the information used for revision plausibly entails that Ann is a six-year-old mammal), so our non-redundancy condition, whether taken to be a constraint on truth, or on felicity, would make the correct predictions here.²⁵ The problem is that, if the non-redundancy condition were a semantic constraint on truth, then (22) would be predicted to be false, and LIMITED CLOSURE would have an instance which is false in this context. By contrast, if non-redundancy is merely a constraint on felicity, we can preserve LIMITED CLOSURE while still explaining the unacceptability of (22).²⁶

²⁵Here we are assuming that Pete does not believe that all six-year-old entities can learn to play tennis, so that throughout our discussion of the example, there has been an implicit restriction to six-year-old mammals (and presumably six-year-old humans).

²⁶Some of our informants judge (23) to be unacceptable in the following context (though some also judge it to be felicitous, thus the %’ symbol):

Horses: We are at the races and overhear Pete say ‘Only Banjax and Tracer have a chance of winning—the rest of the horses are injured’. We then learn that Banjax has been given a tranquilizer, so he certainly won’t win. Sally

6 Exportation

In this section, we connect our theory of revisionist reports to discussions of exportation. We show how the proposal we've presented accounts straightforwardly for some classic cases that fall under this heading, and consider the relationship between our theory of the phenomenon and discussions of "singular thought".

Consider (4) ('The general thinks Shorty should charge first') in the context of *Shorty* from §1. The relevant proposition that the general believe*s (i.e. that is entailed by $B_{G,w}$, where "G" is for "the general") is the proposition that the shortest man should go first. If we assume that the salient question Q in the context of *Shorty* is *Who is the shortest man?*, then the proposition that Shorty should go first is entailed by $\text{Rev}_{G,w}(Q(w))$, and (4) is thus predicted to be true on our account.

Here is a second example, also from Sosa:

Pyro: [C]onsider...the case of a prominent citizen of Metropolis who suffers from pyromania. Impelled by his pyromania, he disguises himself from time to time in order to start some fires, and becomes known to the community as "The Metropolis Pyromaniac."

Let us use 'the Metropolis Pyromaniac' as a definite description, and let us suppose that the name of this person is 'Fiery'. Consider now Fiery's sometime collaborator. She hears the police chief on TV say that the Metropolis

takes Pete to have good betting instincts, so she comes up to us and asks 'Do either of you know who Pete thinks will win the race?'. I reply:

(23) % Pete thinks Tracer will win.

This might be taken to motivate a more constrained theory than we have presented here. More specifically, it might suggest that the correct theory should be restricted to the behavior of determiners, e.g. 'every', 'some', and should not work, as ours does, by allowing revision of the beliefs of the ascriber wholesale. The rough idea is that determiner phrases, e.g. 'every six-year-old' can be replaced by names, e.g. 'Ann', under suitable conditions. We think that such an account deserves to be explored further. However, we also think that there are challenges for theories in this mould. For one thing, it is unclear how such an account could handle Holguín's (12) ('Peggy knows where the treasure is buried') from fn.17, since this seems to involve revising Peggy's conditional knowledge rather than replacing a determiner. For another, it seems we can say 'Mary thinks James or Jeff will win the race' in the context of *Running*, but 'James or Jeff' isn't standardly taken to be a name, or even a plural term. Finally, it is unclear to us how such an account could explain a fact we discuss at the end of the next section, namely: why attitude verbs such as 'dream' and 'suppose' allow *de re* readings of embedded determiner-phrases, but do not allow what are in our terms "revisionist readings", e.g. (28) ('John dreamed that Bill was a bald man in his 90s') below.

Pyromaniac has struck again, leaving, as is his wont, the usual signs of his work. Passing by Fiery's house the collaborator may say to Fiery's wife, congratulating her on the extraordinary devilry of her spouse:

(24) Everyone knows that Fiery laid last night's fire.

The obvious variant of Revisionist Belief for knowledge rather than belief predicts that (24) is true if Q is the question *Who is the Metropolis Pyromaniac?*. As Sosa also in effect observed (and is predicted by our theory), this kind of report is highly context-sensitive. For later in the day, Fiery's wife can reassure their son:

(25) No one knows that Fiery laid last night's fire.

