

Scott Soames

36. This account has two precursors. The first is the introduction of structured meanings (characters) in Richard's 'Direct Reference and Ascriptions of Belief'. The second is a somewhat different use of structured intensions suggested by David Kaplan (personal correspondence) in response to Richard. The account in the text is designed as an improvement on those treatments intended to capture certain insights that motivated them.

37. One could have versions of these theories in which semantic contents were not objects of the attitudes, but only by foregoing the strong motivation the attitudes provide for these theories.

38. Strictly speaking, the intensions of directly referential terms in (3b) and (8c) should be constant functions from circumstances to objects, rather than objects themselves. However, this does not affect the issues at hand.

39. The significance of this result is enhanced by the defence, in section V, of the consequence (23) of (A2), (A4), and compositionality. However, it should be noted that analogous results involving the encoding of structure in objects of the attitudes can be established using the weakenings in IV.

40. I leave it open whether there are such expressions.

41. Here I am using '(35a)', '(35c)', and '(36)' not as names, but as abbreviations for the sentences they normally name.

42. An analogous argument can be constructed regarding assertion.

Chapter 40

The Prince and the Phone Booth: Reporting Puzzling Beliefs

Mark Crimmins and John Perry

In Mark Twain's *The Prince and The Pauper*, Tom Canty and Edward Tudor decide to change lives for a day, but fate intervenes, and the change goes on for a considerable period of time. The whole story turns on what people believe and do not believe about the two boys, and intelligent reader, unexposed to recent philosophy of language and might could probably describe the key facts of the story with some confidence. Such a reader might explain why Miles Hendon, a penniless nobleman who encounters a boy dressed in rags, does not bow to the Prince, noting:

(1) Miles Hendon did not believe that he was of royal blood.

And such a reader might ward off the implication that Miles was a fool ignorant by noting that Miles shared the dominant conception Edward Tudor:

(2) Miles Hendon believed that Edward Tudor was of royal blood.

One of our main claims in this paper is that such a reader would be right on both counts. In this we depart from a recent trend to explain the apparent truth of statements like (1) as an illusion generated by pragmatic features of such claims. Accounts of belief reporting given by Jon Barwise and John Perry,¹ Scott Soames,² and Nathan Salmon³ have employed this strategy of denying the accuracy of our strong intuitions about truth and falsity. Here, we shall present an account that does not ignore pragmatic features, but assigns to them a more honorable role. They do not create an illusion, but help to identify the reality the report is about. Our account honors the intuition that claims (1) and (2) are true.

First appeared in *The Journal of Philosophy* 86 (1989): 685–711. Reprinted by permission *The Journal of Philosophy* and the authors.

Since 'Edward Tudor' in (2) and 'he' in (1) both refer to Edward Tudor, this seems to commit us to some version of the doctrine of *opacity*.⁴ Specifically, we are committed to the view that, if our reader were to say either of the following, in the same circumstances, he would be incorrect:

(1') Miles Hendon did not believe that Edward Tudor was of royal blood.

(2') Miles Hendon believed that he was of royal blood.

The doctrine of opacity has been thought incompatible with two others, to which we also are attracted: the first, *direct reference*, is that the utterance of a simple sentence containing names or demonstratives normally expresses a "singular proposition"—a proposition that contains as constituents the individuals referred to, and not any descriptions of or conditions on them; the second, *semantic innocence*, is that the utterances of the embedded sentences in belief reports express just the propositions they would if not embedded, and these propositions are the contents of the ascribed beliefs.⁵

Direct reference and semantic innocence are well-motivated by many considerations in the philosophy of language. But if direct reference and semantic innocence are correct, then it seems that opacity must not be: the substitution of 'Edward Tudor' for 'he' in (1) [or vice-versa in (2)] should be completely legitimate. The name and the demonstrative refer to the same object. There is just one proposition, belief in which is denied by (1) and affirmed by (2), the "singular" proposition, which we shall represent in this way:

«Being of royal blood; Edward Tudor»

The example is typical of many doxastic puzzle cases in the literature—puzzles because they seem to reveal a conflict among the three very plausible doctrines. We hold all three, however.

When we substitute 'Edward Tudor' for 'he', the words change while the proposition expressed by the embedded sentence stays the same. If we think that belief is a relation to propositions and not words, the apparent change in truth value of the whole report seems puzzling. We are likely to focus on the most apparent change, the change in words, as the clue to the mystery.

The most famous doxastic puzzle case, due to Saul Kripke,⁶ has not to do with substitution, however. Kripke describes a case in which a Frenchman Pierre first hears of London, comes to believe it is pretty, then moves to London, and, not connecting it to the city he has heard about (under the French 'Londres'), comes to believe it is not pretty. He does not change his mind about the city he has heard of, but simply does not connect the "two" cities. We have on sentence

(3) Pierre believes that London is pretty.

that we seem to be able to use, when reflecting on different parts of the story, to say something true and to say something false. The words have not changed. What has?

What changes in this case, and in every other doxastic puzzle case, what we are talking about. Pierre has two different notions of London which play very different roles in his beliefs. An assertion of (3) is true if it is about one of them, false if it is about the other. An ordinary doxastic puzzle case uses a change in words to precipitate the change in the subject matter of the utterance. Kripke spells out the details of his case so clearly that our focus gets redirected without a change in the wording of the report. We shall return to these claims about belief reports in the next section.

One of Pierre's beliefs was caused by his acceptance of the stories he heard about London. It has the content that London is pretty, and it leads him to cherish the prospect of someday visiting that city. This belief also causes him to affirm, in French, "Londres est jolie," in discussions about the city he has heard of.

Also, Pierre has a different belief which was caused by his displeasure with his new surroundings, which has the content that London is not pretty and which causes him to affirm, "London is not pretty," in discussions about his home.

It is a commonplace to distinguish these two beliefs. We think it is often not sufficiently appreciated, however, that the beliefs so distinguished are concrete cognitive structures. Focusing on this fact provides the basis for our account of belief and for our solutions to the various doxastic puzzle cases.

These are the key features of our theory of beliefs:

- (i) Beliefs are concrete cognitive structures: they are particulars that belong to an agent, come into existence, endure, and go out of existence
- (ii) Beliefs are related to the world and to other cognitive structure

Mark Crimmins and John Perry

and abilities in a way that allows us to classify them by propositional content.

Beliefs, since they are cognitive particulars or “things in the head,” are not things that are believed; they are not in any sense the objects of belief. The propositions believed are the objects of belief. An agent believes some proposition in virtue of having a belief with that content. Many agents can believe the same proposition, so propositions are public; they also are abstract. Beliefs are neither public nor abstract; they are concrete particulars that belong to agents just like arms, headaches, and bouts of the flu. A belief comes into existence when an agent forms it; it is not the sort of thing that is around for the agent to adopt. Agents believe the same thing, a proposition p , when each has a belief with p as its content. This is not an analysis of reports of “believing the same thing”—which are not always so simple to unpack—but a clarification of what we mean by objects of belief.

