Philos Stud (2017) 174:2819–2843
DOI 10.1007/s11098-016-0812-2

Russellianism unencumbered

Mark McCullagh[1]

Published online: 7 November 2016

Abstract Richard Heck, Jr. has recently (in "Intuition and the substitution argument," *Analytical Philosophy* 2014) reconfigured the debate over Russellianism about proper names. Sidestepping the usual argument, which concerns "intuitions" about substitutions within "that"-clauses, he proposes a new argument based on the claims that (i) beliefs are individuated by their psychological roles and (ii) ordinary language has belief-specifying locutions that reflect that individuation. Focusing on (ii) I argue that contrary to what Heck claims, "that"-clause ascriptions are not the only candidates. In fact there are much better candidates: ascriptions involving direct quotations. I explain how the proposal is novel (it avoids the usual problems with such ascriptions) and how it answers the requirements of Heck's argument. More broadly what Heck's argument brings out is the diversity of resources ordinary language has for specifying beliefs; a defense of Russellianism needn't rest entirely on claims about "that"-clause ascriptions.

Keywords Belief ascription · Russellianism · Quotation

1 Introduction

Objections to Russellianism about proper names have often been based on the apparent failure of intersubstitutability of co-denoting proper names in "that"-clauses, particularly in belief-ascribing sentences. Richard Heck, Jr (2014) has recently reconfigured this debate. He seems to concede—in response largely to the work of Jennifer Saul (1997, 1998, 2007)—that Russellianism can deflect one version of this "substitution argument." The version that Saul says has been "the

✉ Mark McCullagh
mmcculla@uoguelph.ca

[1] Department of Philosophy, University of Guelph, Guelph, ON N1G 2W1, Canada

main argument" (1997, 107) against Russellianism is based on intuitions we have about the truth conditions of such sentences. However, Heck claims that there is an argument that involves no appeal to such intuitions. Following what he argues was Frege's (1892) own line of thought, Heck argues on the basis of a claim about *beliefs*: that they are individuated at least partly by their psychological roles. Since this is so, Heck argues, we should expect to find in ordinary language some sentences that specify beliefs in a way that reflects that individuation. (After all, we're *interested* in beliefs.) For Heck it follows that "that"-clause belief ascriptions do this work. If that is right then Russellianism is false, for a specification of that form cannot do that work if a proper name in it contributes to its truth condition only (as the Russellian claims) by picking out the thing named.

Heck thus sidesteps one of the Russellians' principal moves. We are all used to reading their explanations of why our intuitions about the truth conditions of belief-ascribing sentences are mistaken: we confuse them with implicatures, or with other asserted contents, etc. Heck's argument makes this sort of explanation irrelevant. This significantly changes the debate, since much of the discussion of Russellianism has concerned the plausibility of such explanations.

What does this mean for the overall dialectic? Surprisingly, in one way it makes things easier for Russellians: since intuitions about truth conditions are no longer at issue, Russellians are no longer obliged to explain them away. On the other hand, the argument presents a new challenge to the Russellian. The new challenge is to identify sentences of ordinary English that literally specify beliefs in a way that reflects their individuation: surely there are such. "That"-clause belief ascriptions don't do this, if the Russellian is right about their semantics, so the Russellian must identify other sentences that do.

Heck's intervention thus widens the field of inquiry: the topic now is the belief-specifying resources of ordinary language as a whole, rather than the belief-specifying powers only of "that"-clause ascriptions. The reason this matters is that ordinary language also offers belief-specifying locutions of other kinds. My proposal will be that so-called "quotational" belief ascriptions—surprisingly not discussed by Heck—are remarkably good candidates for meeting Heck's challenge. They let us literally specify beliefs in a way that is more sensitive to their psychological role than are specifications in terms of Russellian propositional contents. There are different purposes for which we talk about beliefs, and one type of locution may serve a given purpose better than another. On some occasions our primary aim is to specify a belief's truth condition; on others it is to specify its psychological role; often it is a mixture of both purposes. Since equivalence of truth condition doesn't correspond to equivalence of psychological role, there's simply no reason to expect there to be one ordinary locution that *best* serves each of these purposes. So a proposal on which the latter purpose is best served by an ordinary language locution other than "that"-clause ascriptions shouldn't come as a surprise.

Here again the dialectical setting matters, since quotational belief ascriptions have a long history in the debate over Russellianism. In the past, what mattered was their role in explaining away our non-Russellian intuitions about "that"-clause belief ascriptions. Thus they were proposed either as *analyses* of such ascriptions or as expressing what is *communicated*—suggested, asserted, conveyed,

conversationally implicated—by them. Both sorts of proposals have their problems. But the role they play in the reconfigured debate is a new one: the proposal is not that quotational ascriptions are *analyses of*, or *suggested by*, "that"-clause ascriptions but simply that their semantic contents do belief-specifying work that the semantic contents of "that"-clause ascriptions cannot do. In light of this the Russellian might even thank Heck for pointing us towards the proper role of these locutions in an overall defense of Russellianism. For the problems facing past proposals concerning these locutions are not problems for the proposal I'll be making.

More broadly, once the debate is broadened—as I think Heck's argument shows it should be—beyond its traditional focus on "that"-clause ascriptions we see them as just one among a variety of locutions we use to talk about beliefs. And since belief-talk answers to different purposes that on occasion pull it in different directions, it's unsurprising that some locutions serve a given purpose better than others do. We see differences of this sort even between the two sorts of locutions that are at issue in this paper. My focus, however, is not on those differences in themselves but on how they relate to the requirements of Heck's argument.

2 Heck's argument: the specification challenge

Heck's argument rests on the claim that beliefs are individuated according to their psychological roles. If belief a plays a role in causing some action or judgement, and belief b does not play that role, then $a \neq b$: they are distinct beliefs. On its own, of course, this claim follows just from the indiscernibility of identicals. What makes it interesting are its interactions with other claims. For there are cases in which its antecedent is true and Russellians will assign a and b the same content. Heck illustrates this using a scenario based on the familiar Superman story (Siegel and Shuster 1938):

> Lois is sitting at her desk at the offices of the *Daily Planet* and talking on the phone to Clark, who is about to cover a school board meeting. A few minutes later, Jimmy Olson pokes his head through the door and tells her, "Someone just said they saw Superman at the school board meeting!" Excitedly, Lois hangs up on Clark, jumps up from her desk, and rushes out the door. (10)

Heck argues that "when Jimmy said what he did, Lois acquired a new belief, one she did not previously have: It is Lois's having this new belief that explains why she jumped up from her desk and rushed to the school board meeting" (12). While speaking with Clark, who is at the meeting, she has one belief (about his being at the meeting)—"the C-belief"—and upon getting the news from Jimmy she acquires another, "the S-belief." (Heck introduces the terms as neutral vocabulary in order not to presuppose the truth of any "that"-clause belief ascriptions.) Heck is explicit about his reliance on the aforementioned claim about the individuation of beliefs: the C-belief and the S-belief "are different mental states, since they play different roles in Lois's psychology" (13).

What does Russellianism say about this scenario? In this debate "Russellianism" is a claim about semantics. It is the claim that attitude-ascribing verbs ("believes," "thinks" etc.) relate thinkers to the semantic values of their complement "that"-clauses, and the value of such a clause "depends, so far as any names that might occur in [it] are concerned, only upon their reference[s]" (2). On this view, Lois's having *either* of the two beliefs makes *both* of the following sentences true:

bel-Clark Lois believes that Clark Kent is at the school board meeting.
bel-Superman Lois believes that Superman is at the school board meeting.

To many this has seemed an obviously false implication of the view. One way for the Russellian to defend it is to say that the strongly felt difference between these "that"-clause belief ascriptions—henceforth, **t-ascriptions**—is not a difference in truth condition (Salmon 1986, 1989; Soames 1987a, b, 2002, 2005; Berg 1988). It is a difference, rather, in what utterances of those sentences *assert* or *suggest* or *implicate*. (I'll use "convey" as a placeholder term for these; the contrast is with what is literally expressed.)