The acceptability of (25) can be explained by our account so long as in this context the true answer to the relevant question Q does not entail that Fiery is the Metropolis Pyromaniac.

Our account thus has the resources to produce intuitive results in canonical cases discussed in the literature on exportation.

Discussions of the examples we've considered are usually taken to bear on questions about the character of people's mental states, and in particular on the nature of what are sometimes called "singular thoughts" or "*de re* beliefs". On one standard gloss an individual has a *de re* belief about x if they believe that x is F , for some F . It is thought that, when a person has such a *de re* belief, it follows that there is something which they believe to be F . Philosophers have taken an interest in these examples in large part because of their connection to a set of issues in the philosophy of mind surrounding the nature of singular thought.

But if our account is correct, these examples do not bear on what we understand to be the main issue in this debate. It is true in the first case, for instance, that if the police believe that the Metropolis Pyromaniac did it, they believe that Fiery did. But this is not because at all worlds consistent with what the police believe*, Fiery did it.²⁷ Rather, 'the police

²⁷Many discussions of singular thought employ the notion of a "singular proposition", i.e. a proposition that is *about* an individual in a special sense of "aboutness". It's not clear that this notion makes good sense in the possible worlds framework in which we're working, since in this framework the proposition that if Fiery is a pyromaniac, then Fiery is a pyromaniac is identical to the proposition that if it is raining then it is raining, but the former might look to be a singular proposition *about* Fiery whereas the latter undoubtedly is not. Still, if one favored a different theory of propositions, one could rephrase the point we've just made using the notion of a singular proposition: exportation holds in this instance, even though the police do not believe* a singular proposition about Fiery.

believe Fiery did it' is true in context because revisionist reporting is admissible here. We can truly report the ascriber as believing a proposition even though they do not believe* that proposition. While this tells us something important about the relation of believing (i.e. the relation expressed by 'believe' in context), the correct conclusion to draw seems to be not that we need a new theory of belief or singular thought in the philosophy of mind. Rather, we should conclude that the question about the nature of singular thought in the philosophy of mind should not be understood in terms of the relation of believing; it should instead be understood in terms of the relation of believing*. We will return to the importance of this distinction between the relation of believing and the relation of believing* in the conclusion.²⁸

Our account of these cases extends to a range of data concerning predicates, which have not always been recognized as related to the canonical cases of exportation. Consider the following example from Sudo (2014):

[Mary says to herself that John is the same denomination as Bill. She thinks Bill may be Protestant or Presbyterian or Catholic. But we know that Bill is Catholic.]

(26) Mary thinks John is Catholic.

Our theory of revisionist reporting handles this example straightforwardly. Provided Q is the question *Is Bill Catholic?*, then since Mary thinks that John is the same denomination as Bill, the theory predicts the relevant true reading here.

To repeat, we have not attempted to present a theory of what questions tend to be supplied for which utterances. In the absence of such a theory, our account of these data is not fully predictive: flexibility is inversely related to strength. But by showing that these reports can fit into a single semantic framework, we have at least made a step toward a more predictive theory.

We end this subsection by presenting some rather striking evidence related to exportation that fits well with our just-so story about the function of revisionist reports, giving further support to the idea that examples of export should be handled by the same mechanism that handles revisionist reports. The evidence involves an intriguing contrast briefly mentioned by Percus & Sauerland (2003, fn.19, 238). The contrast can be brought out by considering the following scenarios, and the reports that follow them:

Reviewer Belief: John sent his paper to a journal a few weeks ago. He has no idea who his reviewer is, but we overhear him

²⁸Goodman (2018) and Openshaw (2018) also consider ascriptions such as (1), and argue that these reports can be true without, e.g. Pete having any singular thoughts about Ann. We are obviously sympathetic to this claim, but unlike Goodman and Openshaw, we have provided a general semantic theory which vindicates this conclusion.

say ‘Everyone who reviews for that journal is really ancient, and they’re all bald as well. So, the person reviewing my paper is a bald man in his 90s’. As it happens, Bill—a man in his 20s—is actually John’s reviewer. Knowing this, I say (27):

(27) John thinks that Bill is a bald man in his 90s.