To countenance beliefs as particulars is not to deny that there are interesting systems of abstract objects which might be used to classify them, such as meanings, Fregean senses, intensions, characters, or the like. But in addition to having these abstract features, beliefs, like other concrete particulars, have lots of other features, both intrinsic and relational, many of which can in some cases be relevant to explaining how we talk about beliefs in belief reports. In particular, we often exploit facts about the causes and effects of beliefs, a point to which we shall return.

There are a number of reasons to allow ourselves to speak of particular beliefs, rather than just of a belief relation between a person and an abstract object of some kind. There is, first, the attraction of having entities that can occupy causal roles with respect to perception, reasoning, and action. As Jerry Fodor and others have argued at length, structured concrete particulars or “token” mental entities go a long way toward explaining the roles of belief, desire, and so on in cognition. There is, second, the fact that the most plausible statements of materialist intuitions about the mind are formulated in terms of particular mental entities. And, third, there is the problem that belief puzzles repeatedly have emphasized: it seems that, for any natural way of classifying beliefs with abstract objects, we can find examples in which a single agent, at a single time, is belief-related to one such abstract object twice over. These are cases, we would like to be able to say, in which an agent, at a time, has two beliefs classified by the same sense, meaning, or whatever. Classifying beliefs only with abstract meanings, senses, and so on is like classifying drops of

water only with intrinsic properties. Kant argued against Leibniz the intrinsic properties of particulars will not always provide us with sufficient material for their individuation. Kant took it as obvious that there can be two exactly similar drops of water; the puzzle cases make it clear that there can be two beliefs sharing the abstract features which one or another theory of belief claims to be central.⁷

Beliefs, then, are particulars that bear complex causal relations to an agent’s perceptions, actions, and other cognitive structures and abilities. The story of the causal properties of beliefs will be closely bound to the story of how and why beliefs can be classified with propositional content. A belief constrains an agent’s reasoning and action in a way that is conducive, if the belief’s content is true, to the agent’s getting what she wants.

The ground-level facts behind belief are simply the facts of agent having beliefs. There is a basic relation $B(a, b, t)$ that holds of an agent, belief, and a time just in case b is a belief that belongs to the agent a at time t .

Normally, a belief has a propositional content. So there is a partial function $\text{Content}(b, t)$ that, for a belief b and time t at which b exists yields the content of b . The content of a belief will be determined by the “internal” structural properties of the belief plus its real connections to things and circumstances in the world and to the agent’s other cognitive structures and abilities.

If an agent a at time t has as an object of belief the proposition p , then there is a belief b such that:

$$B(a, b, t) \ \& \ \text{Content}(b, t) = p.$$

So much is all that is really needed for a theory of belief adequate for broad explanation of the doxastic puzzle cases, and so we are tempted to stick with just the minimal theory of beliefs given so far. The minimal theory is compatible with a wide class of views about beliefs, about propositions (or contents), and about central issues in theories of representation, practical reasoning, and inference. The crucial features of the semantics we give for belief reports, and the resulting solutions for the troubling cases, are therefore to some degree theory-neutral. But we want to present a slightly more detailed, if still simple-minded, theory of belief which satisfies the demands of the minimal theory and which yields sufficiently rich account of just how the puzzling belief reports work.

Beliefs are structured entities that contain ideas and notions as constituents. Ideas and notions, like beliefs, are on our view concrete cog-

itive particulars. So there is no such thing as agents having the same idea or notion, but only similar ones. Admittedly, the technical use we make of these terms involves a departure from what we ordinarily say about 'ideas' and 'notions', or at least represents a choice among the many different ordinary uses of these terms. On our use of the terms, there are no notions and ideas that agents do not have, any more than there are headaches that no one has. The difference between notions and ideas is the difference between an agent's "ways of thinking" about individuals versus properties. The properties and things of which ideas and notions are ideas and notions we call their *contents*. We shall explain in a moment how the contents of ideas and notions help determine the contents of beliefs.

What determines the content of an idea or notion? For example, what is it about Myles's notion of the poorly-dressed boy which causes it to be a notion of Edward as opposed to another boy? The crucial fact is that it was Edward with whom Myles was confronted when he formed this notion. Edward played the right part in the causal origin of the notion; the notion was formed in order to keep track of information about Edward—that is what makes him its content. So the content of an idea can depend on its external properties, like facts about its origin. The very same notion might have been a notion of a different person, had someone other than Edward figured in its origin.

There is a close parallel between this view of the contents of ideas and causal views of the semantics of names. A speaker can refer to an individual with a name, it is held, because that individual figured, in the right way, in the speaker's adoption of the name as a tool of reference.⁸

The content of an idea is not always fixed once for all by facts about the circumstances of the idea's origin. Some ideas are *context-sensitive*, in that their contents change with changes in the agent's circumstances. The context sensitivity of ideas is analogous to that of demonstratives in language. David Kaplan⁹ has proposed that there is associated with each demonstrative a *character*, a function that specifies how the content of a demonstrative depends on the circumstances surrounding its use. The content of a use of the word 'you', for example, is the person being addressed in the circumstances of the utterance. Analogously, an agent *a* may have an idea I_{addr} of "being the one I am addressing." The property which is the content of this idea changes with changes in circumstances as follows:

In any circumstances in which a person *b* is being addressed by *a*, the content of *a*'s idea I_{addr} is the property of being *b*.

Undoubtedly, each of us has a "you" idea, the content of which is determined functionally in this way. We do not share ideas, but we have it with the same *semantic role*. An idea's semantic role is the function that determines the idea's content based on the agent's circumstances. Semantic roles for ideas are a bit like characters of expressions; some ideas have semantic roles that are context-sensitive, others have semantic roles that are constant functions—their contents do not vary with changes in context.

So there are two ways in which an agent's external circumstances might be relevant to determining the content of an idea. First, the facts surrounding the origin of the idea may fix its content once for all. Second, the idea's semantic role may be sensitive to changes in the agent's circumstances—the content of the idea may vary from occasion to occasion. So an idea may exhibit origin sensitivity, context sensitivity, or both.

Myles's idea of red is certainly not context-sensitive. It may be deemed origin-sensitive, whether one supposes that his idea stands for red nately, or because of some original assignment of ideas of colors early in Myles's life. Myles's idea of being past, in contrast stands for different properties as his life unfolds; at each time *t*, the idea stands for the property of occurring before *t*. This idea is certainly context sensitive, and it may not be origin sensitive. And Myles's notion of Prince Edward formed upon hearing of the newborn Prince, is origin-sensitive, but it is not context-sensitive.

Notions are the things in the mind that stand for things in the world. A notion is a part of each of a collection of beliefs¹⁰ (and of other mental structures, such as desires and intentions) that are internally about the same thing. This is not a definition of 'notion', but just a central fact about notions—sharing a notion is what it is for beliefs to be internally about the same thing. An agent may occasionally (and will in many of the examples) have several notions of a single individual. This can happen in two ways. First, in cases of misrecognition and "failure to place," an agent may have two notions of an individual which he does not link together; such an agent is guilty of no internal inconsistency. But also an agent can retain two notions of an individual, while linking them, in a way one does when one recognizes that "two" of one's acquaintances is actually a single individual. Why might two notions be retained when

such a recognition takes place? One reason for this would be to allow the possibility of easy revision in case the “recognition” was in error. But an agent can also burn his bridges and merge two notions into a single notion. Two beliefs, then, can be internally about the same thing in two ways: by sharing a notion, and by containing notions that are linked.