In her contributions to this debate—to which Heck's article is proximally a reply—Saul (1997, 1998, 2007) has argued that in this respect t-ascriptions are, contrary to the traditional take, on an equal footing with many "simple" sentences that are not about mental states or speech acts (e.g. "Clark Kent went into the phone booth and Superman came out"). Saul characterizes both types of case as ones in which our "intuitions" about the truth conditions of sentences conflict with what Russellianism (in the case of t-ascriptions) and standard semantic treatments (in the case of simple sentences) tell us about those truth conditions. She proposes, very plausibly, that the conflict in the latter case be explained by appeal to what utterances of these sentences *convey* rather than what they literally express, and infers that the same sort of explanation works in the former case as well. If it does, then the conflict with "intuitions" is not a problem for Russellianism.

Heck does not directly dispute this argument of Saul's. Rather he sidesteps it, for *his* argument against Russellianism involves no claim about intuitions. What it involves instead is a claim about ordinary speakers' need for belief-specifying locutions that reflect the individuation of beliefs. To continue in terms of the working example, Heck's claim is that ordinary speakers, given their interest in explaining Lois's actions, must have some way to *distinctly specify* the C-belief and the S-belief. Heck now challenges the Russellian: What are these specifications? "How is it possible, in English (and other natural languages), to express that Lois has the S-belief as opposed to the C-belief" (19)? I'll call this the **Specification Challenge**. Heck will claim that t-ascriptions are the only candidates for meeting it. Since they can't do the requisite specifying if Russellianism is true, Heck concludes that Russellianism is not true. After elaborating some aspects of the Challenge and explaining how Heck presses it against Russellianism, I will propose other locutions as candidates for meeting it.

Several points about the Specification Challenge are worth making explicit.

One is that since the individuation of beliefs by their psychological roles is extremely fine-grained, it's doubtful that any natural-language locutions can *fully*

reflect that individuation. Consider a scenario somewhat different from the one Heck describes. Suppose that—as we might put it—Lois mistakenly takes there to be two Supermans, having received information on different occasions about "Superman" along with some evidence that there were two heroes with that name; and suppose that—as she might put it, taking there to be two—she believes of each that he can fly. (This is a variant of Kripke's (1979) famous Paderewski example.) These are different beliefs: all it takes for them to be distinct is some difference in their psychological roles, and their places in what Lois takes to be different evidentiary chains is such a difference. Yet our ordinary practice of using t-ascriptions seems to support applying the one t-ascription "Lois believes that Superman can fly" to each (at least when they are considered in isolation).

This doesn't show that the Specification Challenge cannot be met, just that it shouldn't be interpreted as requiring too much. Heck is careful to note this. He says that he is not after a specification that is true *only* of the S-belief (in our case); what he is after is a specification that *distinguishes* it from the C-belief. All this requires is that there be "a sentence for whose truth Lois's having the S-belief is sufficient, but for whose truth her having the C-belief is not sufficient" (2014, 20 n. 28). And he is "happy to concede that it is not utterly implausible that English should lack the resources needed to express suitable explanations of people's behavior in some cases" (25) (including, arguably, cases such as the one described in the previous paragraph).

On the Russellians' view, as we noted, neither *bel-Clark* nor *bel-Superman* does the distinguishing work Heck requires. So even though the Specification Challenge does not, absurdly, require what we might call *maximal* specifications—ones that can in principle be satisfied by only one belief each—it does require something that t-ascriptions don't deliver, if Russellianism is true.

Another point is that the Challenge concerns the truth conditions of sentences rather than what utterances of them convey. If the Challenge admitted of satisfaction in the latter way then we would be back on familiar ground: the Russellian could invoke one of their accounts (cited above) according to which there is some *conveyed* content that is true of the S-belief but not of the C-belief.

But the Challenge requires sentences that do the specifying work literally. Suppose a Russellian offered an account of how *bel-Superman* conveys some content *A* and *bel-Clark* conveys some different content *B*, saying that this explains the felt difference between *bel-Superman* and *bel-Clark*. Heck would simply press the Challenge again, asking: Are these *conveyed* contents ones that can be literally expressed by sentences "in English"? And if so, do they specify beliefs in a way that meets the Challenge? If the answers to both these questions are Yes, then the Specification Challenge would be met by those sentences, making the specificatory powers of *bel-Superman* and *bel-Clark* irrelevant (to the question of whether the Challenge can be met). But if the answer to either question is No, then the Challenge is not met.

Regarding the first of those two questions: there are two ways in which there might fail to be sentences "in English" that literally express the contents *A* and *B*. One is that expressing those contents requires vocabulary that no English speaker has. But another, more relevant in this debate, is that expressing those contents

requires special technical vocabulary from philosophy or psychology, which Heck would not count as being "in English" given his focus on "ordinary speakers" (19, 27). The point of the reference to ordinary speakers is to exclude belief-specifications that involve some special technical vocabulary, even if specifications of that sort are—as is possible—*better* at specifying beliefs than any specifications using ordinary speakers' locutions. The relevance of this is that (as both Braun (1998, 567–568) and Heck (2014, 22) object, following Rieber (1994) and Schiffer (1992, 513–518)) several Russellian accounts hold that the contents that are conveyed by ordinary t-ascriptions are ones whose expression requires technical philosophical concepts such as a "propositional guise" or a "way of grasping a proposition," related in belief by a ternary relation (designated "BEL" by Salmon (1986, 117)). Saul (1998), who takes herself to be developing Salmon's and Soames's accounts in their 1980s works, claims that t-ascriptions implicate contents about "guises" (383–386). No sentence about "BEL" or "guises" or "ways of grasping" satisfies the Challenge.

(To exclude sentences involving such vocabulary as candidates for meeting the Challenge is not to denigrate them. Heck acknowledges the possibility of "introduc[ing] new vocabulary that allowed us to talk about modes of presentation or mental files (or what have you), so that the difference between Lois's having the C-belief and her having the S-belief could be expressed by well-informed philosophers and cognitive scientists" (19).[1])

Suppose then that the answer to the first of the questions is Yes. That is, suppose that (to continue with our example) *bel-Clark* and *bel-Superman convey* different contents *A* and *B* that *can* be expressed in ordinary English—say, with sentences **sent-A** and **sent-B**. The Challenge now concerns *these* sentences. Do *sent-A* and *sent-B* specify beliefs in a way that adequately reflects facts about their invididuation? If the answer is No, then we are no further ahead, as concerns the Challenge, than we were with *bel-Clark* and *bel-Superman*; we would merely have put in their place another pair of sentences that fail to meet it.

Scott Soames's more recent (2002, 2005) account of t-ascriptions stands in this relation to the Challenge—that is, it answers Yes to the first question and No to the second. On that account, assertive utterances of t-ascriptions often result in the *assertion* of contents that are augmentations of the literally expressed contents of the t-ascriptions. Since the augmentation can go differently even for t-ascriptions whose "that"-clauses pick out the same Russellian proposition, Soames can plausibly claim to have explained how they can generate differing intuitions about their truth conditions. Now these asserted contents are ones that on Soames's view can themselves be literally expressed in t-ascriptions that involve no special technical

[1] For example, Pryor (2016), citing Heck (2012) as a like-minded precursor, has recently proposed a "mental graphs" vocabulary for modelling the contents of beliefs and other mental states, arguing that "Frege Problems in thought are best modeled using graph-theoretic machinery" (1). This vocabulary lets us distinctly specify even members of a large set of beliefs that not only have the same Russellian contents but fit the same t-ascriptions—situations akin to the two-Superman-beliefs scenario I sketched above. The illustrative scenario Pryor uses is one in which "Alice believes her local baseball team has several Bobs on it. In fact, they only have one Bob...[and] the memories, information, and attitudes she associates with 'each' of them is the same—or as close to this as possible" (3).

vocabulary (2002, 210–35); so the answer to the first question is Yes. But since the original problem was the specificatory inadequacy of t-ascriptions as the Russellian construes them, the answer to the second question is No: the augmented contents are themselves specified in the very way that is problematically coarse-grained, from the point of view of the Challenge. Thus the overall position fails to meet the Specification Challenge, even if it does offer an explanation of non-Russellian "intuitions" about the truth conditions of t-ascriptions. (This illustrates another way in which Heck's argument changes the debate.)