Reviewer Dream: John sent his paper to a journal a few weeks ago. He has no idea who his reviewer is, but we overhear him say ‘I had the most vivid dream last night. I was watching the person reviewing my paper; he was a bald man in his 90s’. As it happens, Bill—a man in his 20s— is actually John’s reviewer. Knowing this, I say (28):

(28) # John dreamed that Bill was a bald man in his 90s.

The puzzle is that (27) is acceptable but (28) is not. That is, exportation seems to be possible with some attitude verbs, e.g. ‘believe’, but not others, e.g. ‘dream’. Also consider a similar example with ‘suppose’:

Reviewer Suppose: John sent his paper to a journal a few weeks ago. He has no idea who is reviewer is, but we overhear him say ‘Suppose older, bald reviewers are more lenient. And suppose that the person reviewing my paper is a bald man in his 90s’. As it happens, Bill—a man in his 20s— is actually John’s reviewer. Knowing this, I say (29):

(29) # John is supposing that Bill is a bald man in his 90s.

We submit that our account of revisionist reporting can explain these contrasts. First, observe that ‘dream’ and ‘suppose’ do not allow for revisionist reports in general:

Tennis Dream: Ann is a six-year-old girl whom Pete, an expert in tennis pedagogy, has never met and whose existence he is unaware of. We overhear Pete say ‘Last night I had a dream about tennis; I dreamed that every six-year-old could learn to play tennis in ten lessons’. I say:

(30) # Pete dreamed that Ann could learn to play tennis in ten lessons.

That is, only certain attitude verbs allow for revisionist reporting. It seems that attitude verbs which denote mental states or events which are closely tied to conscious experience (dreaming, supposing), as opposed to states or events associated more closely with outward performance (saying) or which an agent can be in while asleep (hoping, wanting, knowing, believing) do not admit of revisionist reporting. This generalization can be explained by reflecting on what, according to us, the overall communicative goal or aim of revisionist reports is. If, as we have suggested, revisionist reports are used as a shorthand to communicate relevant aspects of the ascriber's dispositions, then the possibility of revisionist reporting may be connected to the character of the underlying attitude being reported. It is comparatively natural to extend a practice of reporting states like believing, knowing or wanting, to encompass a person's dispositions to be in various further related states. By contrast it is not natural to extend a practice of reporting a conscious episode a person is undergoing, to encompass what conscious episodes they are disposed to undergo in other circumstances. A great deal more would have to be said to make this idea precise. But we think it is plausible that the restriction on revisionist reporting for reports featuring these verbs can be naturally seen to flow from the purpose of revisionist reporting as we understand it.

7 *De re* attitude reports

In this section we discuss two further examples which illustrate some limits on the empirical ambitions of the theory we have presented in this paper.²⁹ The first—the transparent/opaque ambiguity—we believe is not an example of revisionist reporting. The second—“double vision”—can be thought of as an example of revisionist reporting when the relevant cases are spelled out in one way, but should not be thought of as such an example if they are spelled out in a different, equally natural, way.

There is a general point which will underlie our discussion of these examples: that the categorization of examples connected to attitude reports is theory-dependent. We have already seen that, from the perspective of our theory, revisionist reporting and export are importantly “the same phenomenon”, while double vision and export are not. We think this is an independently plausible, attractive pattern of predictions. But from the perspective of other theories, things might look quite different. For instance, the theory of Percus & Sauerland (2003) naturally suggests an account on which export and double vision are handled by the same mechanism. Our goal here is simply to say a little about how the land lies as seen from the vantage point of our theory. We think the picture which emerges is independently

²⁹This section can be skipped without loss of continuity.

interesting. But in sketching this picture, we are not presenting an argument for our theory. We think that judgments about which phenomena are pretheoretically understood to be the same or different should not on their own carry serious weight in favor of or against a given theory.

Our first example concerns an ambiguity which arises when determiners occur in the scope of intensional operators. Such sentences are ambiguous between what we will call “opaque” and “transparent” readings, a distinction which is also often described as a distinction between *de dicto* and *de re* reports, or “narrow scope” and “wide scope” readings. For instance, consider (31):

[Ralph thinks that Mary lost the race, but she actually won it.]

(31) Ralph believes that the winner lost.

This sentence has a salient true and a salient false reading in this scenario. The true reading is the “transparent”, “*de re*” or “wide-scope” reading; the false one is the “opaque”, “*de dicto*” or “narrow-scope” reading.