For the purposes of this paper, we assume that each belief involves a single k -ary idea and a sequence of k notions.¹¹ To represent the structure of such a belief, we write:

$$\text{Structure}(b) = \langle \text{Idea}^k, \text{Notion}_1, \dots, \text{Notion}_k \rangle.$$

Each belief has as its content the proposition that the objects its notions are of have the property or stand in the relation, that its idea is of:

$$\text{Content}(b, t) = \langle \langle \text{Of}(\text{Idea}^k, t); \text{Of}(\text{Notion}_1, t), \dots, \text{Of}(\text{Notion}_k, t) \rangle \rangle.$$

The structures of beliefs are individuated not simply by the ideas and notions involved in them, but also by which argument places of the ideas the various notions fill. Thus the order of the notions in our representation of the structure of the belief reflects an assignment of notions to the argument places of the associated idea.

To be clear about the relation between beliefs and their contents, we need to introduce some new concepts.

A belief b associates an idea I with a notion n at an argument place pl :

$$\text{Associates}(b, I, n, pl)$$

The belief that Tom fired Mary and the belief that Mary fired Tom differ in which places are associated with which notions, even though the ideas and notions involved are the same.

An argument place of an idea is intimately connected with an argument role of the relation which is the content of the idea, and so with an argument role in the content of the beliefs of which the idea forms a part.¹² If we were to consider complex cases, spelling out this relationship might be a matter of some delicacy, but we shall take it to be straightforward here. We shall say that an argument place pl_I of an idea I generates an argument role r_p of a proposition p (an example below will make this clearer):

$$\text{Generates}(pl_I, r_p)$$

Finally, a notion is responsible for which object occupies an argument role of the content of a belief, when the belief associates it with an idea at

the argument place which generates the argument role in the content of the belief:

$$\text{Responsible}(n, r, b) \Leftrightarrow_{\text{def}} \exists I, pl \text{ Associates}(b, I, n, pl), \text{ and Generates}(pl, r)$$

When a notion in a belief is responsible for filling an argument role of belief's content, it fills the role with its own content, the object of which the notion is a notion.

To give an example: Arthur's belief that Yvain smote Kay involves Arthur's idea for smiting, I_s , and his two notions of Yvain and Kay, n_Y and n_K . The idea I_s has two argument places, one (pl_+) for the smiter and one (pl_-) for the smitten. In Arthur's belief (call it b), a notion n_Y is associated with argument place pl_+ of I_s , and n_K is associated with pl_- . The content of b is the proposition p , where:

$$p = \langle \langle \text{Smote}; \text{Yvain}, \text{Kay} \rangle \rangle$$

The relation “smote” has two argument roles, one (r_+) for the smiter and one (r_-) for the smitten; these are also argument roles of the proposition p . In p , Yvain fills r_+ and Kay fills r_- of the “smote” relation. Since associates pl_+ with n_Y , and pl_- (the smiter in I_s) generates r_+ (the smite in p), we say that, in b , n_Y is responsible for filling r_+ in p . Arthur's notion of Yvain is responsible in b for determining who fills the argument role r_+ in p . And n_Y provides its content, Yvain, to fill that argument role. Figure 1 should make this clear.

Notions and ideas are key figures in our commonsense “folk” model of cognition. The recurring appearance in philosophy of such things as concepts, senses, ways of thinking, names in a language of thought, mental file folders, and other such devices reflects a firm intuition about the mind, namely, that having beliefs about an individual means having beliefs involving an internal something that is one's cognitive “fix” on the individual. As we have said, we think the correct way to express this intuition demands reference to cognitive particulars that are involved in belief desires, and so on. Now, this leaves a great deal open about just what kinds of thing our notions and ideas are. For all we have said, notions and ideas might be—or might have been—particular words in a language of thought, physical objects like file folders, or things with more of a dispositional character, like the process underlying the disposition of an agent to have a specific “pattern of neural activation” in certain circumstances. And, whichever of these kinds of thing our notions and ideas are




Figure 1
Arthur's belief and its content

they certainly may be classifiable with senses, property clusters, intentions, and so on. We want our "notions and ideas" to capture what is in common among all these very different models of cognition: there are things shared by different beliefs which explain the internal way in which beliefs must be about the same object or property.

On this theory, one can have two beliefs with exactly the same content or with diametrically opposed contents, such that there is no significant causal relation between them—because they involve different notions. This is a feature of all of the problematic examples that we shall consider. There is nothing particularly puzzling about this—and in fact there is nothing particularly puzzling about any of the examples that we discuss, so long as we simply consider the beliefs and not the reporting of them. Nevertheless, it is a good idea to go over the examples in some detail, for it is these details which our semantic account pays more attention to than others of which we know.

Consider the Prince and the Pauper. Miles Hendon has two notions of Edward Tudor. They have quite different circumstances of origin. One Miles has had for a long time. It is associated (in his beliefs) with his ideas as being a Prince of England, being named 'Edward Tudor', being rich, not being a pauper, not looking like a pauper, not being likely to run into (me) on an average day, and the like. The beliefs with this notion as a constituent influence Miles's behavior when confronted with ordinary sorts of information about Edward Tudor. When he reads an article in the

Times, for example, it is beliefs with this notion as a constituent which is affected.

His other notion was formed when he saw Edward being set upon an angry mob—angry because Edward, dressed in rags, had been proclaiming himself to be Prince. This notion is associated with ideas of being out of his mind, being dressed like a pauper, and not being of royal blood. The beliefs involving this notion, and not those involving his old notion of Edward Tudor, influence Miles's behavior toward Edward and Edward's assertions during the period he is associated with him as a comrade, until that point, toward the end of the story, when Miles merges his two notions and comes to believe that Edward the pauper is Edward the Prince.

Perhaps the ultimate doxastic puzzle case is Mark Richard's puzzle about the woman in the telephone booth:

Consider *A*—a man stipulated to be intelligent, rational, a competent speaker of English, etc.—who both sees a woman, across the street, in a phone booth, and speaking to a woman through a phone. He does not realize that the woman whom he is speaking to—*B*, to give her a name—is the woman he sees. He perceives her to be in some danger—a run-away steamroller, say, is bearing down upon her phone booth. *A* waves at the woman; he says nothing into the phone (ibid. p. 439).

The man has two distinct, unlinked notions of the woman. Via one, he believes that she is in danger. This is the notion which arose in virtue of his visual perception of her, and which is associated with an idea of being in grave danger. It is this notion which is involved in the beliefs that motivate his waving out the window. The second notion is an older one, assuming the woman is an old acquaintance. It is associated with an idea of being the person addressed, and not associated with ideas of being the person seen or being in danger. Hence, the beliefs involving this notion do not motivate a warning.

Let us return to Kripke's case. Pierre has the same misfortune as Miles and the man on the phone: he has two notions of the same thing. He has one notion of London which is linked to his memories of the stories and to his use of the word 'Londres'. He has another, unconnected notion of London which is influenced by his perceptions and memories about his present surroundings and which influences his use of the word 'London'. He has a belief associating the former notion with his idea of being pretty but has no belief associating the latter notion with this idea. In fact, Pierre associates an idea of being ugly with the latter notion.

