

CONTEXT IN THE ATTITUDES

There is a growing consensus that the truth conditions of propositional attitude reports are 'radically' context sensitive, in that, as Quine put it, '... even if we eternalize the contained sentence [in an attitude sentence] and also rid the containing sentence of such sources of truth-value variation as inadequate descriptions, indicator words, and the like, still the whole may remain capable of varying in truth value from occasion to occasion'¹ This, anyway, is what we seem forced to accept if we believe, as I will assume, that the attitude reports (including negated reports) that seem overwhelmingly natural in cases like Kripke's Pierre/-London and Paderewski examples are indeed true, rather than, as some 'Naive Russellians' would have it, in some cases false but pragmatically satisfactory, and that those that seem false are indeed false, and are not simply misleading.

I wish first to motivate very briefly two points about the kind of context sensitive semantics needed for attitude reports, namely that reports are about referents and about mental representations; then I will compare two proposals for treating the attitudes, both of which capture the two points in question.² The first proposal, due to Mark Richard,³ employs an unusual notion of proposition, and traces the attitudes' contextuality to variation in just which relations are expressed by attitude verbs. The second proposal, due to John Perry and myself,⁴ employs a more traditional Russellian notion of proposition, and locates the contextuality not in the content of attitude verbs or of any other overt syntactic elements,

¹ *Word and Object*, p. 218. An eternal(ized) sentence is one whose significance and truth value do not vary from context to context.

² The present comparison should not be mistaken for a thorough review of accounts that take attitude statements to be about both referents and representations (see, for example: Nick Asher, 'Belief in Discourse Representation Theory', *Journal of Philosophical Logic* 15, 1986, pp. 137–189, and Hans Kamp, 'Prolegomena to a Structural Theory of Belief and Other Attitudes,' in C. A. Anderson and J. Owens (eds.), *Propositional Attitudes*, (Palo Alto, CSLI, 1990).

³ *Propositional Attitudes: An Essay on Thoughts and How we Ascribe Them*, (MIT Press, Cambridge, 1990)

⁴ 'The Prince and the Phone Booth: Reporting Puzzling Beliefs', *Journal of Philosophy* 86, 1989, pp. 685–711.

but in the presence of entities to which the speaker refers only tacitly. These features reveal key differences in views which, as I hope to show, are nonetheless close neighbors in a promising part of theoretical space.

1. REFERENTS AND REPRESENTATIONS

If the puzzling attitude reports that seem true are true, then while such statements are semantically about the objects, properties and relations normally referred to in uses of the expressions in their content sentences, such statements also constrain and sometimes uniquely specify the *token mental representations* with which the agent allegedly represents these referents. To support the first point, I will draw from the work of the Naive Russellians; for the second I will examine the attitude puzzle cases so common in the literature and suggest a simple principle about how they can be generated.

The Naive Russellians, including Nathan Salmon,⁵ Scott Soames,⁶ and past time-segments of Mark Richard⁷ and Jon Barwise and John Perry,⁸ have argued forcefully for the principle of ‘semantic innocence’, to the effect that the expressions embedded in attitude statements refer to the same objects, properties and relations they refer to in other uses. I will sketch some of their arguments.

Most obvious and forceful are arguments from initial plausibility. To say that Susan believes that Bush is popular is to relate Susan to Bush – it is to specify something that Susan believes about Bush. And this is as much to make a claim about Bush as about Susan. The intuition is enhanced by noticing parallels between the way singular terms within attitude reports are treated, and the way constant terms are treated in simple referential formal systems.⁹ We feel free to quantify existentially over individuals referred to within the scope of attitude operators, and to refer back with pronouns inside the that-clause to an individual we have mentioned outside it. If Susan believes that Bush is popular, then there is someone, namely the President, such that Susan believes that he is popular. The pretheoretical intuition that the initial report is about Bush (which translates into a theoretical intuition that Bush is a propositional

⁵ *Frege's Puzzle*, (MIT Press, Cambridge, 1986).

⁶ ‘Direct Reference, Propositional Attitudes and Semantic Content’, *Philosophical Topics* 15, 1987, pp. 47–87.

⁷ ‘Direct Reference and Ascriptions of Belief’, *Journal of Philosophical Logic* 12, 1983, pp. 425–452.

⁸ *Situations and Attitudes*, (MIT Press, Cambridge, 1983).