Our theory cannot account for the true reading of this sentence. If Ralph’s beliefs are revised by a proposition which does not entail that Mary won, then the report will be straightforwardly false. But if Ralph’s beliefs are revised by a proposition which entails that Mary won, then his revised belief set will no longer entail that she lost.³⁰

The second type of example that we wish to discuss are so-called cases of “double vision”. Here the situation is a bit less clear; our theory predicts that there are at least two different kinds of cases. The basic example comes from Quine (1956). Ralph sees Ortcutt wandering by the docks in shady circumstances, and concludes he is a spy, so (32) is acceptable:

(32) Ralph believes that Ortcutt is a spy.

However Ralph also sees Ortcutt on TV as the mayor his town, and he thinks the mayor of his town is not a spy. Thus, (33) is also acceptable:

(33) Ralph believes that Ortcutt is not a spy.

So as not to use the overloaded term “*de re*”, we will call the phenomenon associated with this example *double vision*.

³⁰Although we won’t expound on this point here, it is important to note that the standard treatment of this kind of example, e.g. appealing to world pronouns in the syntax (Percus, 2000; Keshet, 2008; Schwarz, 2012; Elbourne, 2013), can’t be extended to handle revisionist reporting.

Some theories of the belief* relation will predict that, in the case as described here, Ralph both believe*s the proposition that Ortcutt is a spy, and also believe*s the proposition that Ortcutt is not a spy. If the case is understood in this way, our theory has nothing to say; our theory predicts that the semantic phenomenon exhibited in these examples is distinct from that of exportation and revisionist reporting.

But some will think that this is not the appropriate way of understanding what Ralph believe*s in this case. They will hold, first, that Ralph believe*s that the person he saw by the docks is a spy, and that he believe*s that the person he saw on TV is not a spy. But they will also hold, second, that either Ralph does not believe* that Ortcutt is a spy, or he does not believe* that Ortcutt is not a spy, or that he does not believe*s either of these claims. In this case, our theory does have something to say. And indeed, whatever one thinks of this way of describing Quine's original case, there are variants on the case that make this view of Ralph's belief*s quite natural. Suppose for instance that Alfredo tells Ralph reliably that the person Alfredo saw by the docks was a spy, and Ralph comes to believe this proposition. Suppose later that Beth tells Ralph reliably that the mayor of Beth's town is not a spy, and Ralph comes to believe this proposition. If we know that Ortcutt is the person Alfredo saw and the mayor of Beth's town, and Ralph has no other salient ways of thinking about this individual, then both (32) and (33) may be used truly in our context. But in this new case it is extremely natural to say that Ralph believe*s neither the proposition that Ortcutt is a spy, nor the proposition that Ortcutt is not a spy.

When the case is regimented in this way our theory can accommodate the basic data, and thus makes it seem as though the data are of a piece with revisionist reporting and exportation. In a context where Q is the question *Who did Alfredo see at the docks?*, (32) will be predicted to be true. In a context where Q is the question *Who is the mayor of Beth's town?*, then (33) will be true. But as this description already suggests, there is one important wrinkle: our theory cannot predict that these two sentences are both true in the same context. Making that prediction would require revising Ralph's beliefs by propositions that would render his beliefs inconsistent; if Ralph's beliefs were revised by both pieces of information at once, it is natural to think he would give up one of the beliefs he held initially, e.g. that he would stop believing that the person Alfredo saw was a spy.

This means that, even in the case where Ralph acquires his beliefs on the basis of testimony, our theory cannot be used to predict the truth of:

(34) Ralph believes that Ortcutt is a spy and Ortcutt is not a spy.

Since the verb is the relevantly context-sensitive element, and there is only one occurrence of it, we can only use a single Q in revising/strengthening

this report.

A variety of responses are available; we'll just discuss two. One is that this report is not true in the relevant scenario. Our own judgments about this example are unstable, and we don't think it would be a grave problem not to predict a true reading of this sentence. This position is compatible with holding that (34) could be true when the underlying situation with belief* is understood according to the first regimentation above. A second response is to hope that, given a mechanism for predicting the truth of (32) and (33) when the underlying situation with belief* is understood as in the first regimentation, then the combination of revision and that other mechanism could be used to predict the truth of (34). We don't have a preference among these responses. As we have said, our goal here has simply been to document the ways in which our theory divides up the territory.