Our basic idea is simple: a belief report claims that an agent has a belief with a certain content. But the basic idea, unembellished, will not allow us to hold the family of views we want to defend. For (2) and (2') would claim that Miles Hendon had at least one belief with the content

«Being of royal blood: Edward Tudor»

while (1) and (1') would deny this—thus contradicting our truth intuitions and the doctrine of opacity.

But our embellishment is also simple. When we report beliefs, there is always some further condition that a belief with the specified content is claimed to meet. The belief report is true, only if a belief meeting that further condition has the right content. What may be novel is our insistence that this additional requirement is part of the proposition expressed by the belief report. Thus, it is a condition on the truth, not merely the felicity, of the report.

Consider (1). In context, (1) provides an explanation of why Miles Hendon did not treat someone he was looking at in a certain way—a way that would have been compulsory for Miles, given the status of that person. We are interested in the content of only those beliefs which motivated Miles's behavior, the beliefs which involve the notion of Edward which arose when Miles saw him being threatened and which explain Miles's treatment of him. The existence of such a notion is clear from the description of the incident. We know that Miles is perceiving Edward and interacting with him on the basis of what he, Miles, perceives. Our view is that, in reporting beliefs, we quite often are talking about such notions, although our belief reports do not explicitly mention them. The general solution to the puzzles is to allow a condition on particular beliefs, over and above a content condition, to be part of the claim made. The version of this strategy we shall pursue here is to take this further condition always to be a specification of the notions that are supposed to be involved in the ascribed belief.

We shall say that a notion that a belief report is about is an *unarticulated constituent* of the content of the report—it is a propositional constituent that is not explicitly mentioned. We shall distinguish another kind of belief report, and say more about the notion of unarticulated constituents in a moment. But first let us see what the semantics of this sort of belief report looks like.

From our account of beliefs, we have the following concepts:

$B(a, b, t)$: b is a belief that belongs to agent a at time t .

Content $(b, t) = p$: p is the content of belief b at time t .

Responsible (n, r, b) : $\exists I, pI$ Associates (b, I, n, pI) , and Generates $(pI, r$
form: We take a belief report to be an utterance u of a belief sentence of

A believes that S

where A is a singular term and S is a sentence. We assume a semantics the use of the embedded sentence, so that $Con(u_S)$ (the content of u_S the proposition expressed by the utterance of u corresponding to S Where u is a belief report at t which is about notions n_1, \dots, n_k, z $p = Con(u_S)$).

$Con(u) = \exists b[B(a, b, t) \wedge Content(b, t) = p \wedge \bigwedge_{n, r, p} Responsible(n, r, t$

The claim made by the belief report is that the agent a has a belief with the content p , involving the notions n_1, \dots, n_k (in a certain way) This claim entails the proposition that a has a belief with the content but the truth of that proposition is not sufficient for the truth of the report—the report says more than that about the ascribed belief.

We shall say in such cases that the notions that the belief report about are *provided* by the utterance and its context. Note that the provided constituents of the report's content are not existentially quantified. Let us see how this theory works with Miles, Edward, and our intelligent reader. We take our reader to be talking about n_{vis} , the notion Miles acquires of Edward from visually perceiving him on the occasion of the rescue. $Con(u_S)$ is just the proposition «Being of royal blood; Edward». So our reader is saying with (1) that there is no belief that associates Miles's idea of the property of being of royal blood with Miles's notion n_{vis} . He is not contradicting any proposition that Miles has some other notion of Edward Tudor which is so associated.

And, in fact, a proposition of this latter kind might be just what our reader intends to claim with (2). Imagine the case in which he reads the Miles Hendon is shouting, while treating Edward as a mad fool, «Prince Edward is a man of royal blood, you fool, who would not dress in rags.» Our reader might intend to say, of the notion involved in the beliefs that motivate this behavior, that it both is of Edward Tudor and is associated with the idea of being of royal blood.

If so, our reader would surely be consistent, direct, and innocent. On the one hand, the proposition he in turn denies and affirms Miles's belief in is just the singular proposition that contemporary theories of direct reference assign to the utterances of "Edward Tudor is of royal blood" and "He is of royal blood" in the described contexts. On the other, the denial and affirmation are completely consistent.

We have claimed that, in belief reports, an n -ary relation is reported with an n -minus-one-place predicate. On our account, the complex relation invoked in belief reports is a four-place relation: an agent believes a proposition at a time relative to a sequence of notions. But there is no argument place in the 'believes' predicate for the sequence of notions. The notions are unarticulated constituents of the content of the report.

Propositions have constituents. The proposition that Yvain smote Kay has Kay as a constituent—Kay himself is in that claim. When Arthur says, "Yvain smote Kay," there is no great mystery about why Kay, rather than someone else, is part of the claim Arthur makes: Arthur uses the name 'Kay', which, as he uses it, refers to Kay. Kay is the content of Arthur's utterance of 'Kay'. This is what it is to be an articulated constituent of the content of a statement.

It is very common in natural languages for a statement to exploit unarticulated constituents. When we consider the conditions under which such a statement is true, we find it expresses a proposition that has more constituents in it than can be traced to expressions in the sentence that was spoken. Each constituent of the content which is not itself the content of some expression in the sentence is an unarticulated constituent of the content of the statement.

We report the weather, for example, as if raining and snowing and sleeting and dark of night were properties of times, but they are one and all relations between times and places. If I say, "it is raining," you understand me as claiming that it rains at that time at some place the context supplies. It often is, but need not be, the place of utterance. If I am talking to a friend in Kansas City on the phone, or watching news reports about the continuing floods in Berkeley, you may understand me to be talking about those places rather than the place where we both are.

The phenomenon of unarticulated constituency is similar to that of indexicality in the reliance on context. But the two phenomena should not

be conflated. If we say, "It's raining here," an expression in our statement identifies the place. The place is articulated in a context-sensitive way. In the case of indexicals, expression and context share in the job of identifying the constituent, according to the conventional meaning or character of the indexical. In a case of underarticulation, there is no expression to determine the constituent in this way.

It would be misleading, however, to say that, in the case of unarticulated constituents, the context alone does the job. The whole utterance—the context and the words uttered—is relevant to identifying the unarticulated constituent. Thus, a change in wording can affect the unarticulated constituent, even though it is not a change in an expression that designates that constituent. Suppose I am in Palo Alto talking on the phone to someone in London; we both know that it is morning in Palo Alto and evening in London. If I say, "It's exactly 11 A.M.," I will be taken to be talking about the time in Palo Alto; if I had said, in the same context, "It's exactly 8 P.M.," I would be taken to be talking about the time in London.

The important principle to be learned is that a change in wording can precipitate a change in propositional constituents, even when the words do not stand for the constituents.

Unarticulated constituency is one example of the incrementality of language. In the circumstances of an utterance, there always is a great deal of common knowledge and mutual expectation that can and must be exploited if communication is to take place. It is the function of the expression uttered to provide just the last bit of information needed by the hearer to ascertain the intended claim, exploiting this rich background. What is obvious in context we do not belabor in syntax—we do not articulate it.