We should consider some specifications that Heck himself uses, that are not t-ascriptions; Heck's implicit attitude towards them may tell us more about his conception of the Challenge. Given how he uses the Challenge against Russellianism (which I explain below), it is somewhat surprising that he himself uses ordinary English to tell his readers what the S-belief and the C-belief are:

> [The S-belief is] the belief that is responsible for Lois's getting up from her desk and rushing to the school board meeting. ...[The C-belief] is the belief that does not, in these circumstances, cause her to rush to the meeting but that might, in other circumstances, be responsible for her going to the meeting if, say, she needed to find another reporter to help her cover a breaking story. (13)

Presumably these descriptions Heck uses do not themselves meet the Challenge, otherwise the fuss over t-ascriptions would be beside the point. But *why* not? They pick out the beliefs literally, not by what they convey; and they involve no special technical vocabulary. Perhaps the relevant failing is their context-dependence: one might reasonably maintain that they fail to specify beliefs independently of context of utterance, with their reliance on phrases like "these circumstances."

But I think there is another reason why Heck might not want to count these specifications as meeting the Challenge, for there is an impressive feature that t-ascriptions evidently have and Heck's descriptions evidently lack. T-ascriptions *encode truth conditions*, whereas the sorts of partial specifications of psychological role that Heck uses in telling us what the S-belief and the C-belief are, do not come anywhere close to doing so.

When I say that a specification "encodes truth conditions" I mean that anyone who understands the specification comes thereby to know which truth condition a belief must have in order to fit it. (By "knowing which," in such cases, I mean knowing a biconditional whose right side expresses the condition, e.g. "the S-belief is true if and only if..."). Just knowing that a belief has the truth condition *of that very belief*, for example, doesn't count as "knowing which" truth condition it has.) In this sense t-clause specifications do encode truth conditions. Stephen Schiffer pressed just this point in objecting to Davidson's (1968) account of indirect discourse:

> If in uttering "Sam said that flounders snore" you assert truly that Sam said that flounders snore...[then in understanding your report] I know, especially, that what he said is true just in case flounders snore. (Schiffer 1987, 133)

Schiffer's objection was that you can understand the Davidsonian paraphrase *without* knowing that; the paraphrase fails because unlike the "said that" statement, it doesn't encode a truth condition.[2]

Truth-condition-encoding locutions are at one end of a range, at the other end of which are specifications such as "the belief that caused Lois to get up from her desk." *Just* knowing that the S-belief caused Lois to get up from her desk isn't knowing anything about its truth condition. (At best one could infer a truth condition if, in addition, one had exhaustive knowledge of Lois's situation and all her other beliefs and desires. The same could be said, however, of specifications such as "the first belief Fred formed on Thursday.") Heck doesn't count specifications useless in that way as satisfying the requirements of the Specification Challenge.[3]

Finally it is clear that the Challenge requires that there be *sentences* that do the specifying work, not just *utterances* of sentences. Again, this is needed in order to rule out trivializing proposals, e.g. that utterances of (in the simplest case) "that belief" count as specifications meeting the Challenge. (This is not to say that Heck thinks of belief ascriptions as having the same truth condition in every context of utterance. Sentences of the form "S believes that p" arguably vary in how de-re like they are, from one context of utterance to another (3), concerning this or that component of the "that"-clause. But Heck's simplifying assumption is that this sort of variation doesn't affect the question at issue. It'll be good enough to have a pair of sentences that in many common contexts will do the distinguishing, in terms of their truth conditions—as uttered in those contexts—between two beliefs.)

What the Specification Challenge requires, then, in sum, are: sentences composed of ordinary vocabulary, that literally specify beliefs sufficiently finely to distinguish between beliefs in relatively non-exotic cases such as Lois's, and which do better than supplying *no* information about truth conditions.

3 Options for responding to the challenge

Someone who accepts Heck's framing of the debate has three options: (1) accept that the Challenge cannot be met; (2) maintain the Challenge can be met by t-ascriptions, so Russellianism is false; (3) maintain that the Challenge can be met by sentences other than t-ascriptions, so it has no implication concerning their semantics.

[2] Note that the problem here arises most vividly on the *hearer's* side not the speaker's. Presumably a speaker would assert a t-ascription only if they do know the truth condition of the complement sentence in the "that"-clause. Schiffer's point is that on Davidson's account, a *hearer's* understanding of the ascription *doesn't* require that they know that, contrary to what we know about what is involved in understanding t-ascriptions. I'm grateful to a referee for querying this point.

[3] Why not exclude *all* locutions that aren't truth-condition-encoding, not the just the ones at the opposite extreme? Because that woud invite the rejoinder that the Specification Challenge is stacking the deck in favor of t-ascriptions or locutions that simply *state* truth conditions. And it would be unmotivated, inasmuch as the thought behind the Challenge is that it's *psychological roles* that belief specifications need to capture.

Heck rejects option (1), writing, "I have no argument to offer...other than to observe that languages grow and change to serve the expressive and communicative needs of their users, and so that it would be extremely surprising if there simply was, in ordinary English, no way to express such a significant difference as the one we are discussing" (19). For this reason I too put no stock in option (1).

Option (2) is Heck's position. As we have noted, Russellians cannot agree that the t-ascriptions *bel-Clark* and *bel-Superman* meet the Challenge, for their view entails that those sentences "must be true or false together" (18). So if those t-ascriptions *do* meet the Challenge then Russellianism is false.

Option (3) gets surprisingly little consideration in Heck's discussion. He repeatedly tells us that the only candidates he can see for meeting the Challenge, as concerns the Lois scenario, are t-ascriptions:

> it is so difficult to see...what other sentence, if not [*bel-Superman*], might express that Lois has the S-belief specifically...(21).
>
> I have no idea what sentence other than [*bel-Superman*] might express that Lois has the S-belief specifically. (22, see also 27).
>
> If it is not ["Lois remained at her desk because she did not believe that Superman was at the school board meeting"] that expresses that Lois remained at her desk because she lacked the S-belief, then what sentence of English would? (24)

Indeed at one point he implies that option (3) is somehow ruled out *in principle*: "the semantic doctrines constitutive of Russellianism make it impossible in principle for any sentence of English ever to distinguish between Lois's having the C-belief and her having the S-belief" (27).

But despite Heck's mention of Russellianism in that passage, it's worth emphasizing that the *only* implications of Russellianism that have been at issue in the debate, and in Heck's paper, are ones concerning the semantics of t-ascriptions. We know that if Russellianism is true then *those* sentences cannot satisfy the Challenge. But no similar implication concerning other sentences has been established.

This focus on t-ascriptions should give us pause. Should we expect there to be only one type of sentence that can be used to specify beliefs? Beliefs have many aspects: they have truth conditions; they have psychological roles; they have histories; there are facts about how widely shared a belief-content is (with beliefs held by others), etc. In light of this it would actually be astonishing if there were only one type of sentence that could be used to specify them. We should expect, rather, that there are different kinds of sentences that do the specifying work that ordinary speakers are interested in doing—different kinds of candidates, therefore, for meeting the Specification Challenge. (That there is this variety doesn't show that all such candidates *succeed* in meeting the Challenge; but it does show that they should each be considered. Below I discuss further the issue of the diversity of purposes in belief specification.)

In what follows I will argue that there are belief-specifications of another kind, which succeed where t-ascriptions (by the Russellian's lights) fail. Indeed, they are

ones that have a history in the debate over Russellianism, which makes it even more surprising that Heck doesn't consider them as candidates for meeting the Challenge.

4 Q-ascriptions

4.1 The proposal

What I have in mind are specifications of beliefs in terms of sentences that believers can use to express them in assertions. For example, we can specify the S-belief in this way using the description "the belief that Lois can express by asserting 'Superman is at the school board meeting' " and the C-belief using the description "the belief that Lois can express by asserting 'Clark Kent is at the school board meeting'." I'll call these **q-ascriptions**, since they involve direct quotations. (As I explain below (§ 6.2), for good reason they seldom occur in ordinary speech or writing in this pure form. More usual are mixed t-ascriptive and q-ascriptive forms—"Lois believes that 'Superman' is at the meeting"—or uses of directly quoted sentences, rather than "that"-clauses, as complements to "believes.")