8 Conclusion

In this concluding section, we consider how revisionist reports make trouble for a simplistic view of the connection between the relations expressed by attitude ascriptions and the relations of most interest to philosophers of mind and cognitive science.

Let us call the sorts of mental relations to contents that are of most interest to philosophers of mind and cognitive science *M-relations*, and the relations expressed by attitude verbs in context *L-relations*. If revisionist reports can be true, then M-relations are not always straightforwardly reflected in natural language, as L-relations: in certain contexts, 'believe' and other attitude verbs pick out relations that are distinct from any M-relation. The relation expressed by 'believe' in context can always be obtained from M-relations together with certain facts about the background context, but it is not in general identical to any M-relation.

Philosophers of language and mind have considered a diverse array of possible divergences between M-relations and L-relations. To illustrate, we'll just discuss two prominent examples. First, expressivists in the tradition of Gibbard (1990) have provided a unified semantic theory of attitude ascriptions according to which attitude reports which are similar in their surface form express quite different facts about the mind. On such theories, in both 'John believes murder is wrong' and 'John believes cats are animals', 'believes' expresses the same relation: essentially, a relation between individuals and sets of pairs of possible worlds and hyperplans. But the truth of these ascriptions is underwritten by different M-relations. On a flatfooted way of interpreting Gibbard's model theory, the truth of the first is underwritten by John's standing in a plan-like M-relation to a set of hyperplans, while the

truth of the second is underwritten by his standing in a belief-like M-relation to an object represented by sets of worlds. ‘Believe’ univocally expresses a relation between people and sets of pairs of worlds and hyperplans, but very different M-relations underwrite the holding of this L-relation.³¹

Second, philosophers and semanticists have also countenanced divergences between L-relations and M-relations as part of their treatment of issues related to “Frege’s puzzle” (Frege, 1892). For instance, on an influential “neo-Russellian” semantics for attitude reports (associated with Salmon (1986)), there is an underlying M-relation to mental representations, and one bears the L-relation expressed by an attitude verb to a proposition if and only if one bears the relevant M-relation to *some* mental representation which expresses that proposition. On this theory, too, M- and L-relations would be distinct.³²

If revisionist reports can be true, we must countenance a further form of divergence between M-relations and L-relations. On our proposal for handling these reports, the L-relation expressed by an attitude verb in context (for instance what ‘believe’ expresses in context) is obtained by revising the M-relation (what the subject believe*s) by the answer to a salient question. A precedent for this kind of divergence is the style of contextualism about ‘know’ developed in Lewis (1996). Lewis takes as primitive a relation that holds between a person and a proposition when the proposition is entailed by the person’s evidence, and takes ‘know’ to express in context a relation that holds between a person and a proposition whenever the proposition is entailed by the conjunction of the person’s evidence and a contextually salient proposition. Contextualism of this form is typically married with a background view that what we tend to express by ‘know’ is not a deep or important aspect of the mind. If our examples of revisionist reports are true, it would be natural to take this moral to extend quite generally. For it would not just be ‘know’ that expresses a relation which diverges from the most interesting relevant M-relation; the L-relations we express by uses of many

³¹Similar points might be made about expressivist treatments of modals, e.g. Yalcin (2007), Yalcin (2012), Rothschild (2012).

³²Arguably this kind of approach could be traced back to Kaplan (1968), who earlier suggested a more limited version of this kind of divergence, restricted to so-called “*de re*” reports. Crimmins & Perry (1989), Richard (1990), Crimmins (1992) and now Goodman & Lederman (2019) all present contextualist variations on this basic idea.