⁹ See Salmon, *Frege's Puzzle* and Richard, *Propositional Attitudes*.

constituent of the claim made), thus deserves treatment as our default position; fancy arguments are needed rather to deny than to support it.

But we nonetheless have fancy arguments in favor of semantic innocence. For one, it has been noted that we normally feel free, in translating attitude reports, to translate a name with a different coreferential name, where the two names share nothing but a referent; and we never translate a name with one having a different referent.¹⁰ For another, it has been noted that the referent of a name used in an attitude report is among the only features of the reported attitude that cannot be canceled by further conversational contributions. For instance, while it is acceptable to say “Tom believes that Tully denounced Catiline; I do not mean that he would put it that way”, it is semantically inadmissible to say instead “. . . I do not mean that his belief is about the person Tully” (or even “. . . about the person Cicero”).¹¹ For yet another, it has been argued that unless attitude reports are to represent a departure from well-supported principles of compositionality and direct reference, it must be that the semantics for attitude reports is compositionally derivative on the semantics for the embedded content sentences, and so that even in attitude reports, names refer to objects (and property-words to properties, and so on).¹²

None of the arguments can be considered knock-down, and that is not surprising, since they relate to foundational issues in nascent semantic theories, which are not to be judged until they are mature enough to bear responsibility for their faults. But we make our best guesses, and the considerations cited support the hunch that semantic innocence will hold by and large in any satisfactory semantic theory. But keep in mind that semantic innocence does not entail that the *only* aspect of a name that affects the claim made in a belief report is its referent.

It is old hat that, if our intuitions about truth values are on target, many attitude statements report not only the objects and properties the agent’s attitude is claimed to be about, but also something more internal about the agent’s cognitive fixes on the properties and objects. Many views of ‘cognitive fixes’ that take them to be such abstract, public entities as intensions, descriptions or senses (at least on the simplistic understanding of Frege popular in the Colonies), have been tried and found wanting, as

¹⁰ See Alonzo Church, ‘Carnap’s Analysis of Statements of Assertion and Belief’, *Analysis* 10, 1950, pp. 98–99, and Saul Kripke, ‘A Puzzle About Belief’, in A. Margalit (ed.), *Meaning and Use*, (Reidel, Dordrecht, 1979).

¹¹ See, for example, Richard’s ‘Attitude Ascriptions, Semantic Theory, and Pragmatic Evidence’, *Proceedings of the Aristotelian Society* 61, 1987, pp. 243–262.

¹² See Salmon, *Frege’s Puzzle*, and Soames, ‘Direct Reference, Propositional Attitudes and Semantic Content’.

theoretical tools both in philosophical psychology and in semantics for attitude reports. I believe – though many disagree with me – that the arguments for this are extremely strong (if not, again, knock-down); in any event I will not rehearse them here.¹³

I think the most successful accounts we have of cognitive fixes take them to be *token mental representations*, concrete particulars analogous to file folders or concrete nodes in concrete networks (these are to be distinguished from linguistic tokens, such as inscriptions and utterances). A token representation can belong only to a single agent, persists through time, and is at any one time involved in various beliefs, desires, intentions and so on, but it is not individuated by the contents of these attitudes any more than a file folder is individuated by the stuff in it. There are important questions about the details and even the central tenets of this conception of cognitive fixes, but it is not my task to examine these.

If we adopt the view that cognitive fixes are to be explained in terms of representations, then our approach to describing the cognitive facts in the puzzle cases is obvious. These are all cases in which an agent has two unconnected representations of a thing or property. Kripke's Pierre, for instance, has two representations of London, one that is involved in city-seen-in-travel-brochure beliefs, another that is involved in home-town beliefs.

There is only one really crucial feature of Pierre-like cases, in which a belief sentence can be used on two occasions, with no obviously changing references, to make claims with different truth values: our agent must have two unconnected representations of a thing, such that only one is involved in a belief that the thing is such-and-such (pretty, talented, a Catiline denouncer, seen in the evening). This means, if our truth-value intuitions are genuine, that belief reports constrain the representations that may possibly be involved in a belief satisfying the report.