The very idea of relating beliefs to assertions expressing them is of course not new. The relation invoked in a q-ascription does seem close to the relations that Salmon and Soames (in earlier writings) have proposed when trying to articulate what the utterance of a t-ascription conveys. For both writers, it's something about *words*. Soames says that t-ascriptions "often suggest corresponding relations to certain sentences" (Soames 1987a, 68)[4] and Salmon writes, in the same vein:

> there is an established practice of using such a sentence as ["Lois Lane believes that Clark Kent is Superman"], which contains the uninteresting proposition that Lois Lane believes the singular proposition about Superman that he is him, to convey furthermore that Lois Lane agrees to this proposition *when she takes it in the way it is presented to her by the very sentence "Clark Kent is Superman"* (assuming she understands this sentence). (Salmon 1989, 249–250)

But these views don't directly address the Specification Challenge, for two reasons. One is that they are views about what t-ascriptions *convey* rather than what they literally express. Another is that when we look at Salmon's account especially, we find that it works with notions that aren't explicitly part of ordinary discourse about beliefs. (This shows up in the quoted passage with his reference to "the way it is presented;" the notion of a "way" in which a Russellian proposition is "presented" to a subject is not a notion for which we have an ordinary locution.) As was discussed above, no such account meets the Specification Challenge.

So the proposal I'm making is a new one, even though it does involve one idea familiar from previous defenses of Russellianism. I turn now to its elaboration and defense. By way of elaboration it must be explained how q-ascriptions work: what

[4] Soames's (2002, 2005) view is rather different, as discussed above, p. 6.

their truth conditions are, *how* they specify beliefs. For there are reasons to worry that they don't work. One is encouraged by a widely-endorsed argument; another is a worry about what the "can" comes to; another is a worry about whether "expressing" is a relation we can use; finally there is a worry about whether q-ascriptions are any use in specifying beliefs of nonlinguistic creatures. By way of defense it must be explained how specifications of this sort meet the specific requirements imposed by the Challenge (§ 5), and how this proposal is reflected in our ordinary belief-specifying discourse (§ 6).

4.2 General worries

4.2.1 The Bigelow argument

Belief specifications using direct quotation have a long history in the debate over the semantics of t-ascriptions—and thus in the debate over Russellianism, since that debate has mostly focused on the semantics of t-ascriptions. Carnap (1947) famously proposed specifications of that sort as *analyses* of t-ascriptions. Nowadays they are generally considered unsuited to that role. The argument for this (due to John Bigelow) is simple. Suppose that Marian believes that Robin wins (and in case it needs adding, that Marian doesn't believe that Robin loses). Now consider the sentence

M Marian believes that Robin wins.

In English, sentence *M* is true. Now consider a variant of English differing only in that "wins" means what "loses" means in English.[5] In this language, sentence *M* is false, given our supposition. Consider now the proposed quotational paraphrases of *M* in English and the variant language. In either language the paraphrase is a sentence along the lines of "Marian stands in relation *R* to 'Robin wins'." Given that the only difference between English and the variant language is over the meaning of "wins," there is no difference between that sentence's meanings in the two languages since "wins" is not used, only quoted. Thus while *M* is true in English and false in the variant language, the quotational paraphrase of *M* has the same truth value in both languages. So the quotational paraphrase of *M* fails about as badly as a proposed analysis can fail (Bigelow 1978, 109–111; Cresswell 1980, 24–25).[6] Note that the details of relation *R* don't matter to the argument. It could be

[5] The argument here requires that whatever it is a direct quotation refers to, it is something that can be interpreted differently: it is words that are directly quoted, not words with their meanings somehow built in. As Max Cresswell puts it, in his exposition: "we can speak of the same language but with a different meaning assignment" (1980, 19). I'll work with this common assumption as well. Doing so only makes things harder for me, since denying the assumption would allow me to reject the argument underwriting the worry about quotational belief specifications that I'm taking it upon myself to address.

[6] Bigelow says that his argument is "very like" that given by Church (1950), who in turn credited Langford (1937). (Stephen Schiffer's objection to Davidson's (1968) account of indirect discourse (discussed above, p. 7), is of the same general sort.) Seymour (1992) attempts a rehabilitation of quotational analyses. He tries to avoid Bigelow's argument by holding that "the sentences that occur within quotes in a substitutional formula [which he proposes as the quotational paraphrase of a

being disposed to assert or even *being disposed to assert in English*; the argument goes through. It threatens an entire family of proposals.

One prominent objection to this style of argument is due to Field (2001). He defends a quotational approach to content specifications and objects to translation-based arguments on the grounds that when we translate a quotational attribution, we *also translate* the words that are quoted in it for the purpose of content-specification—just as when we translate a novel, we translate the directly quoted dialogue in it. Field would read *M* on the lines of "Marian accepts a sentence playing the role in her psychology that 'Robin wins' plays in mine" (163). Translated into the variant of English in the way Field recommends, this becomes (following Sellars (1962, 35–36)): "Marian accepts a sentence playing the role in her psychology that 'Robin loses' plays in mine." And this does have the same truth condition as the corresponding t-ascription in the variant language. So on Field's view the Bigelow argument fails because it mistakes how translation works (in the cases in question).

Whether Field's response to the argument succeeds is a large topic, and this is not the place to decide conflicting claims about the nature of translation. For my purpose it is enough to note that the Bigelow argument is a challenge to quotational belief-specifications only when they are put forward as *analyses* of t-ascriptions. It does not show that they *fail to specify beliefs*. The claim that they do specify beliefs is consistent with everything in the argument. The problem that arises when we consider the variant language is not that the quotational paraphrase doesn't specify beliefs. Rather, it's that in the variant language, unlike in English, what it specifies is not what the corresponding t-ascription specifies.

But that feature of q-ascriptions isn't at all a problem for the proposal I am making, which is that q-ascriptions succeed *in meeting the Specification Challenge*. The motivating assumption behind the Challenge is that ordinary language does contain locutions that answer to our interest in specifying beliefs in a way that reflects their individuation. So in this dialectical context it's the *specificatory usefulness* of a belief-specifying locution that matters, not its semantic equivalence (or lack thereof) with some other such locution. It's important to appreciate here the magnitude of Heck's reconfiguration of the debate: no longer is it (as I wrote above) focused on the semantics of t-ascriptions. The inquiry now is about the belief-specifying resources of ordinary language *in general*. So the Bigelow argument isn't even a *prima facie* problem for the proposal I'm making.

4.2.2 The "can"

Another general worry about q-ascriptions concerns the "can": when we say that someone *can* express their belief by asserting sentence *S*, how much are things allowed to differ from actuality, for the purpose of evaluating that claim? One might point out that for any given belief and sentence, it's metaphysically possible that the

Footnote 6 continued
t-ascription] semantically presuppose their semantical rules" (193). It is hard to judge the success of the attempt, since Seymour does not develop the idea of a quoted sentence presupposing its semantical rules.

latter is suitable for use in expressing the former. (Mostly what that involves is just *S*'s having in the believer's language a meaning that matches the content of the belief.) If *that's* what we mean by some q-ascription then it tells us nothing about the content of the belief. (In the Bigelow argument, a speaker of the variant language does at least know (on any relation that is a plausible candidate for *R*) that if the quotational specification is true then "Robin wins" has the same content in Marian's language as one of her beliefs. Here we wouldn't even have that, if the "can" in a q-ascription invokes bare metaphysical possibility.) What modality does "can" bring to bear, when used in a q-ascription?

It has long been recognized that the specific modal meaning that the word "can" expresses depends on the context of its use (Lewis 1976, 150). "In ['John can swim'], *can* indicates the subject's intrinsic ability, while ['Mary can see the ocean'] has to do with the situation in which the subject finds herself" (Portner 2009, 196). In a q-ascription it is a meaning of the latter sort that is expressed.[7] It is a straightforward " 'can' of ability," not of bare metaphysical possibility (let alone of something even less plausible such as epistemic possibility, for which "can" is seldom used (Collins 2009)).

So although it's *possible* for q-ascriptions to be interpreted in a way that makes them almost vacuous, the normal interpretation of the "can" is one that makes them quite useful. And in many cases—those in which the subject *actually does* express the belief in question—no "can" is even needed in a successful q-ascription.

4.2.3 Expressing

Next (under the heading of general worries about q-ascriptions) there is the matter of their invocation of the relation of *expressing*. Invoking this relation in a q-ascription does not introduce an obligation to give a reductive analysis of that relation in terms of other relations; some specific reason would have to be produced for doubting that there is a relation corresponding to the word. As the word is part of our folk-psychological vocabulary rather than from some contested theory, *prima facie* it is reasonable to take there to be a relation it picks out.

Nevertheless one might wonder whether the expressing relation is too liberal, since we speak of people expressing their feelings in song or dance, of political views being expressed in plays or novels, and (perhaps) of people expressing their beliefs by *acting* this way or that (other than in speech acts).