Examples of this general pattern could be multiplied. To mention just two more salient, striking ones: first, on the account of Richard (1990) an utterance of ‘S believes P’ may be true provided S believe*s Q and Q is sufficiently similar to P, where what is “sufficiently similar” may vary depending on the context (cf. also Bach (1997); Graff Fara (2013) makes a similar move, but for different reasons, for desire-reports). Second, and more proximately related to the topic of this paper, some theorists have observed that reports of the form ‘S said P’ can be sensitive to the background information of conversational participants (von Stechow & Zimmerman, 2005; Brasoveanu & Farkas, 2007; Saebo, 2013), in ways broadly reminiscent of revisionist reports.

other attitude verbs in many other contexts would also diverge in striking ways from the M-relations of most interest to the philosophy of mind.³³

Still, our account is consistent with the ideas, first, that people often talk about M-relations and, second, that there is a systematic relationship between L-relations and M-relations. On the first point: on our account, in a context where the question supplied by context is the null question (whose true answer is the tautologous proposition), the L-relation expressed by ‘believe’ is the M-relation of believing*. Indeed, for all we have said, ‘default’ or ‘normal’ contexts supply this null question, so that in such normal contexts attitude verbs do express the relevant M-relations. On the second point: our theory systematically characterizes L-relations in terms of M-relations, given the specification of a contextually salient question. The rules of language allow revisionist reporting, in quite surprising ways. But that does not mean there are no rules at all.

References

- Aloni, Maria. 2005. Individual Concepts in Modal Predicate Logic. *Journal of Philosophical Logic*, **34**(1), 1–64.
- Bach, Kent. 1997. Do Belief Reports Report Beliefs? *Pacific Philosophical Quarterly*, **78**(3), 215–241.
- Blumberg, Kyle, & Holguín, Ben. 2018. Ultra-Liberal Attitude Reports. *Philosophical Studies*, **175**(8), 2043–2062.
- Blumberg, Kyle, & Holguín, Ben. 2019. Embedded Attitudes. *Journal of Semantics*, **36**(3), 377–406.
- Brasoveanu, Adrian, & Farkas, Donka. 2007. *Say* reports, assertion events and meaning dimensions.
- Ciardelli, Ivano, Groenendijk, Jeroen, & Roelofsen, Floris. 2016. *Inquisitive semantics*. Lecture notes.
- Ciardelli, Ivano A., & Roelofsen, Floris. 2015. Inquisitive Dynamic Epistemic Logic. *Synthese*, **192**(6), 1643–1687.
- Crimmins, Mark. 1992. *Talk About Beliefs*. MIT Press.

³³Jeremy Goodman and Matt Mandelkern have independently suggested to us that introspection principles might provide an potential heuristic for distinguishing between M-relations and L-relations. Consider the claim that if a person believes that p , they believe that they believe that p . If revisionist reports are true, it seems we can give straightforward counterexamples to this principle. For instance, Pete believes that Ann can learn to play tennis in ten lessons, although plausibly he does not believe that he believes that Ann can learn to play tennis in ten lessons. But this counterexample, which relies on revisionist reports, does not bear on an analogous principle formulated using M-relations, namely, that if a person believe*s a proposition, then they believe* that they believe* it. A proponent of this second principle, for M-relations, might then use failures of introspection for the L-relation expressed by ‘believe’ in context to diagnose which reports are revisionist reports. We ourselves are not particularly attracted even to the M-relation version of introspection principles, but we think the example provides a helpful illustration of how the contrast between M-relations and L-relations must be handled with care. For discussion of some related points in connection to Lewis’s contextualism, see Holliday (2015) and Salow (2016).