Not only can we distinguish among competing representations with our reports in the Pierre-like cases, it seems clear from the same sorts of intuitions that we can center the discussion around a specific representation so that the report is true only if the agent has a belief involving that specific representation. That is, we can use attitude reports not only to loosely constrain, but even to uniquely specify representations. When a

¹³ Among the considerations I have in mind are Kripke's arguments against description theories of reference (see 'A Puzzle About Belief' and *Naming and Necessity*, Cambridge: Harvard University Press, 1980), and the standard objections to Fregean theory (interpreted as being analogous to description theories). Stephen Schiffer examines and enhances many such arguments in his recent book, *Remnants of Meaning*, (MIT Press, Cambridge, 1987), as does Richard in *Propositional Attitudes*.

specific representation is salient, we can claim that the agent has a belief involving *it* (I think this happens in many of the puzzle cases, including the Pierre case). It can seem odd to suggest that we normally converse about specific cognitive representations, which after all must be strange entities hidden deep within the heads of the agents we discuss. But I think a strong case can be made for taking representations to be key players in ordinary folk psychology. Representations are commonplace: facts about representations underlie our best commonsense explanations of behavior. Representations are epistemically accessible: we commonly know at least as much about each others representations as we know about each others' left knees, and no one thinks that left knees – even ones hidden from view – are beyond our referential reach. Representations are far more systematically of interest to us than are left knees, so it is not outlandish to suppose that we have ways of constraining and of specifying the ones about which we speak.

On any plausible account, the ways we constrain or specify representations will have a great deal to do with context. All of which raises the question of just how context comes into the semantics.

An unsatisfactory idea we can dispense with at the start is to say that context provides (somehow) simply a set of representations (or conditions on representations) that must be involved in a belief if the report is to be true. The reason this is no good is that it will not allow us to distinguish the claim,

- (1) Tom believes that Cicero was louder than Tully,

from the claim,

- (2) Tom believes that Tully was louder than Cicero.

We need to specify in some way which *job* the different representations are claimed to perform in the reported belief. In short, we need for our semantics in some way to assign conditions on representations to the different roles representations must play in the reported belief (these might range from empty conditions to conditions specifying particular representations).

2. TWO APPROACHES TO CONTEXTUALITY

Richard's account can be seen as a very natural proposal on a certain influential approach to contextuality familiar from David Kaplan's work. We define a class of *contexts*, into which arbitrary sentences can be plugged and interpreted; context sensitivity is a difference in the truth conditional

content of a sentence as interpreted in different contexts. On this, ‘doughnut’ conception of context, we will want to be able to say that, in a single context, sentence (3) expresses a true claim while (4) expresses a false one.¹⁴

- (3) Tom believes that Cicero is famous.
 (4) Tom believes that Tully is famous.

Since sentence-in-context semantics by definition takes the claim expressed to be a function only of the sentence uttered and the context in force, the difference in the claims must be keyed by the sole syntactic difference in the sentences uttered, namely, by the name used for Cicero. And this difference will result in claims that place different conditions on the representation of Cicero involved in the reported belief. Our semantics for attitude reports, then, will need to take the name used as determining in a particular context the condition invoked; it will have to be sensitive to more than just the referent of the name (which of course is not to say that the name must refer to something else). The most natural way of accomplishing this is for the context to determine a mapping from names to conditions on representations. This, to a first approximation, is Richard’s proposal.

According to Richard, the semantic value of a that-clause is a structured syntax/world hybrid containing the expressions occurring in the clause paired with their referents (he calls such hybrids both ‘propositions’ and ‘Russellian Annotated Matrices’ or ‘RAMs’). He would write the RAM expressed by ‘that Cicero is famous’ as follows:

$\langle \langle \text{‘Cicero’}, \text{Cicero} \rangle, \langle \text{‘is famous’}, \text{being famous} \rangle \rangle$

The cognitive objects of belief are similar RAMs, containing mental representations instead of names. Tom, who has representations ρ_{Cicero} and ρ_{Tully} of Cicero, stands in the *cognitively* basic belief relation *Bel* to the following RAM:

$\langle \langle \rho_{\text{Cicero}}, \text{Cicero} \rangle, \langle \rho_{\text{famous}}, \text{being famous} \rangle \rangle$,

but does not *Bel*-believe the similar RAM containing ρ_{Tully} .