[7] In what Portner says is "the 'standard theory' of modality within formal semantics" (47), that of Kratzer (1977, 1981), context contributes a "conversational background" *f*, a function from worlds to sets of propositions; and modals are evaluated, at a world *w*, as quantifying either universally or existentially over the set of worlds in which all the propositions in *f*(*w*) are true. So our "can" concerning Lois, evaluated at a world *w*, introduces existential quantification over all the worlds in which the propositions capturing the conversational background—in this case, "the situation in which [Lois] finds herself" (in *w*)—are true. Equivalently: it is the expressive options Lois has in worlds *in which her situation is the same as in w* that matter to the truth (at *w*) of a q-ascription concerning her. So when evaluating the modal at the actual world, we needn't worry about farfetched scenarios in which Lois's words have changed their meanings, or she is mute, or her belief-expressing mental machinery has gone awry, etc. (Of course we do need to worry about them when *evaluating at* worlds in which they occur, but that's fine: we get the right truth value by doing so.)

To this I reply that liberality in itself is not a problem since we are concerned with cases in which *a belief* and *the assertion of a sentence* are the relata. So it is neither here nor there, that expression relates relata of other types as well: from the fact that a belief can be related (via expression) to (say) someone's dancing a certain way, it does not follow that the *overall* relation between the dancing and the belief thereby expressed is the same as that between an assertion and the belief thereby expressed. In the latter case the expressing relation obtains along with that of equivalence of truth condition; in the former case it obtains without the latter relation's obtaining. So we are not somehow *assimilating* assertions to dances, when we specify them in q-ascriptions as relata of a relation of which dances (let's suppose) are also relata.

Even without a reductive account of the expressing relation, we can state a couple of plausible necessary conditions on an assertion's expressing a belief. One is that the belief play a central role in causing the assertion. This expressive role is remarkably *stable*, by which I mean that unlike other actions that a belief causes, a belief's role in causing an assertion that expresses it is much less hostage to the comings and goings of auxiliary beliefs than its role in causing actions of other kinds. Consider a voter marking a ballot: whether the belief *that A is the best candidate* causes the believer to vote for *A* depends on her views about strategic voting, for example. But those views make no difference to the expressive relation's obtaining between that belief and her assertions of "*A* is the best candidate." When we try to think of auxiliary beliefs that do make such a difference we think of beliefs about the propriety, say, of speaking up on such a topic, or beliefs about the meanings of words (which affect her choice of sentence to assert). But in most cases beliefs of the former sort aren't there to hinder expression, since most topics aren't controversial or possibly offensive; and in almost all cases beliefs of the latter sort are very stable for the agent over time—rarely is a speaker's choice of words to express some belief hindered by a worry that her words have changed their meanings since she last used them. All this is to say that although changes in other beliefs typically do change a belief's causal role in relation to action *in general*, its role in *causing expressive assertions* is much less vulnerable to changes of that kind. The words with which a belief can be expressed, in some agent, don't change due to the sorts of changes in other beliefs that often occur in the normal course of events. This makes them unusually useful as items by reference to which we can specify beliefs; hence it is unsurprising that in ordinary language we do this, as opposed to specifying beliefs in terms of their relations to behaviour of other kinds.

Another plausible necessary condition on an assertion's expressing a belief is that the speaker takes them to have the same truth condition. This leaves room for some flexibility, which we might want to leave room for inasmuch as the notion of expression arguably admits of degree: one sentence can better express a belief than another.

4.2.4 Inexpressible beliefs

Finally another general worry about q-ascriptions is that they can't be used to specify beliefs that are linguistically inexpressible, either because the believer is a

pre- or non-linguistic creature or because, say, the believer lacks the vocabulary to make his belief explicit in an assertion.

One overly quick response to this would be to insist that the non-linguistic among us don't really have beliefs. (For what it's worth, though, cognitive psychologists do seem to lean in the direction of restricting the use of "belief" to describe language users: "most researchers now prefer not to use such highly cognitive language in describing the representations and computations of young infants because the term 'belief' can be taken to imply a language-like format" (Carey 2009, 104).) It would also be controversial to maintain that every belief that does reside in a language user is one to which that believer can give expression in speech. The "can" would run up against the possibility of inhibitory mechanisms (if the belief is, say, offensive or shameful) as well as facts about variations in articulateness.

But the typical Frege case—our central concern here—is one in which the believer is able to articulate their beliefs in speech. Very often it is part of the story that the believer does give expression to the relevant beliefs by asserting sentences. Indeed, Kripke's oft-discussed (1979) paper, the point of which is to show how the sort of implication often taken to be problematic for Russellianism can be obtained without reliance on that doctrine, relies explicitly, in its presentation of various scenarios, on statements about what sentences the believer would *assert* (439).

I conclude that neither the Bigelow argument, worries about the use of "can" in q-ascriptions, concerns about the liberality of the ordinary-language notion of expression, nor concerns about beliefs inexpressible in language give us a reason to reject the proposal that q-ascriptions are belief specifications worth considering in relation to Heck's Specification Challenge. All that this means, however, is that they're not non-starters; what matters next is how they fare in relation to the aspects of the Challenge that we identified earlier.

5 How q-ascriptions meet the specification challenge

Recall our elaboration of the Specification Challenge: it asks for sentences composed of ordinary vocabulary, that literally specify beliefs sufficiently finely to distinguish beliefs with the same Russellian contents, in relatively non-exotic cases such as Lois's, and which do better than supplying no information about truth conditions. Do q-ascriptions meet these requirements? Let's go back to Heck's Lois scenario. Our two q-ascriptions are:

q-Superman Lois has a belief that she can express by asserting "Superman is at the school board meeting."

q-Clark Lois has a belief that she can express by asserting "Clark Kent is at the school board meeting."

These sentences are formed from ordinary vocabulary, used in the ordinary way; so there is no prospect of their being disqualified as candidate locutions on the ground of their failing to comply with the part of the Challenge concerning the ordinariness of the sentences in question. The more substantial questions concern, first, their

usefulness in "distinguishing" beliefs that t-ascriptions, as conceived by the Russellian, don't distinguish; and second, how close they come to "encoding" the truth conditions of the beliefs they specify.

5.1 The distinguishing powers of q-ascriptions

For the truth of *q-Superman* Lois's having the S-belief suffices, since it *is* part of the psychological role of that belief that it can cause (in the way we're picking out by invoking the expressing relation) assertions by her of "Superman is at the school board meeting." The C-belief does not do that, so Lois's having the C-belief does not suffice for the truth of *q-Superman*. That q-ascription therefore satisfies Heck's stated criterion for "distinguishing" between the two beliefs.

We should also consider other types of cases in order to get a better picture of the distinguishing powers of q-ascriptions. The Lois situation as Heck tells it is a classic Frege case: she uses two co-denoting simple expressions that she thinks are not co-denoting. But there are variant cases. One that might seem challenging to the q-ascription approach is a case of someone who uses *one* simple expression that she thinks denotes differently on different occasions of use. Recall the case I introduced above (§ 2), in which Lois mistakenly takes there to be two Supermans, believing of each that he can fly. Call these two beliefs "the S1-belief" and "the S2-belief." They have the same Russellian contents, so t-ascriptions, as conceived by the Russellian, won't distinguish them. Can q-ascriptions do so?

What we would need, in order to distinguish the S1-belief from the S2-belief, is a q-ascription made true by Lois's having the S1-belief but not made true by her having the S2-belief. Ideally it would be one whose quoted sentence has the same Russellian content as "Superman can fly." Since the proper name she uses in expressing S1 is the same as that she uses in expressing S2, quotation of expressive assertions seems unlikely to distinguish the two. So this case is much more challenging for my proposal than the case in Heck's scenario; the constraints seem unsatisfiable.

But let's look more carefully. Since Lois takes there to be some differences between (what she takes to be) the two Supermans, there will be some differences between what she believes about (what she takes to be) one of them and what she believes about (what she takes to be) the other. For concreteness' sake, suppose that Lois believes that one of the Supermans is famous for having rescued a cat and the other is famous for having rescued a dog, and that to her, the S1-belief concerns the former and the S2-belief the latter.

Now ask how she can express her S1-belief and her S2-belief. Here the conversational context matters. When, as could very easily happen, both (of what seem to her to be the two) Supermans are conversationally salient, the demands of conversational helpfulness will motivate her, when she decides to express the S1-belief, to say something other than "Superman can fly," for she will believe that her audience needs to be told *which* Superman she is talking about.