This is by no means to transgress the intuition of the systematicity of language which is commonly reflected in principles of "compositionality." Since we finite creatures are able to make and understand a potential infinity of claims, there must be systematic features of our statements which explain our infinite abilities in something like a combinatorial fashion—in terms of our more finite abilities to understand the contributions of specific features of statements toward the claims made. But there is no reason to assume that these features of statements must all involve syntactic expressions. It is just as systematic for a form of speech, like a belief report or a report of rain, to call for a propositional constituent that meets, say, certain conditions of relevance and salience, as it

description of, nor a condition on, a place. In the same way, on our view a belief reporter refers to an agent's notions. We have chosen not to take this way in our official account only to avoid being read as claiming that notions are referred to by the reporter's words.

A difficult issue facing all views of direct reference, and ours in particular, is the need to make sense of intuitions about truth and falsity: cases of reference failure. This problem is especially acute for our account in some cases of denials. Consider the following example. A blind man facing in the direction of a distant building. Someone, unaware of the man's blindness, says, "He believes that that building is far away." Or normally would take this report to be about the notion the man has as a result of his current visual perception of the building. The speaker is trying to refer (though not with a word) to such a notion, to provide such a notion for the report to be about. But of course there is no such notion in this case. Is this report false, or, owing to a failure of tacit reference, does it fail to express a proposition? Certainly, we ordinarily would respond not by saying, "You have failed to express a proposition," but "He doesn't believe that"—and we have the strong intuition that this denial would be true.

Compare the following case: an astronaut on the moon¹⁸ says, "It's three o'clock." Typically, this sentence would be used to express the claim that it is three o'clock in Z, where Z is the time zone in which the utterance takes place. The confused astronaut thinks that there are time zones on the moon, and he intends to claim that it is three o'clock in "Z," which is the time zone he is in. But there is no such time zone. So he fails to express a proposition. We feel no qualms, however, about denying his claim: "It's not three o'clock. There are no time zones on the moon, you know."

The present difficulties are often discussed in connection with "negative existential" claims. But the same issues arise with respect to all sorts of denials in which the speaker believes there to be reference failure. A child who sincerely asserts, "Santa will come tonight," fails to refer, and therefore, on most direct reference accounts, fails to express a proposition. But the parent who responds, "Santa will not come tonight," explains that there is no Santa, makes what seems to be a true claim, despite the fact that the use of 'Santa' does not refer.

Note that these examples would present no trouble for description theories of reference. For if in these cases the original speakers are seen not as attempting to provide a specific thing to be a propositional co-

is for a form of speech to have a syntactic expression *stand for* a propositional constituent.¹⁷

Consider our practices of reporting velocity. A claim that an object is moving at a certain velocity makes sense only if it is understood with respect to what the velocity is to be assessed. We say that velocity is relative to an observer, or a frame of reference—we must count something as stationary. But we articulate this additional parameter of velocity claims only when it is not obvious what is to count as stationary. We have in English a number of general-purpose constructions for articulating commonly suppressed constituents of a claim. We say, 'with respect to ...' or 'relative to ...' or 'in the sense that ...'. The more likely the unarticulated constituent is to be unclear, the more likely it is that we have a natural way to articulate it.

In the case of belief reports, in which notions are unarticulated, we do have rough and ready ways to clarify just which notions we mean to talk about. We say, for instance, that Miles believes that Edward is a peasant in one way—the way related to the boy in front of him, not in the way related to the Prince. Or we add to the report, "that is, he thinks the boy in front of him, who really is Edward, is a peasant." Or we specify how Miles would or would not "put" his belief. Or we allude to the evidence which led Miles to form the belief, or to the actions it would be likely to bring about. Each of these devices can succeed in distinguishing among the two notions which in context can seem equally relevant, thus eliminating possible confusion about which notion we mean to talk about.

We do not, of course, have a very direct way of specifying the notions we mean to talk about in belief reports. This is due to the fact that it almost always obvious which notion a speaker is talking about. Where it is not, we either use one of the devices just mentioned, or leave the language of belief reporting altogether and talk instead about what the agent would say or would do.

Unarticulated constituency and direct reference are of a single stripe. In fact, if we take the term 'reference' in the ordinary sense in which it does not require a referring expression, unarticulated constituency can be seen to result from a kind of direct reference—perhaps "tacit" reference. When a speaker claims that "it's raining," she is referring to a place and not to a

stituent, but merely as claiming that there is a thing meeting a certain condition (being the generous elf known as 'Santa', being the local time zone, or being the man's perceptual notion of the building), then the claims are straightforwardly false and the denials are true.

The descriptorial theories have even more than this kind of extensional correctness going for them; it is because the cited conditions—call them *providing conditions*—are not satisfied that the denials are true. In the child's use of 'Santa', the providing condition, of being the generous elf known as 'Santa', plays a central semantic role, even though it is not the referent of the child's use of the name. It is a condition that the child expects to be filled as a precondition of successful reference. He expects to refer successfully to a thing in virtue of its meeting the providing condition. His supposed ability to refer to a thing by using the name 'Santa' depends on the condition's being satisfied. Similarly, the astronaut takes it that he can talk directly about a time zone, that he can provide one, because it meets the providing condition of being the local time zone. And, we claim, the belief reporter expects to be able to talk directly about a notion because it satisfies the condition of being the man's perceptual notion of the building.¹⁹

A normal, successful case of direct reference involves a speaker referring to an object in virtue of that object's satisfying a providing condition. Reference failure involves failure of a presupposition, namely, the presupposition that a providing condition is satisfied. Now, expressions like proper names and underarticulated phrases that normally invoke devices of direct reference are sometimes used where there is no presupposition that the relevant providing conditions are satisfied. The denials in the cases of the blind man, the astronaut, and Santa are like this. In each of these denials, the speaker does not presuppose that there is a thing meeting the providing condition that is invoked by the utterance. Instead, we claim, the speaker *raises the providing condition to constituency*—he talks about the condition itself rather than about a supposed thing that meets it. The providing condition now plays a semantic role—as a constituent of the proposition expressed in the denial—more central than its usual auxiliary role of providing a propositional constituent.

In particular, the claim expressed by 'Santa will not come tonight' (in the described circumstances)²⁰ is to the effect that there is no generous elf known as "Santa" who will come tonight. And the proposition expressed by 'It is not three o'clock' is that there is no local time zone such that it is three o'clock there. And the content of 'He does not believe that that

building is far away' is the claim that there is no perceptual notion of a building such that the man has a belief involving that notion, with the content that the building is far away. The denials are thus true, and the truth is consistent with our claim that the assertions they deny strictly speaking fail to make claims.

Of course, for each of the original, claimless assertions, there is a proposition closely related to the kind of proposition the speaker intends to express, which we can for most purposes charitably treat as the content of the statement. Specifically, we can take the speaker to have expressed the claim that there is a thing meeting the invoked providing condition such that so-and-so. In fact, the speaker of such an assertion is pre-assumptively committed to this proposition, in virtue of his commitment to the presuppositions that must be satisfied if he is to make a successful claim the way he intends.