In a true belief report, the RAM the speaker expresses with his that-clause must *represent* a RAM the agent *Bel*-believes. For any function *f* and RAMs *p* and *q*, if *q* is the result of replacing the first member of each

¹⁴ We will want to say this only if we take doughnut contexts to correspond to the intuitive notion of context that we employ in generating the intuitions about the truth values of hypothetical utterances. More on this below.




Fig. 1. How context gets in: Richard's account.

ordered pair in p with what f assigns to it, we say that p represents q under f , or $Rep(p, q, f)$. The context c of an attitude report determines an admissibility restriction, r_c , on mappings from name-types to representations; if f obeys the restriction, we say $Obeys(f, r_c)$. In a belief report, only certain representations are allowed to play the same role in a believed RAM as a given name-type plays in the RAM expressed by the that-clause. In Tom's case, we can map 'Cicero' to ρ_{Cicero} but not to ρ_{Tully} .

In a context c , the verb 'believes' has as its content the following *ascriptively* basic belief-relation ($Believes_c$) between agents A , RAMs p and mappings f (we ignore time):

$$\exists q [Rep(p, q, f) \ \& \ Obeys(f, r_c) \ \& \ Bel(A, q)].$$

That is, A *Bel*-believes some RAM to which f admissibly (on r_c) maps p . In contexts that provide different admissibility restrictions, then, 'believes' stands for different belief-relations. The content of a typical belief report 'A believes that S ', in context c , where ' S ' expresses RAM p_S , is the following:

$$\exists f \ Believes_c(A, p_S, f).$$

Thus, on Richard's account, the contextuality in belief reports is traced to the particular relation expressed in context by the attitude verb. A particular $Believes_c$ relation constrains which that-clauses can represent which of the agent's RAMs, by constraining which expressions can be used to represent – not refer to – which of the agent's mental representations. This is why we can say something true with (1) but not with (2); because of the particular $Believes_c$ relation we express in a normal context, the names 'Cicero' and 'Tully' can only represent ρ_{Cicero} and ρ_{Tully} , respectively. Figure 1 gives a schematic idea of the information about Tom's belief expressed in (1).

A consequence is that the aspects of context relevant to the contextuality

in which we are interested cannot change in mid that-clause. This is because context comes in only once, in determining the content of the attitude verb. Since what context provides is sensitive only to name-types, any two uses of the same name-type in a that-clause must be mapped to the same representation if the report is true. For instance, “Ann believes that Larry is here and Larry is reading”, can be true only if Ann has a belief involving the same representation of Larry twice over; in normal contexts, that seems right. But I think there are cases in which this result conflicts with strong intuitions. Consider the natural use of,

- (5) He’s falling for it; Cyril believes that John is John’s father.

On Richard’s account, unless Cyril is cognitively deficient or a believer in science fiction this statement is false.

Some comments about an example from Scott Soames suggest that Richard would hold that (5) is strictly speaking false, but pragmatically okay.¹⁵ I will suggest below that this kind of move is mistaken for the same reason Richard rejects similar moves in the naive Russellian approach on which coreferential substitution is always strictly speaking valid: if we do not need to abandon truth intuitions, we should hang on to them. Another possibility is to hold that what look like different occurrences of the name ‘John’ are really occurrences of different homophonic name-types. But surely this would be an objectionable proliferation of names; we have in this case none of the usual reasons for positing distinct homophonic name-types.

Since idiosyncratically different relations are expressed by ‘believes’ in different contexts, another consequence is that it is unlikely for agents in different contexts to talk or think about the same ascriptive belief-relation. Since I obviously can *truly* attribute to you the belief that Pierre believes that so and so, the truth conditions of such an iterated ascription cannot require what is unlikely: that the relation stood for by ‘believes’ in my context is the same as in yours. Thus, the embedded use of ‘believes’ must form an exception to the rule that expressions inside that-clauses contribute to the content of the clause what they would normally refer to in the context. For this reason, Richard holds that, to the contents of that-clauses in which it occurs, the verb ‘believes’ contributes not one of the contextually determined three-place belief-relations it otherwise expresses, but rather its *character* – the context-independent four-place relation *Believes^{char}* got by abstracting over contexts in a belief-relation.¹⁶

¹⁵ *Propositional Attitudes*, pp. 218–219.

¹⁶ *Propositional Attitudes*, pp. 245–246.