- Crimmins, Mark, & Perry, John. 1989. The Prince and the Phone Booth: Reporting Puzzling Beliefs. *Journal of Philosophy*, **86**(12), 685.
- Crnić, Luka. 2011. *Getting even*. Ph.D. thesis, MIT.
- Dorr, Cian, & Hawthorne, John. 2013. Embedding Epistemic Modals. *Mind*, **122**(488), 867–914.
- Dorr, Cian, Goodman, Jeremy, & Hawthorne, John. 2014. Knowing Against the Odds. *Philosophical Studies*, **170**(2), 277–287.
- Drucker, Daniel. 2019. Policy Externalism. *Philosophy and Phenomenological Research*, **98**(2), 261–285.
- Elbourne, Paul. 2013. *Definite Descriptions*. OUP Oxford.
- Elliott, Edward. 2019. *Unawareness and Implicit Belief: Possible Worlds Models of Informational Content*. Unpublished MS.
- Fox, Danny. 2007. *Presupposition and Implicature in Compositional Semantics*. London: Palgrave Macmillan UK. Chap. Free Choice and the Theory of Scalar Implicatures, pages 71–120.
- Frege, Gottlob. 1892. Über Sinn und Bedeutung. *Zeitschrift für Philosophie und philosophische Kritik*, **100**, 25–50. Translated as “On Sense and Reference” by Max Black in *Philosophical Review*, 57.3: 209–230, 1948.
- Fritz, Peter, & Lederman, Harvey. 2015. Standard State Space Models of Unawareness. *Theoretical Aspects of Rationality and Knowledge*, **15**.
- Gibbard, Allan. 1990. *Wise Choices, Apt Feelings: A Theory of Normative Judgment*. Harvard University Press.
- Goodman, Jeremy, & Lederman, Harvey. 2019. Perspectivism. *Noûs*.
- Goodman, Rachel. 2018. Do Acquaintance Theorists Have an Attitude Problem? *Australasian Journal of Philosophy*, **96**(1), 67–86.
- Graff Fara, Delia. 2013. Specifying Desires. *Noûs*, **47**(2), 250–272.
- Groenendijk, Jeroen, & Stokhof, Martin. 1984. *Studies on the semantics of questions and the pragmatics of answers*. Ph.D. thesis, Institute for Logic, Language & Computation (ILLC), University of Amsterdam.
- Grove, Adam. 1988. Two Modellings for Theory Change. *Journal of Philosophical Logic*, **17**(2), 157–170.
- Grzankowski, Alex. 2012. Not All Attitudes Are Propositional. *European Journal of Philosophy*, 374–391.
- Hamblin, Charles L. 1973. Questions in Montague English. *Foundations of Language*, **10**(1), 41–53.
- Hansson, Sven Ove. 2017. Logic of Belief Revision. In: Zalta, Edward N. (ed), *The Stanford Encyclopedia of Philosophy*, winter 2017 edn. Metaphysics Research Lab, Stanford University.
- Hawthorne, John, & Manley, David. 2012. *The Reference Book*. Oxford University Press.
- Hintikka, Jaakko. 1962. *Knowledge and Belief*. Ithaca: Cornell University Press.
- Holguín, Ben. 2018. *Knowledge by Constraint*. Unpublished manuscript.
- Holliday, Wesley H. 2015. Epistemic Closure and Epistemic Logic I: Relevant Alternatives and Subjunctivism. *Journal of Philosophical Logic*, **44**(1), 1–62.
- Jerzak, Ethan. forthcoming. Two Ways to Want? *Journal of Philosophy*, 1–37.
- Kaplan, David. 1968. Quantifying In. *Synthese*, **19**(1–2), 178–214.
- Keshet, Ezra. 2008. *Good Intentions: Paving Two Roads to a Theory of the De re/De dicto Distinction: Massachusetts Institute of Technology Ph. D.* Ph.D. thesis, dissertation.
- Keshet, Ezra. 2010. Split Intensionality: A New Scope Theory of de Re and de Dicto. *Linguistics and Philosophy*, **33**(4), 251–283.
- Kets, Willemien. 2012. *Bounded reasoning and higher-order uncertainty*. Unpublished MS.
- Kripke, Saul A. 2011. Unrestricted exportation and some morals for the philosophy of language. *Pages 322–350 of: Philosophical Troubles*. Oxford University Press.
- Križ, Manuel. 2015. Homogeneity, Non-Maximality, and all. *Journal of Semantics*, **33**(3), 493–539.
- Lewis, David. 1979. Attitudes de Dicto and de Se. *Philosophical Review*, **88**(4), 513–543.