Above we analyzed our reader's utterance of (2) in an imaginary case which Miles has been shouting about the Prince. In fact, Miles was not shouting, "Edward Tudor is of royal blood," at the time he encountered the boy. The reader actually has no specific actions on Miles's part which he can tie such a notion of Prince Edward. It is obvious from the general tenor of the novel, however, that Miles would have such a notion. Every full-witted adult in England at the time has a notion of Prince Edward—one they acquired shortly after he was born which motivates their behavior in regard to the Prince of Wales, such as their use of the phrase 'Prince Edward', their decorum when the royal procession goes by and the like. Our reader may not be able to pick out anything very specific in Miles's behavior to serve as evidence that he has such a normal notion of the Prince. But he has every right to suppose that he has one.²¹

It may seem implausible to suppose that our reader, in using (2), can directly provide a notion for the report to be about, since the reader is not directly acquainted with such a notion. If this intuition is right—our assumption we shall question in a minute—our machinery gives us a natural way of respecting it: this is a case in which, instead of a notion, providing condition becomes a propositional constituent. What our reader is claiming with (2) is that there is some normal notion via which Miles believes that Edward is royal; that is, the condition of being a normal notion of the Prince is the unarticulated constituent. The report, in this construal, is an example of a second kind of belief report—in which notions are not provided, but instead are constrained by provided con-

ditions; the report is about those conditions in the sense of “about” appropriate to propositional constituents.

For this (supposed) second kind of belief report we can give the following account. Where μ is a belief report at τ which is about conditions C_1, \dots, C_k , and $p = \text{Content}(\mu_S)$, where μ_S is the subutterance of μ corresponding to the object sentence S ,

$$\text{Com}(\mu) = \exists b [B(a, b, t) \wedge \text{Content}(b, t) = p \\ \wedge \exists n_1, \dots, n_k \bigwedge_{r_i \text{ in } p} (C_i(n_i) \wedge \text{Responsible}(n_i, r_i, b))]$$

So we have room in our framework for two sorts of belief report, corresponding to whether notions are themselves *provided* or merely *constrained* by conditions. Supposing for now that there really are two kinds of belief report, how can we know, for a given report, of which kind it is? One way, surely, is to look at what would happen if the appropriate notions were to fail to exist. If the report would then be false, then it is a case of notion constraint rather than provision; if the report would fail to make a claim, then it is a case of (attempted) notion provision.

Of course, we have seen how, in a case where an attempt to provide a notion fails, a proposition closely related to what the speaker is trying to express takes center stage. This is the false proposition to the effect that the agent has a notion which meets the invoked providing condition and which is involved in a belief with such-and-such content. Given this fact, our intuitions about whether a belief report fails to make a claim or is simply false are in the same boat as our intuitions about the truth value of the child’s claim that Santa is coming. The falsity of the closely related propositions, plus the truth of the natural denials of these statements, may well obscure intuitions about the truth of the original claims.

In this paper, we adopt officially the position that there really are belief reports of the second kind (which are about conditions rather than notions). Given our points about providing conditions and propositions to which speakers are pre-assertively committed in cases of direct reference, however, a plausible case can be mounted for the view that, in all successful belief reports, specific notions are provided for the report to be about.²²

Assuming, now, that there are two classes of belief reports, there is no reason to suspect that all reports will fall clearly into one camp or the other. For example, if our reader simply assumes that Miles must have a normal notion of King Henry and expects his audience to do the same, then it makes little difference whether he claims that Miles has a belief

involving that notion (notion provision) or just a belief involving a normal notion (notion constraint). Since it makes little difference, our reader need not go to any pains to indicate which of the claims he is making; report simply can land between the two claims.²³

V

With that out of the way, let us turn to our examples.

First, a recap of the semantics of (1) and (2). We shall treat (1) as a case of notion provision. The provided notion is Miles’s notion of Edward which is connected with his perception of and actions toward Edward the mob incident. The reader claims that Miles does not have a belief involving that notion, with the content that Edward is of royal blood.

With (2), the reader provides a condition on notions, the condition being a normal notion of Prince Edward. The reader claims that Miles has a belief involving some normal notion of Edward, with the content that Edward is of royal blood.

In the Pierre case, the sentence (3) gets used in two reports, first in discussion of Pierre’s initial acquaintance with London through stories then in a discussion about Pierre’s thoughts of his adopted home. Call these reports μ_3 and μ'_3 . Pierre actually has two notions of London, or relevant to each discussion; call the first n and the second n' . The notion meets the condition C of being a notion germane to the discussion (Pierre’s reaction to the stories; the notion n' meets the condition C' of being a notion germane to the discussion of Pierre’s new home.

If one of the two analyses is uniquely correct for μ_3 and μ'_3 , it is perhaps the account in terms of notion constraint. The speaker of the former report is claiming that Pierre has a belief involving some notion germane to the current conversation about the stories, with the content that London is pretty. The speaker of the latter report requires that the belief involve some notion relevant to the conversation about Pierre’s new home.

If the circumstances of μ_3 and μ'_3 are such as to make the notions n and n' clear and present to the speakers and their audiences, then the analysis should be in terms of notion provision. If this is the case, then the speaker of μ_3 claims that Pierre has a belief involving the notion n with the content that London is pretty; the speaker of μ'_3 claims that Pierre has a belief involving the notion n' with the content that London is pretty. If the circumstances of the two reports are less clear-cut, as they often are, then, a

noted earlier, the claims made by the speakers might fall between those offered by the notion-provision and notion-constraint accounts. There just might be no saying.

Note, though, that any of these analyses constitutes a solution to the puzzle. The claim made in κ_3 is simply true, and the claim made in κ_5 is simply false.

Kripke presents the puzzle as arising from a few very plausible principles about belief reports, including:

Disquotation: If a normal English speaker, on reflection, sincerely assents to ' p ', then he believes that p (*op. cit.*, pp. 248–9 [885–6 in this vol.]).

Translation: If a sentence of one language expresses a truth in that language, then any translation of it into any other language also expresses a truth (*op. cit.*, p. 250 [886–7]).

On our account of belief reporting, neither of these principles is at all plausible in general. Each principle presupposes that it is belief sentences that are true or false. On our view, a single sentence, like (3), can be used in both true and false reports. Kripke assumes that, because of the lack of obviously context-sensitive words, (3) can be considered more or less "eternal." But words are not the only sources of context-sensitivity; the presence of unarticulated constituents also can widen the gap between a sentence and the proposition expressed by a statement of it. And that is what happens in the Pierre case.

Richard lists three sentences considered as uttered by A watching B in the phone booth:

- (4) I believe she is in danger.
- (5) I believe you are in danger.
- (6) The man watching you believes you are in danger.

A uses (4), clearly, to make a true report. His notion n_{vis} of B which stems from his view out the window, which is associated with his idea of being in peril, and which causes his waving is supplied. It is claimed that A has a belief involving n_{vis} with the content that B is in danger. He in fact has such a belief.

The man would not sincerely use (5) over the phone; if sincere, he certainly would deny (5). The natural intuition, we think, is that a use of (5) in the described circumstances would make a false claim. (It is this reaction which Richard sets out to prove mistaken. The very possibility of our semantics shows that his proof is in error.)