English does have a device that lets her *say* which, and to do so—importantly for us, given the just-stated constraints—in a way that doesn't affect the Russellian content of the sentence she uses. The device of *nonrestrictive relative clauses* is

perfect for just this expressive function. Lois can express the S1-belief, but not the S2-belief, by asserting

Superman, who rescued the cat, can fly.

Clearly this doesn't express the S2-belief, since from her point of view that belief is about the Superman who rescued the dog. But does it express the S1-belief? To answer this we should see what difference the nonrestrictive relative clause—which she inserts as an identificatory aid to her audience—makes to the truth condition of the sentence.

On most semanticists' treatments, it makes no difference: "The basic intuition that most authors have expressed about the semantics of nonrestrictive modification is that nonrestrictive modifiers are not semantically composed at all with the phrases they modify" (Heim and Kratzer 1998, 64). In an influential article Bach (1999) proposed to account for this by saying that a sentence containing such a clause expresses *two* propositions, not one. For example, the sentence

Ann's computer, which she bought in 1992, crashes frequently

expresses the two propositions *that Ann's computer crashes frequently* and *that she bought it in 1992*, "but it does not express their conjunction" (Bach 1999, 351). One of the expressed propositions, the one corresponding to the overall sentence minus the nonrestrictive clause, is more salient, or somehow primary (353). Although Bach doesn't put his point in terms of assertions, one way to do so is to say that the assertion of one *sentence* amounts to the assertion of two *propositions*. How does this bear on the question of which belief is thereby expressed?

A natural thing to say is that the belief whose content is that of the primary proposition is the one expressed by the assertion, and that the secondary proposition functions less as the expression of a belief (although the speaker does of course *have* a belief corresponding to it) and more as a communicative aid to the expression of the belief corresponding to the primary proposition.

With this apparatus clarified we can now say that Lois's use of "Superman, who rescued the cat, can fly" expresses the same belief that her use of "Superman can fly" does, when the latter is used by her for the Superman who rescued the cat. In other words, those two sentences *both* express the S1-belief, despite the former's containing the extra identifying words—words that semanticists tell us make no difference to the sentence's content. It's just that asserting the sentence with the extra clause expresses it *better* than an assertion of the other does, in a conversational context in which Lois believes that both Supermans are salient. That q-ascription with the extra clause, therefore, does the distinguishing Heck requires, in the case under consideration.[8]

[8] This solution would not handle cases, if such are possible, of the sort that Pryor discusses (see above, n. 1): Alice's believing that there are multiple Bobs on her baseball team, without believing anything distinguishing about any of them. For she would have nothing distinguishing to *say* about any of them in a relative clause. As cases of this sort are much more exotic than Frege cases or even Paderewski cases, I assume that Heck would not count it against a proposal that it fails to offer a specificatory solution that works in this sort of case. (I'm grateful to a referee for pointing this out.)

Could a proponent of t-ascriptions make this move as well? Do

Lois believes that Superman, who rescued the cat, can fly

and

Lois believes that Superman, who rescued the dog, can fly

distinguish between Lois's S1- and S2-beliefs? No, since the nonrestrictive relative clauses make no contribution to the semantic values of the "that"-clauses. (The *quotation* of one does make a difference, in a q-ascription, to *its* semantics, which is what makes the q-ascriptions approach work in these cases.) Furthermore what they express in these cases, on the usual sort of line about nonrestrictive relative clauses, are commitments not by Lois but *by the belief ascriber* concerning Superman; so their uselessness in specifying Lois's beliefs is over-determined.

In a good range of cases, then, q-ascriptions do useful distinguishing work that isn't done by t-ascriptions (as the Russellian conceives of them). Unlike t-ascriptions they allow us to distinguish beliefs with the same Russellian content not only in standard Frege cases such as the one Heck presents but even in relatively exotic Frege cases that are superficially more challenging for q-ascriptions.

5.2 Q-ascriptions and truth conditions

Above I explained why the Bigelow argument doesn't undercut the claim that q-ascriptions specify beliefs. What it undercuts, rather, is a claim about their relation to t-ascriptions, and my proposal involves no such claim. Nevertheless there is a worry about my proposal which the Bigelow argument does encourage, since that argument also brings out the fact that q-ascriptions don't encode truth conditions (in the sense explained above, § 2).

To see this, put yourself in the position of a speaker of the variant language described in the Bigelow argument. You understand the q-ascription "Marian has a belief she can express by asserting 'Robin wins'." This understanding doesn't suffice for knowing the truth condition of the belief(s) it specifies. For even if you know that Marian speaks English (which the q-ascription doesn't tell you), and you speak it yourself (which isn't required for you to understand the q-ascription as a sentence of the English-variant), getting from understanding the q-ascription to knowing the truth condition of the specified belief(s) would require that you use that information in an inference, along with what the q-ascription itself tells you. Which is just to say that the q-ascription doesn't *encode* the truth condition. In short, while the q-ascription doesn't encode the truth condition, someone who understands it, and has the two pieces of knowledge just described, is *one inference away* from knowing the truth condition.

Since our present concern—guided by the Specification Challenge—is with the belief-specifying resources available in English to ordinary speakers, the situations that interest us are typically ones in which it is common knowledge that the subject of the q-ascription speaks English, and that the ascriber and her audience do also. Presumably our ordinary locutions developed in response to ordinary needs felt in ordinary situations. In cases of this typical sort, being one inference away from

knowledge of the specified belief(s)' truth condition is for practical purposes just as good as having it encoded in the ascription, since the required auxiliary premises are, in these cases, known. (Note that it is not just its being one inference away that makes for the usefulness of a q-ascription in the typical case, but the presence of the required auxiliary knowledge needed for that inference. One could, after all, be one inference away in a very different sort of case: one in which the required auxiliary premises are *not* known, perhaps are even unknowable on current evidence.)[9]

So the proposal that q-ascriptions satisfy the requirements of the Specification Challenge fares well: they are well-formed, having truth conditions that do specify beliefs; they can be used quite ordinarily to distinguish beliefs having the same Russellian content not only in the sort of case Heck presents in framing his challenge but also in cases that are superficially much more challenging for the q-ascription approach; and they come very near to—for practical purposes, in most uses, just one trivial inference away from—encoding truth conditions, which puts q-ascriptions *much* closer, in terms of usefulness, to t-ascriptions than to the sort of *ad hoc* descriptions in terms of causal roles that Heck himself used (in telling us what the S-belief and the C-belief are).

6 Russellianism, q-ascriptions and ordinary discourse

My main claim in this paper has been that q-ascriptions meet the Specification Challenge: they use ordinary vocabulary to literally specify beliefs in a way that reflects more of their psychological roles than t-ascriptions do—enough, in particular, to be helpful in Frege cases. To this one might reasonably respond, Why then aren't q-ascriptions the dominant mode of belief specification? Is it because our ordinary practice is somehow deficient?

We can articulate the point as follows. Suppose that Russellianism is true; that (as Heck argues) t-ascriptions with Russellian truth conditions do not satisfy the requirements of the Specification Challenge; and that (as I have argued) q-ascriptions do fully satisfy those requirements. Now consider our choice of locutions, when we wish to speak about beliefs. If our goal in such discourse is always to satisfy the requirements of the Specification Challenge, then it seems we should always use q-ascriptions, for they meet those requirements and t-ascriptions don't. (Here I'm simplifying by taking it that the choice is just between t-ascriptions and q-ascriptions.) But this is far from being the case. Very often—perhaps even usually, although a corpus survey would be needed to back up that claim—we use t-ascriptions.

[9] One might wonder whether this introduces a regress of t-ascriptions, if knowing a sentence's truth condition comes from knowing its meaning. For that knowledge is itself knowledge of what *meaning* t-ascriptions state: "Superman can fly' means *that Superman can fly*. But we don't here have the same challenge, for coarse-grained (i.e. Russellian) meaning ascriptions are enough to get to knowledge of truth conditions: "that Superman can fly" and "that Clark Kent can fly" specify the same truth condition. (I'm grateful to a referee for posing this question.)