That this seems an *ad hoc* exception to a plausible general rule is one objection; others are that the RAMs expressed by that-clauses – Richard’s *propositions* – determine incomplete truth conditions (they are true or false only relative to a context), and that it becomes too easy for an iterated ascription to be true, as I will now explain.

My use in context *c* of,

(6) Tom believes that Pierre believes that London is pretty,

is assigned the following logical form:

$$\exists f[\text{Believes}_c(\text{Tom}, p, f)],$$

where *p* is the RAM expressed by my that-clause, “Pierre believes that London is pretty”, namely, the following RAM:

(7) $\langle\langle(\text{treatment of ‘}\exists g’)\rangle\rangle, \langle\langle\text{‘Pierre’}, \textit{Pierre}\rangle\rangle, \langle\langle\text{‘believes’}, \textit{Believes}^{\textit{char}}\rangle\rangle, \langle\langle\text{‘London’}, \textit{London}\rangle\rangle, \langle\langle\text{‘is pretty’}, \textit{being pretty}\rangle\rangle, \langle\langle(\text{treatment of ‘}g’)\rangle\rangle.$

Assume Tom knows that Pierre has two representations of London, and that Tom himself, not being confused, has only a single representation of London. Corresponding to *Pierre’s* two representations of London there should be two readings of my (6): I may mean that Tom’s belief is about Pierre’s travel-brochure representation, or Pierre’s home-town representation. But there is no way for Richard’s contextual machinery to distinguish these claims, because there is no constituent in the RAM (7) expressed by my that-clause, of which *Tom* has two representations.

In response to this difficulty, one might suggest that Tom has two different representations of ascriptive belief relations (*Believes_c*) that build in different restrictions. But this certainly does not seem necessary: if Tom knows about Pierre’s representations, then he can form beliefs (one of which we could report with (6)) simply about what RAMs Pierre *Bel*-believes, without ever representing any *Believes_c* relation between Tom and his own RAMs.

Richard also has available to him a scope ambiguity in the placement of the quantifier over the mapping *f*, but this will in general be of no help for the current difficulty. The difference between talking about the brochure and home-town kinds of beliefs is clearly not itself a difference in scope, so if the ambiguity is to help it must be that on both readings, the quantifier ‘ $\exists g$ ’ has wide scope, so that our claim is to the effect that there is a mapping such that Tom has a belief about *it*. This would allow us to distinguish Tom’s two kinds of beliefs only if we could restrict which of his representations of *mappings* can be represented by the variable ‘*g*.’

But Richard believes ('indeed, I insist') that there can never be restrictions on what representations variables can represent.¹⁷

Another response might go as follows.¹⁸ Recall that p itself is true or false only relative to a context. If the contextual admissibility restrictions allow only Pierre's travel-brochure representation (ρ_{tb}) to map to London, then the RAM p is true, in our context, if and only if Pierre believes that London is pretty in the way involving ρ_{tb} . In another ascriptive context, the truth of p would require that Pierre's belief involve his home-town representation (ρ_{ht}). So, with a change in the ascriptive context, the truth conditions of the RAM that we allege Tom to believe, change in just the right way to handle the different readings. But this reply is mistaken. What cognitive RAMs Tom *Bel*-believes does not change with our ascriptive context. Since the envisioned contextual shift does not affect Tom's *Bel*-beliefs, and does not involve changes in which of Tom's representations can represent constituents of p , it cannot lead to a change in truth value of our use of (6).

The iterated ascription problem, and that of John and John's father, are in no sense utterly decisive objections to Richard's account. But certainly if we can retain Richard's insights while resolving these problems, we ought to do so. And I think we can.

We can get from Richard's account to Perry's and mine by making three changes. First, we abandon the indexicality of attitude verbs. The information about representations is not built into the verbs' content, nor the content of any sentential expressions (overt ones, anyway), but comes about through *tacit* reference. Whatever entities are needed to provide this information (such as admissibility restrictions) are to be first-class propositional constituents in the claims made by attitude reports. We can take 'believes' to express the following relation among an agent A , a RAM p , a mapping, and a restriction r :

$$\exists q[\text{Rep}(p, q, f) \ \& \ \text{Obeys}(f, r) \ \& \ \text{Bel}(A, q)].$$

This helps get us around the problem with iterated reports. If the belief we attribute to Tom is about not only Pierre, the belief relation and a RAM, but also about a specific admissibility restriction, then we get our different readings from different restrictions: according to Tom's beliefs, Pierre may believe the proposition that London is pretty under one ad-

¹⁷ *Propositional Attitudes*, p. 152. Even if we allow such restrictions, it is implausible to suppose, for the same reasons as before, that Tom must have a token representation of a mapping for the report to be true.