The Prince and the Phone Booth

The set-up for (6) is as follows. B sees a man, A , in a building across street waving frantically. Amused, she says (over the phone), "the man watching me believes that I'm in danger." Echoing her, A utters (Surely, B 's claim is true. And if so, A 's use of (6), which is in explicit agreement with her, is true also.

So we hold that the use of (5) is false while that of (6) is true. But how can this be? The two reports are uttered by the same person in the same circumstances, they ascribe beliefs to the same agent, and they use precisely the same embedded sentence, understood in the same way! The only difference is the way in which the man is referred to—in the one case with 'I', in the other with 'the man watching you'.

Difference enough, we think. The pragmatic principle of self-ascription applies to (5) but not to (6):

Self-Ascription: An utterance of 'I believe that . . . τ . . . ' provides (or, is about) the notion that is connected to the speaker's use of ' τ '.

Using 'I' in (5), A thus directs attention to the notion (n_{phone}) that linked to his use, in (5), of 'you'—the notion of B which is associated with his idea of being the one he is addressing²⁴ and not associated with his idea of being in danger. So A 's use of (5) makes the claim that he has belief involving n_{phone} , which has the content that B is in danger. He in fact has no such belief.

In (6), A is discussing those beliefs of the man watching B , that is, of himself, which explain the frantic gestures directed at B . So he claims the man has a belief involving n_{vis} , the notion linked to his perception of the window and his gestures of warning, which has the content that B is in danger. In fact, A has such a belief.

Richard's case is especially interesting because it shows how a contextual shift can be brought about by a change in wording outside of the embedded sentence in a belief report. This gives added force to our analysis of substitution worries: the wording changes in the usual cases of reluctance to substitute are responsible, not for changes in meaning or explicitly specified content, but for changes in what is provided by context for the reports to be about.

Our semantics allows that, for a given belief sentence, absolutely any of the agent's notions may be provided—there is no semantic restriction or what notions may be provided in a use of a given sentence. But there are many pragmatic principles, like self-ascription, that constrain which notions can be provided in the normal case. It is semantically but not

Mark Crimmins and John Perry

pragmatically possible for a use of 'I believe I am not me' or (normally) 'S, but I do not believe that S' to be true. Although it is semantically possible, in W. V. O. Quine's example, for an utterance of 'Tom believes that Cicero is Tully' to express a true proposition (say, if Tom's 'Cicero' notion is provided twice over), there may be no very natural use of that sentence which in fact expresses the proposition (although surely we can concoct a Richard-ish example to put this point in doubt). In the normal case, the use of different names for Cicero serves as a strong, though perhaps defeasible, indication that the names have some importance to what's being said over and above just standing for Cicero. Such a difference in names requires a sufficient reason—in this case, a difference in which notions are being provided to play the corresponding roles in the ascribed belief.

The relation of the present proposal to Fregean semantics for belief reports should be relatively clear. The broad similarity consists in the agreement that a belief report specifies, in addition to simply which objects the agent is claimed to have a belief about, also just how the agent is cognitively connected to those objects. On our account, the report specifies (or constrains) the particular notions allegedly involved in the belief. On a Fregean account, "senses" are specified.

Two crucial differences separate the accounts. First, we stress the particularity and unsharability of notions. Since notions are full-fledged particulars immersed in the causal order, they have a great array of different features that we can exploit to provide them in our belief reports. They are involved in beliefs, associated (sometimes) with words, formed in specific circumstances, connected to preceptual situations, reasonings, and actions; they survive the formation and abandonment of beliefs in which they are involved; and so on. We can use each of these kinds of fact to give us a handle on a notion, a way of picking it out. This frees us from a problem often noted about the Fregean strategy: it appears that, on most natural construals of what senses are, we often do not know just what sense an agent attaches to an object (we do not grasp it), and so we cannot know just what we are attributing to the agent with a belief report, which, after all, must be about the agent's ways of thinking.

As we have said, there is nothing in our view incompatible with something like Fregean senses, considered as entities which we can use

to classify an agent's notions. A Fregean might well take our talk "notions" as an account of what it takes for an agent to grasp a sense—agents grasp senses in virtue of having appropriate notions.

The second departure from a Fregean account is in our claim that an agent's way of thinking about things (her notions), though they are specified in a belief report, are not the referents of the words occurring in the embedded sentence. This difference becomes especially important in the analysis of certain kinds of reports: those with content sentences containing devices of underarticulation, and those with no content sentence at all, but which instead are completed with the likes of 'what you said 'the same thing', and 'Church's Thesis'.

The account of belief reports sketched here closes some doors. If, as we claim, a single belief sentence can be used in both true and false report then there can be no simple logic of such sentences. The simplest possible rule,

A believes that S

A believes that S

does not hold in general, as we learn from Kripke's puzzle.

Even a logic of belief sentences restricted to a single context will prove difficult.²⁵ Although a relativized version of the above rule will certainly hold, this one,

A believes that S (relative to c)

A = B (relative to c)

B believes that S (relative to c)

will not, as we learn from Richard's puzzle.

Also closed is the prospect of a strictly compositional semantics for belief sentences. The semantic values of the subexpressions in a belief report, on our analysis, do not provide all the materials for the semantic value of the report itself. Notions and conditions on notions are not articulated, but end up in the contents of reports; so the semantics for belief reports is in an important way noncompositional.

In addition, our account denies what some have seen as a primary desideratum for theories of belief: that a belief report claims simply that a binary relation holds between an agent and an object of belief.

And, perhaps worst of all, we have given an account on which it appears to be next to impossible to give a complete, systematic account of which claims are made by which belief reports. We have claimed that belief reports are context-sensitive, that they invoke unarticulated constituents, without offering any general method for determining what the relevant contextual factors are, and how they give rise to these unarticulated constituents of belief reports.

Tempted as we are to view each of the above results as an insight rather than a drawback, we realize that we have abandoned many of the issues and goals commonly pursued in this area. But we think the account opens many doors as well.

Whereas there is little possibility of an interesting logic of belief sentences, the logic of beliefs, notions, and ideas is available. Such issues as logical and analytic closure of belief, explicit versus implicit belief, and inferential issues in belief change really belong to the logic of beliefs rather than to the logic of belief sentences. We can explore the logic of the relations we have seen as underlying our ordinary talk about beliefs—but this logic will not be a logic of ordinary language.

Of course, we have explained very little about what beliefs, notions, and ideas are. But we think our partial account of them raises obvious questions in theories of representation, action, perception, and the metaphysics of mind.

Our semantics is not compositional, but there is system in the non-compositional mayhem. The ways in which notions and conditions on notions are provided have yet to be explored to any great extent. But the discussions of the belief puzzles suggest several directions from which to look at these mechanisms.

Last, the move to unarticulated constituents emphasizes the importance of pragmatic facts about language to the study of what seem like purely semantic issues. In order to express claims, we exploit a tremendous variety of facts, conventions, and circumstances, of which the meanings and referents of our terms form just a part. So it is a mistake to relegate pragmatics to matters of felicity and implicature. In the case of belief reports, it is central to understanding content and truth.