The worry, then, is that my proposal justifies an implausible claim about our ordinary discourse. My response to this comes at the worry from each side. First I explain how the assumptions in the line of thought just sketched are implausibly strong. Then I show that q-ascriptions have a greater role in our actual practice than philosophers typically acknowledge. Combined, these considerations greatly undercut the worry that there's a problematic mismatch between my claims about q-ascriptions and the facts about our ordinary talk about beliefs. Moreover, as I note at the end of the paper, my proposal's implications concerning such practice are not obviously less problematic than are those of Heck's proposal.

6.1 Problems with the worry

First, some points about the line of thought behind the worry: it is less secure than it seems.

It's not the case that our goal in talking about beliefs is *always* to satisfy the requirements of the Specification Challenge. It is a commonplace observation that there are *two* large purposes served by our talk about each other's beliefs: we want to understand each other's actions, and we want to gather, store and pass along information. David Lewis puts the point by saying that one purpose of belief-talk is that of "revealing how the subject's actions serve his desire according to his belief and how his belief evolves under the impact of his experience" and that the other is that of describing him as "a partner in cooperative work and as a link in channels for information" (1986, 58–59).[10] Seldom is our discourse focused solely on one of these goals. Often we are interested to some degree both in a belief's truth condition and its psychological role. So it is implausible to maintain, as in the line of thought sketched above, that specificity of the sort required by the Specification Challenge is *always* our goal in talking about some belief. In many cases our interest is more in the information-conveying role of beliefs, which explains the oft-noted fact that we are very often ready to accept any of several distinct t-ascriptions relating a subject to the same Russellian proposition (see e.g. Wettstein 1986, 205–207; Saul 1998, 366).

In reply one might claim that what matters is that the specificity we often require, while less than that required by the Specification Challenge, still is greater than what t-ascriptions offer if their semantics is Russellian. Here another error must be avoided, which is to think that t-ascriptions with Russellian semantics are explanatorily inert. Someone's standing in the *belief* relation to a Russellian proposition is a psychological fact about them. So a t-ascription *stating* that she stands in that relation to a complex entity of that sort has a psychological upshot. A claim that someone represents in thought the *taller-than* relation, for example, will, against a background assumption that they do so in the typical way, have clear implications about inferences they will make to and from beliefs whose Russellian contents have that relation as a component. Such a background assumption is not

[10] See also Stalnaker's (1984, 4–5) discussion of the "pragmatic picture" and the "linguistic picture" of mental representation.

implausible. Crimmins (1992, 94–97) has argued that such psychological typicality is widespread: "there are many cases in which, among members of a certain group, everyone has a normal notion of some individual" (97); moreover, "it seems much more difficult to imagine cases of multiple, unconnected ideas of the same property or relation than to imagine multiple, unconnected notions of individuals" (95).[11] Crimmins doesn't present a developed theory of what "normal notions" are. But if there is something to his thought, then t-ascriptions in their typical uses would go *well beyond* being minimally non-inert. A full investigation of this topic would take us far afield, into empirical social psychology. At any rate it is far from obvious that the explanatory upshot to a t-ascription (with Russellian semantics), given in a conversational setting in which there are common-ground assumptions about the typicality of how the subject relates to components of the Russellian proposition in question, is too meagre to serve our interest, in that setting, in psychological explanation.

Finally it should be noted that the line of thought behind the worry neglects the possibility of *asserted* contents doing explanatory work, even in cases where the asserted contents are not literally expressed. Soames (2002, 2005), in particular, has recently used in defence of Russellianism the idea that these can diverge significantly. While his proposal doesn't involve the idea that q-ascriptive contents are asserted by t-ascriptions, it is not an unreasonable idea to pursue. (Indeed in his earlier work Soames did, as we noted, propose that claims about words (with which the subject might express the specified belief) are *suggested by* t-ascriptions.) That is a way in which we might reconcile a relative paucity of q-ascriptive *sentences* being uttered, with the claim that often q-ascriptive *contents* are asserted (and play roles in explanations). Here is not the place to defend this idea. What is relevant here is that the line of thought behind the worry requires an argument *against* it; this is another respect in which that line of thought is less compelling than it can seem.

The line of thought behind the worry was that there is a mismatch between what my proposal entails about our ordinary discourse, and what is actually the case with that discourse. Having discussed that line of thought I'll now tackle the issue from the other side, arguing that there is more q-ascriptive ordinary discourse than philosophers usually acknowledge.

6.2 Varieties of q-ascriptions in ordinary use

When we look at actual belief-specifying discourse—poking around in a searchable corpus is one way to do this—we find that t-ascriptions are just one among multiple often-used forms. While actual uses of the canonical form of a q-ascription—"*X* has a belief she can express by asserting ⌜*p*⌝"—are rare to nonexistent, their *contents*

[11] It is indeed remarkable that the Frege cases philosophers discuss always concern substitution of codenoting terms or noun phrases, rather than verb phrases, quantifiers or other parts of speech. Of course this fact about us, if it is a fact, is contingent; but our practice does develop in response to contingencies, if they are widespread.

are, I think, expressed quite often in sentences of the form "*X* believes ⌜*p*⌝."[12] (In introducing q-ascriptions I put forward the more explicit, canonical form in order to disaggregate that content, making it easier to defend the legitimacy of it piece by piece. As we saw, there are several components to it that require separate defense.) And sentences of this sort *are* common. A query for "believe" or "believes" followed by a quoted sentence gives several thousand instances in the Corpus of Contemporary American English (Davies 2008). Representative are:

> You may believe "I just don't understand money," when in fact, it's high finance that you don't understand. (*Essence* magazine, 1996)
>
> Eight in ten of those who believe "government should do more to solve problems" voted for Obama (*USA Today*, 2012)
>
> This boundary is certainly a real issue for faculty who assign grades to students, but Connell believes "the student/faculty member relationship is different than the student/librarian relationship." (*Reference & User Services Quarterly*, 2014)

Each of these, I think, readily admits of paraphrase into the q-ascriptive canonical form.[13]

Another very common form is what we might call *mixed* t- and q-ascriptions. By this I mean t-ascriptions parts of which are direct quotations. Something like this is familiar from the recent literature on so-called "mixed quotation," an example of which is

> Alice said that life "is difficult to understand." (Cappelen and Lepore 1997, 429)

As the literature on mixed quotation shows, its semantics is a tricky business. But there's no dispute over whether there are such quotations; in fact it is an extremely common form in contemporary writing.

One might balk at the very idea of a corresponding form for ascribing beliefs, on the grounds that while in a mixed quotation it's pretty clear that some *utterance of the directly quoted words* is at issue, having a belief involves no utterance of words—so what would the directly quoted part of such a belief ascription be quoting? Nevertheless, this is a common form of belief specification. A search on "believes that" followed (within 8 words) by a direct quotation yields many hundreds of very ordinary-looking results. An example:

[12] I will not try to justify this claim here, as doing so would require a full inquiry into the lexical semantics of "believes." We would have to decide, for example, whether we want the semantics of "Susan believes John" to run along exactly the same lines as "Susan believes 'Robin wins'."

[13] Not all philosophers have neglected such uses. Geach (1967) vigorously defended the legitimacy of what he called "the *oratio recta* construction"—his example being

> James believes "My wife's fear is 'I have cancer'."

He claimed that "Philosophers have a curious prejudice against" this form even though it is "common in all vernaculars" (167).

> George Kennan, a well-known diplomat who formulated U.S. policy toward Russia after World War II, believes that NATO expansion is a "fateful error of American policy" that will "inflame" Russia. (*USA Today*, 1997)

Construing the direct quotation as working the same as in a q-ascription of the canonical form allows us to make sense of such uses. What seems to be going on here is that a belief is being specified *partly by a specification of words the subject used to express it* in some assertion. But it is also used in describing hypothetical scenarios, so it's not limited to use in reporting on beliefs that actually are expressed:

> If, for example, the player believes that a "winner" tries harder than a "loser," then casting him/herself into the category of a "winner" may lead to the causal attribution that good effort caused success. (*Journal of Sport Behavior*, 1991)

This is what we would expect, given our inclusion of the "can" of ability in our canonical q-ascriptive form. The form is also used to attribute beliefs to groups:

> The group believes that apes are "conscious" and so deserve legal protection of their right to life and freedom from imprisonment and torture. (*Psychology Today*, 1999)

We should expect mixed forms to be more common than the forms in which a whole sentence is quoted. The reason is that it is unlikely that *every component* of an attributed content is one whose psychological role is worth specifying in the extra-careful, q-ascriptive way. In terms of Frege cases, the point is that the vast majority of Frege cases will concern only *one* component of the attributed content: Hesperus/Phosphorus; Superman/Clark Kent. As Crimmins noted (discussed above), it strains plausibility to concoct pairs in which, say, both the noun phrase and the verb phrase exhibit Frege-case behavior.