¹⁸ Thanks to two anonymous referees for suggesting that I consider this.

missibility restriction (say, one that forces Tom's single representation of London to map to Pierre's travel-brochure representation), but not under another restriction (one that forces it to map to the home-town representation). Allowing the information about representations to arise from tacitly referred to propositional constituents also rids us of the awkwardness of holding that propositions themselves are true or false only relative to admissibility restrictions and rids us of the exception to the rule about expressions contributing what they stand for – all without sacrificing anything of value.¹⁹

Having made this change, cases like “Cyril thinks John is John's father” can be accommodated with a second revision. There is no reason to think that the context of a belief report must provide a single, monolithic entity (like Richard's admissibility restriction on mappings) that includes all the information about representations needed for the report. Instead, we can hold that context provides something corresponding to the first use of ‘John’ and something different corresponding to the second use. To each use of a name in a that-clause, we will say, context associates a condition on the representation that that *use* of the name can represent – the condition associated with a use of a name constrains the representation that must play in the agent's attitude the same role that the use plays in the that-clause (or perhaps rather the same role the RAM pair contributed by the use plays in the RAM expressed by the that-clause). The two uses of ‘John’ are associated with different conditions on representations: the first representation must be connected to Cyril's perception of John; the second must be connected to a more stable, ordinary set of beliefs about John. The difference in these conditions is keyed not by different names but simply by the different roles the uses of ‘John’ play in expressing the proposition (or RAM); the hearer can figure this out easily, so the speaker does not need to use different *terms* for John to flag the difference in conditions on representations. For each different role an expression-value pair plays in the RAM expressed by a that-clause, then, context may provide a condition on the agent's representation that must play the same role in her belief, if the report is to be true.

It might be unintuitive to formalize this in a semantics in which arbitrary sentences are evaluated at ‘doughnut’ contexts, since it is unintuitive to think that contexts contain anything to do with roles of things that might be uttered in them. But the ‘doughnut’ notion of context is unintuitive in

¹⁹ This is not quite correct – we do sacrifice the implausibility of having to claim that Tom has a mental representation of a restriction on mappings between his own and Pierre's representations. But this drawback vanishes in the change that is about to come.

this way in lots of cases where utterances themselves influence semantically relevant features of their context. The problem is that the intuitive notion of a context in which many different sentences can be used does not include in it all contextual material that may be relevant to interpreting a sentence used in it – because the sentence itself can contribute such material. The awkwardness is avoidable: instead of defining contexts and evaluating sentences relative to them, we can simply evaluate possible utterances based on their various features – including their ‘context-modifying’ features.

Having come this far, we notice that the usefulness of the name-types in that-clause RAMs is gone. They allowed us to mark the jobs different representations were supposed to perform in the ascribed belief. But name-types were too coarse-grained for this, as the John/John’s father case shows, so we key instead to roles that expression-value pairs play in RAMS. But for every role played by such a pair in a RAM there is a corresponding role in the simpler Russellian proposition got by replacing each expression-value pair in the RAM with just the value. So we can eliminate that-clause RAMs in favor of our old friends, Russellian propositions. And in place of the RAMs that are *Bel*-believed, we can take simpler entities, namely *thoughts*, which are what you get from mental RAMs by replacing each representation-value pair with just the representation. A thought has the same structure as its content, a Russellian proposition.²⁰ For any entity filling a particular role in the content of a thought, there is a representation *playing a parallel role* in the thought; say that the representation is *responsible* (in the thought) for the parallel role in the content proposition.

The account we end with is roughly that proposed by Perry and myself. Consider *u*, a statement of ‘*A* believes that *S*’, where *S* expresses Russellian proposition *p*. The report is true just in case there is a thought τ such that:

- (i) the agent *Bel*-believes τ ,
- (ii) τ has content *p* (and thus is of the same structure as *p*), and,
- (iii) for each role *r* that an entity plays in *p*. there is a representation ρ that both is responsible in τ for *r*, and meets the condition C_r that is tacitly referred to in connection with role *r*.