Notes

This work was supported in part by the System Development Foundation through a grant to the Center for the Study of Language and Information. We would like to thank the Philosophy of Situation Theory group at CSLI; special thanks to David Israel.

1. *Situations and Attitudes* (Cambridge: MIT, 1983), pp. 253–264.
2. “Substitutivity,” in *Essays in Honor of Richard Cartwright* (Cambridge: 1987), pp. 99–132, and “Direct Reference and Propositional Attitudes,” Almog, Perry, Wettstein, eds., *Themes from Kaplan* (New York: Oxford, 1986).
3. *Frege’s Puzzle* (Cambridge: MIT, 1986).
4. *Opacity* is the claim that substitution of corefering names and demonstrata in belief reports does not necessarily preserve the truth of those reports (descriptions are another matter; it is not nearly as controversial that substitut descriptions for a corefering name can influence the truth value of a belief report). What “substitution” comes to with respect to utterances (belief reports), as opposed to sentences (belief sentences), is not at all obvious. Our simple notions of substitutivity, opacity, and so on are really useful only if sentences are (wrongly) taken as the bearers of truth and content. Here, we shall adopt an informal notion of substitution in belief reports, such that the reports (1) and (1’), as well as (2) and (2’) are related by substitution.
5. For an important qualification, see n. 14 below.
6. “A Puzzle about Belief,” in A. Margalit, ed., *Meaning and Use* (Dordrecht: Reidel, 1979), pp. 239–283 [chap. 38 in this vol.].
7. For a fuller defense of the particularity of beliefs, see Crimmins, *Talk and Beliefs* (Cambridge: MIT, 1992).
8. A speaker can adopt a name (like ‘John’) more than once, to refer to what are different individuals. Each such adoption creates a type of use to which a speaker may put the name. So a causal analysis of names should look no names themselves, but at types of uses of names, as the things for which refer is determined causally. An agent may use ‘John’ to refer either to John Dupre or to John Etchemendy. What individuates these distinct types of uses of the name ‘John’? One answer is that the types of uses of ‘John’ are tied to distinct notions of the agent.
9. See “Demonstratives,” 1977 manuscript reprinted in *Themes from Kaplan*.
10. There is no mystery as to how a single thing can be a part of many different things at the same time (and at different times). One may, for example, be a member of many different committees or clubs.
11. This is to consider only beliefs of a certain kind of composition. In a more thorough presentation, a discussion of other kinds of belief structures, perhaps including general beliefs and complex beliefs, might be called for—although logical connectives and quantifiers can be accommodated within this simple structure. Also, we have chosen to ignore in this paper many subtleties of time reference.
12. Roughly, an argument role of a relation is also an argument role of a proposition (at least) when, in that proposition, the role of the relation is occupied by an object.
13. “Direct Reference and Ascriptions of Belief,” *Journal of Philosophical Logic* xii (1983): 425–452.

14. In accord with our simple version of "semantic innocence," we assume throughout that a belief report specifies the content of the ascribed belief by providing a sentence with the same content, as uttered in the report. The puzzle cases that we consider seem to be ones for which this assumption is correct. There are good reasons, however, to think that things do not always work this way. One way of analyzing, "Barbara believes that the Twin Towers are over a foot tall," would involve quantification over contents of beliefs. Other cases of reporting implicit and tacit beliefs might well work similarly. Another case in which a proposition might be "quantified out" is in the use of, "He believes that Russell's Fairy will make him rich," knowing full well that the embedded sentence does not express any proposition (if in fact it does not). These and other cases make us wary of insisting that a content proposition is always specified in a belief report. The present strategy can be extended in relatively simple ways to account for such cases.

15. Yet another simplification: we ignore the fact that many uses of singular terms, including terms in the subject position of belief sentences, are not directly referential. "Attributive" uses of definite descriptions really should be handled differently. Note also that we really should treat the idea in a belief in the same way we treat notions here: though the puzzles considered here do not turn on this, others certainly do.

16. Here one major difference from the "official" belief-report semantics in ch. 10 of Perry's *Situations and Attitudes* (p. 256) is apparent. There, a belief report is true if the agent has any belief with the specified content. There is a further crucial difference that is not so obvious. Barwise and Perry countenance beliefs as real, concrete things, as we do here. But these beliefs are represented as situations of an agent being related to an anchored belief schema. Belief schemas are abstract objects in which what we have called notions are represented by indeterminates. Although the way this all works is quite complicated, in the end beliefs are individuated by belief schemas—abstract objects—and the things in the world to which the indeterminates in the schemas are anchored. But indeterminates are not notions, and, we think, relations to anchored belief schemas are not quite fine-grained enough to individuate beliefs in the ways needed for belief reports. So we suggest two major changes to the account in *Situations and Attitudes*: we give ourselves the theoretical machinery to talk about notions and ideas directly; we then claim that these things are among the subject matter of belief reports (via the mechanism of unarticulated constituents), and are not merely quantified over.

17. For more on unarticulated constituents, see Perry, "Thought without Representation," *Proceedings of the Aristotelian Society*, suppl. vol. LX (1986): 263–283. There, a systematic semantics for some underarticulated constructions is given, which is connected to a recursive model of syntax in the usual way.

18. John Etchemendy brought up this version of Wittgenstein's example.

19. Just which providing conditions are invoked in a given case depends on a wide range of circumstances. Also, there usually is more than one such condition for

a given use of a term. Providing conditions for a use u of 'here' by speaker A at location l , for instance, include the conditions of being the location of the utterance of u , being where A is, and being l . In the 'Santa' case, we have the conditions: being the referent of the utterance of 'Santa', being the relevant thing known as 'Santa', being the generous elf known as 'Santa', and so on.

20. In circumstances where it is presupposed that the providing condition is in the denial expresses just the negation of the proposition (if there is one) expressed by the corresponding assertion.

21. What counts as being a normal notion certainly depends, we think, not only on what is common in a community, but also on other aspects of the background of the discourse, including facts about what is relevant to the goals of the course. We would expect an account of "being a normal notion" to exhibit many of the same features as an appropriate account of "knowing who b is," which certainly is background-sensitive in many ways. See, for example, Boer and Lycan's *Knowing Who* (Cambridge: MIT, 1985).

22. This view is argued in Crimmins, *Talk about Beliefs*, though in the end argument rests precariously on the fact that none of the examples considered are natural candidates for reports of the "second kind" seems clearly to be as required. There is another way, also, in which a report can land between the notion-provision and the notion-constraint types of report. It is not hard to concoct one in which one notion is provided and another is constrained: a natural construal of our reader's report "Miles believed that he (Edward in rags) was less noble than Prince Edward" might go along these lines. So in the general case, both notion-provision and conditions may be provided; there are no difficulties in formalizing this also in the lines already given for the pure notion-provision and pure notion-constraint analyses.

24. More precisely, with his idea I_{agent} , the idea that has the context sensitive semantic role picking out the person being addressed.

25. Here we mean 'context' in a sense such that various different statements could be made in the same context. One way of taking our claims in this paper would be as denying the general usefulness in semantics of such a restricted notion of context. Taken this way, we have claimed that such things as the words used in a statement can affect the semantically relevant parts of the statement's context.