- Lewis, David K. 1973. *Counterfactuals*. Blackwell.
- Lewis, David K. 1996. Elusive Knowledge. *Australasian Journal of Philosophy*, **74**(4), 549–567.
- Maier, Emar. 2009. Presupposing Acquaintance: A Unified Semantics for de Dicto, de Re and de Se Belief Reports. *Linguistics and Philosophy*, **32**(5), 429–474.
- Malamud, Sophia A. 2012. The meaning of plural definites: A decision-theoretic approach. *Semantics and Pragmatics*, **5**(3), 1–58.
- Mayr, Clemens, & Romoli, Jacopo. 2016. A puzzle for theories of redundancy: Exhaustification, incrementality, and the notion of local context. *Semantics and Pragmatics*, **9**(7), 1–48.
- Montague, Michelle. 2007. Against Propositionalism. *Noûs*, **41**(3), 503–518.
- Openshaw, James. 2018. Singular Thoughts and de Re Attitude Reports. *Mind and Language*, **33**(4), 415–437.
- Percus, Orin. 2000. Constraints on some other variables in syntax. *Natural language semantics*, **8**(3), 173–229.
- Percus, Orin, & Sauerland, Uli. 2003. On the LFs of attitude reports. In: *Proceedings of the Conference 'Sub7 - Sinn Und Bedeutung'*. University of Konstanz.
- Pryor, Jim. 2004. An Epistemic Theory of Acquaintance. In: *UCSB Conference on Content and Concepts*. 2004.
- Quine, W. V. 1956. Quantifiers and Propositional Attitudes. *Journal of Philosophy*, **53**(5), 177–187.
- Recanati, François. 2012. *Mental Files*. Oxford University Press.
- Richard, Mark. 1990. *Propositional attitudes: An essay on thoughts and how we ascribe them*. Cambridge University Press.
- Roberts, Craige. 2004. *Context in Dynamic Interpretation*. John Wiley & Sons, Ltd. Chap. 9, pages 197–220.
- Roberts, Craige. 2012. Information structure in discourse: Towards an integrated formal theory of pragmatics. *Semantics and Pragmatics*, **5**(6), 1–69.
- Rothschild, Daniel. 2012. Expressing Credences. *Proceedings of the Aristotelian Society*, **112**(1pt1), 99–114.
- Saebo, K. J. 2013. Reports of Specific Indefinites. *Journal of Semantics*, **30**(3), 267–314.
- Salmon, Nathan U. 1986. *Frege's Puzzle*. Ridgeview.
- Salow, Bernhard. 2016. Lewis on Iterated Knowledge. *Philosophical Studies*, **173**(6), 1571–1590.
- Schlenker, Philippe. 2008. Be Articulate: A Pragmatic Theory of Presupposition Projection. *Theoretical Linguistics*, **34**(3), 157–212.
- Schwager, Magdalena. 2009. Speaking of qualities. *Pages 395–412 of: Semantics and Linguistic Theory*, vol. 19.
- Schwarz, Florian. 2012. Situation pronouns in determiner phrases. *Natural Language Semantics*, **20**(4), 431–475.
- Singh, Raj. 2011. Maximize Presupposition! and local contexts. *Natural Language Semantics*, **19**(2), 149–168.
- Soames, Scott. 1987. Direct Reference, Propositional Attitudes, and Semantic Content. *Philosophical Topics*, **15**(1), 47–87.
- Sosa, Ernest. 1970. Propositional Attitudes De Dicto and De Re. *Journal of Philosophy*, **67**(21), 883–896.
- Stalnaker, Robert. 1978. Assertion. *Syntax and Semantics (New York Academic Press)*, **9**, 315–332.
- Stalnaker, Robert. 1984. *Inquiry*. Cambridge University Press.
- Sudo, Yasutada. 2014. On *De Re* predicates. In: Santana-LaBarge, Robert E. (ed), *Proceedings of the 31st West Coast Conference on Formal Linguistics*.
- von Fintel, Kai. 1999. NPI Licensing, Strawson Entailment, and Context Dependency. *Journal of Semantics*, **16**(2), 97–148.
- von Stechow, Arnim, & Zimmerman, Ede. 2005. A problem for a compositional account of de re attitudes. In: *Reference and Quantification: The Partee Effect*. CSLI Publications.
- Yalcin, Seth. 2007. Epistemic Modals. *Mind*, **116**(464), 983–1026.
- Yalcin, Seth. 2011. Nonfactualism About Epistemic Modality. In: Egan, Andy, & Weatherson, B. (eds), *Epistemic Modality*. Oxford University Press.
- Yalcin, Seth. 2012. Bayesian Expressivism. *Proceedings of the Aristotelian Society*, **112**(2pt2), 123–160.
- Yalcin, Seth. 2016. Belief as Question-Sensitive. *Philosophy and Phenomenological Research*, 23–47.