What we should expect to see then, in the typical Frege case, is the q-ascriptive apparatus focused on the one component making it a Frege case. Recall the Lois example. What makes hers a Frege case is that she unknowingly thinks of *Superman* in two ways—not that she does so while also unknowingly thinking of *flying* in two ways, or unknowingly thinking of *being at the school board meeting* in two ways. So an ascription such as

> Lois believes that "Superman" is at the school board meeting

is a very natural, and very easily processed, way to do the requisite t-ascriptional and q-ascriptional work *at once*, using the more complex logical form of mixed belief ascription.[14]

Recall the worry I described at the start of this section. I have argued that the line of thought behind it rests on implausibly strong assumptions. Weakening them

[14] I say "complex" because mixed quotations are not conjunctions of direct and indirect quotations. See Geurts and Maier (2003), criticizing Potts (2007) (which was circulating in manuscript several years before being published) on this point. That exchange concerns mixed quotation, but Geurts's and Maier's reasoning applies also to mixed belief ascription. (I have proposed an account of the logical form of mixed quotation in McCullagh (2007).)

increases plausibility, but correspondingly weakens the conclusion we can draw from them. Thus weakened, it conflicts less with the data about how prevalent q-ascriptions are. Combined with the fact that q-ascriptions are more prevalent than philosophers usually acknowledge, it's far from clear that there's any mismatch between what my proposal entails about ordinary discourse and what we observe about it.

Indeed, the worry can now be turned on Heck's own view. Heck insists that t-ascriptions do meet the Specification Challenge: they specify beliefs in as fine-grained a way as we could reasonably ask. The worry *this* claim invites, now, is that it leaves us unable to explain our use of belief-specifying locutions of other forms—in particular, why much of it is q-ascriptive. Such an argument against Heck's view would have to be for another day. All I wish to have argued here is that Heck's Specification Challenge broadens the inquiry beyond t-ascriptions, and that a Russellianism thus unencumbered can meet it.

Acknowledgements I'm more than usually grateful to an anonymous referee, who prompted significant improvements to the discussion in § 6.

References

Bach, K. (1999). The myth of conventional implicature. *Linguistics and Philosophy*, *22*, 327–366.
Berg, J. (1988). The pragmatics of substitutivity. *Linguistics and Philosophy*, *11*, 355–370.
Bigelow, J. (1978). Believing in semantics. *Linguistics and Philosophy*, *2*, 101–144.
Braun, D. (1998). Understanding belief reports. *Philosophical Review*, *107*, 555–595.
Cappelen, H., & Lepore, E. (1997). Varieties of quotation. *Mind*, *106*, 429–450.
Carey, S. (2009). *The origin of concepts*. Oxford: Oxford University Press.
Carnap, R. (1947). *Meaning and necessity*. Chicago: University of Chicago Press.
Church, A. (1950). On Carnap's analysis of statements of assertion and belief. *Analysis*, *10*, 97–99.
Collins, P. (2009). *Modals and quasi-modals in English*. Amsterdam/New York: Rodopi.
Cresswell, M. (1980). Quotational theories of propositional attitudes. *Journal of Philosophical Logic*, *9*, 17–40.
Crimmins, M. (1992). *Talk about beliefs*. Cambridge, MA: MIT Press.
Davidson, D. (1968). On saying that. *Synthese*, *19*, 130–146. Reprinted in *Inquiries into truth and interpretation* (pp. 93–108). Oxford: Clarendon Press.
Davies, M. (2008). The corpus of contemporary American English. Online resource. http://corpus.byu.edu/coca/.
Field, H. (2001). Attributions of meaning and content. In *Truth and the absence of fact* (pp. 157–174). Oxford: Clarendon Press.
Frege, G. (1892). Über sinn und bedeutung. *Zeitschrift für Philosophie und philosophische Kritik100*, 25–50. Translated as On sense and meaning. In P. Geach & M. Black (Eds.), *Translations from the philosophical writings of Gottlob Frege* (pp. 56–78). Oxford: Blackwell.
Geach, P. (1967). The identity of propositions. In *Logic matters* (pp. 166–174). Oxford: Blackwell.
Geurts, B., & Maier, E. (2003). Quotation in context. *Belgian Journal of Linguistics*, *17*, 109–128.
Heck, R. (2012). Solving Frege's puzzle. *Journal of Philosophy*, *109*, 132–174.
Heck, R. (2014). Intuition and the substitution argument. *Analytic Philosophy*, *55*, 1–30.
Heim, I., & Kratzer, A. (1998). *Semantics in generative grammar*. Oxford: Blackwell.
Kratzer, A. (1977). What 'must' and 'can' must and can mean. *Linguistics and Philosophy*, *1*, 337–355.
Kratzer, A. (1981). The notional category of modality. In H. J. Eikmeyer & H. Rieser (Eds.), *Words, worlds, and contexts* (pp. 38–74). Berlin and New York: de Gruyter.
Kripke, S. (1979). A puzzle about belief. In A. Margalit (Ed), *Meaning and use* (pp. 239–283). Dordrecht: D. Reidel. Reprinted in Salmon and Soames (1988), pp. 102–148.

Langford, C. H. (1937). Review of the significs of pasigraphic systems: A contribution to the psychology of the mathematical thought process by E. W. Beth. *Journal of Symbolic Logic*, *2*, 53–54.
Lewis, D. (1976). The paradoxes of time travel. *Australasian Journal of Philosophy*, *13*, 145–152.
Lewis, D. (1986). *On the plurality of worlds*. Oxford: Blackwell.
McCullagh, M. (2007). Understanding mixed quotation. *Mind*, *116*, 927–946.
Portner, P. (2009). *Modality*. Oxford: Oxford University Press.
Potts, C. (2007). The dimensions of quotation. In C. Barker & P. Jacobson (Eds.), *Direct compositionality* (pp. 405–431). Oxford: Oxford University Press.
Pryor, J. (2016). Mental graphs. *Review of Philosophy and Psychology, 7*, 309–341.
Rieber, S. (1994). Review of talk about beliefs, by Mark Crimmins. *Philosophical Psychology*, *7*, 395–397.
Salmon, N. (1986). *Frege's puzzle*. Cambridge, MA: MIT Press.
Salmon, N. (1989). Illogical belief. *Philosophical Perspectives*, *3*, 243–285.
Salmon, N., & Soames, S. (Eds.). (1988). *Propositions and attitudes*. Oxford: Oxford University Press.
Saul, J. (1997). Substitution and simple sentences. *Analysis*, *57*, 102–108.
Saul, J. (1998). The pragmatics of attitude ascription. *Philosophical Studies*, *92*, 364–389.
Saul, J. (2007). *Simple sentences, substitution, and intuitions*. Oxford: Oxford University Press.
Schiffer, S. (1987). *Remnants of meaning*. Cambridge, MA: MIT Press.
Schiffer, S. (1992). Belief ascription. *Journal of Philosophy*, *89*, 499–521.
Sellars, W. (1962). Truth and correspondence. *Journal of Philosophy*, *59*, 29–56.
Seymour, M. (1992). A sentential theory of propositional attitudes. *Journal of Philosophy*, *89*, 181–201.
Siegel, J., & Shuster, J. (1938). Superman. *Action Comics*, *1*, 1–13.
Soames, S. (1987a). Direct reference, propositional attitudes, and semantic content. *Philosophical Topics*, *15*, 47–87.
Soames, S. (1987b). Substitutivity. In J. J. Thomson (Ed.), *On being and saying: Essays for Richard Cartwright* (pp. 99–132). Cambridge, MA: MIT Press.
Soames, S. (2002). *Beyond rigidity: The unfinished semantic agenda of naming and necessity*. Oxford: Oxford University Press.
Soames, S. (2005). Naming and asserting. In Z. Gendler Szabó (Ed.), *Semantics versus pragmatics* (pp. 356–382). Oxford: Clarendon Press.
Stalnaker, R. (1984). *Inquiry*. Cambridge, MA: MIT Press.
Wettstein, H. (1986). Has semantics rested on a mistake? *Journal of Philosophy*, *83*, 185–209.