²⁰ This notion of sameness of structure is probably intuitively clear; spelling it out thoroughly would require a very extended digression.




Fig. 2. How context gets in: the Crimmins-Perry account.

That is,

$$\exists \tau [Bel(A, \tau) \ \& \ ContentOf(p, \tau) \ \& \ \bigwedge_{r \text{ in } p} \exists \rho (Responsible(\rho, \tau, r) \ \& \ Meets(\rho, C_r))].$$

Figure 2 gives a schematic idea of the information about Tom’s belief expressed (on this account) in (1).

This gets around the John/John’s father problem by allowing different conditions on representations to be associated with the different occurrences of John in the proposition that John is John’s father. That is to say, if the proposition is:

$$\langle\langle father; John, John \rangle\rangle,$$

and the two roles filled in it by John are r_1 and r_2 , then context can provide different conditions, C_{r_1} and C_{r_2} , on the representations that must play parallel roles in a believed thought, if the report is true.

The semantics gets around the troubles about embedded ascriptions by taking the information about *how* the agent allegedly represents things to issue in constituents of the proposition expressed by a belief report – constituents that someone must represent if they are to believe what is expressed by the report. Consider again sentence (6). In a context in which Pierre’s travel-brochure representation (ρ_{tb}) of London is uniquely relevantly salient, the embedded utterance of “Pierre believes that London is pretty”, would express the proposition, call it p , that Pierre believes the proposition that London is pretty in a way such that ρ_{tb} is responsible for the subject role. Thus, the proposition Tom is reported to believe in (6) contains ρ_{tb} as a constituent – p is about ρ_{tb} . If we center the discussion around Pierre’s home-town representation (ρ_{ht}), then we attribute to Tom with (6) belief in a proposition about ρ_{ht} . Thus, as required by (at least my) intuitions, (6) can be used to claim in two different contexts that Tom believes two different things about Pierre.

Of course I have not shown that this semantics is wholly adequate. But I hope the close connection with Richard’s work has become sufficiently

clear that many of his excellent arguments in favor of his own view can be seen to apply here as well.²¹

REFERENCES

- Asher, N.: 1986, 'Belief in Discourse Representation Theory', *Journal of Philosophical Logic* **15**, 137–89.
- Barwise, J. and J. Perry: 1983, *Situations and Attitudes*, MIT Press, Cambridge.
- Church, A.: 1950, 'Carnap's Analysis of Statements of Assertion and Belief', *Analysis* **10**, 98–9.
- Crimmins, M. and J. Perry: 1989, 'The Prince and the Phone Booth: Reporting Puzzling Beliefs', *Journal of Philosophy* **86**, 685–711.
- Kamp, H.: 1990, 'Prolegomena to a Structural Theory of Belief and Other Attitudes', in C. A. Anderson and J. Owens (eds.), *Propositional Attitudes*, pp. 27–90. CSLI, Stanford, CA.
- Kripke, S. A.: 1979, 'A Puzzle About Belief', in A. Margalit (ed.), *Meaning and Use*, Reidel, Dordrecht, pp. 239–83.
- Kripke, S. A.: 1980, *Naming and Necessity*, Cambridge: Harvard University Press, Cambridge.
- Quine, W. V.: 1960, *Word and Object*, MIT Press, Cambridge, Mass.
- Richard, M.: 1983, 'Direct Reference and Ascriptions of Belief', *Journal of Philosophical Logic* **12**, 425–52.
- Richard, M.: 1987, 'Attitude Ascriptions, Semantic Theory, and Pragmatic Evidence', *Proceedings of the Aristotelian Society* **87**, 243–62.
- Richard, M.: 1990, *Propositional Attitudes: An Essay on Thoughts and How We Ascribe Them*, Cambridge University Press, Cambridge.
- Salmon, N.: 1986, *Frege's Puzzle*, MIT Press, Cambridge.
- Schiffer, S.: 1987, *The Remnants of Meaning*, MIT Press, Cambridge.
- Soames, S.: 1987, 'Direct Reference, Propositional Attitudes and Semantic Content', *Philosophical Topics* **15**, 44–87.

Department of Philosophy,
Cornell University,
Ithaca, NY 14853
U.S.A.

²¹ My thanks to David Martens, Joe Moore, Mark Richard, Rob Wilson, and referees for helpful discussion and comments.