

Forms and objects of thought

Michael W. Pelczar

Received: 6 February 2006 / Accepted: 16 November 2006 / Published online: 27 February 2007
© Springer Science+Business Media B.V. 2007

Abstract It is generally assumed that if it is possible to believe that *p* without believing that *q*, then there is some difference between the object of the thought that *p* and the object of the thought that *q*. This assumption is challenged in the present paper, opening the way to an account of epistemic opacity that improves on existing accounts, not least because it casts doubt on various arguments that attempt to derive startling ontological conclusions from seemingly innocent epistemic premises.

Keyword Propositional attitudes

Propositional attitudes and the language we use to ascribe them are a perennial source of interest to both linguists and philosophers. Linguistically, the practice of ascribing propositional attitudes defies regimentation of the kind to which it seems possible to subject much of the remainder of human discourse. In particular, our use of such indispensable verbs as “believes” and “wants” appears to violate the otherwise fruitful and well-confirmed principle that the semantic value of a sentence is a function of the semantic values of its component expressions together with their mode of composition. Philosophically, propositional attitudes are of interest as the prime candidates for underived intentionality. Here we find the roots of the compositionality problem in the apparent possibility for a thinker to have one belief to the effect that the world is a certain way without having all the beliefs he could have to the effect that it is that way.

In this paper I survey existing attempts to solve these problems, and propose an alternative solution. What makes this solution an alternative is not that it comports with the available linguistic or psychological data better than existing theories; indeed, the theory I favor—I call it *the multiple relation theory*—seems to be

M. W. Pelczar (✉)
Department of Philosophy,
National University of Singapore, 3 Arts Link,
Singapore 117570, Singapore
e-mail: phimwp@nus.edu.sg

observationally equivalent to at least some of the theories that have been advanced in the past. What is distinctive about the multiple relation theory is rather the theoretic efficiency with which it accounts for the data. According to the multiple relation theory, propositional attitudes do not pose any challenge to formulating a compositional semantics for natural language over and above the challenge already posed by the use in natural language of indexical expressions. So, even though the multiple relation theory does not guarantee compositionality of the strongest sort—the sort that requires the truth conditions of a declarative sentence to be a function of the semantic values of its type-individuated components and their manner of composition—it does have the happy result of reducing the compositionality-resistance of propositional attitude ascriptions to the comparatively well-understood resistance to compositional analysis of sentences containing indexicals.

This semantic simplification has a parallel at the level of mental representation. For, the semantic simplification is effected by treating attitudinal verbs as corresponding each to a family of extensional relations between thinkers and states of affairs, where the latter may be individuated as sets of logically possible worlds (or even more coarsely than that). From this standpoint, intentionality is not so much a feature of intentional states as it is a feature of our practice of ascribing them, where this practice itself may be characterized in purely extensional terms. Those who regard intentionality as the mark of the mental need not therefore regard intentionality as such.

1 The Possibility

Many people believe that Mark Twain wrote *Huckleberry Finn* without realizing that Samuel Clemens wrote it, despite the fact that Twain and Clemens were the same man. There is water in the world if and only if there is H_2O in the world, but this did not prevent Cavendish, who already knew there was water, from discovering that there was H_2O . The solution of a mathematical equation might elude the best minds for centuries, even though its eventual statement is logically equivalent to the triviality that a thing is the same as itself. In general, the fact that p iff q does not prevent a person who believes, knows, or otherwise bears in mind that p from failing to believe, know, or likewise bear in mind that q . I call the possibility for this kind of situation to arise “the Possibility.”

There is continuing philosophical debate over how to explain the Possibility, as well as over what the Possibility tells us about the nature of mental and linguistic content. Partly from an instinct to sidestep such issues, some philosophers deny the Possibility altogether, in spite of the everyday evidence in its favor. According to them, the statement that

- (1) Locke knew there was water.

is logically equivalent to the statement that

- (2) Locke knew there was H_2O .

These philosophers maintain that the only difference between the two knowledge attributions is pragmatic rather than literal. Thus, while we are not (normally)

warranted to assert that Locke knew there was H₂O, the statement that he did know this is literally true. What makes it seem untrue is that the norms governing ordinary conversation cause the statement to *suggest* that Locke was familiar with the results of modern chemistry, without literally saying or entailing this. This supposedly deceives us into mistaking a familiarity with modern chemistry for part of the literal import of the statement that Locke knew there was H₂O.¹

There are two problems with this view. First, it relies on an implausible account of conversational norms. It is doubtful that a norm such as, “When making a knowledge attribution, use language that the individual to whom you are making the attribution would himself use to express the knowledge attributed to him” has the reach or strength to create the standing appearance (on this view illusory) that it is literally false that Locke knew there was H₂O.²

Second, even if conversation were governed by such norms, the whole question of what (1) and (2) literally mean is secondary to what *we* mean by them. At bottom, the Possibility is that *what we mean* when we say that Locke knew there was water is true, despite the fact that *what we mean* when we say that Locke knew there was H₂O is false. Perhaps we overestimate how well our meaning corresponds to that of our words, but this does not create any doubt among ourselves about what we mean by our words. What we mean is normally quite clear. In particular, it is clear that we all normally mean something true by (1) and false by (2), true by “Locke knew that water was water” and false by “Locke knew that water was H₂O,” etc. This is so even if we mean the same thing by “water” and “H₂O.”³

So much for those who deny the Possibility. Those who grant it divide into two camps over the question of what it entails. On one hand, there are those who say that if it is possible for a person to believe that p while not believing that q, then there must be a difference between what he believes in believing that p and what he fails to believe in failing to believe that q. On this view, the belief that Twain wrote *Huckleberry Finn* has a different object from the belief that Clemens wrote it; the knowledge that there is water has a different object from the knowledge that there is H₂O; the act of wondering whether $1 + \frac{1}{2} + \frac{1}{4} + \frac{1}{8} + \dots = 2$ has a different object from the act of wondering whether $2 = 2$; and, in general, for any form of thought, the possibility of having a thought of that form to the effect that p without having a

¹ This is the “Russellian” approach favored by McKay (1981), Salmon (1986), Soames (1987), and, with less emphasis on the role of implicature Braun (1998).

² See Recanati (1993, pp. 328–347) and especially Green (1996).

³ Recent philosophy shows a tendency to subordinate speaker meaning to word meaning, as witness the ongoing debate over how to distinguish what is literally said from what is otherwise conveyed or put forth. This debate is of little philosophical interest except on the doubtful assumption that it matters how far what a philosopher means and is understood to mean coincides with what his words mean, when he says something like “Locke did not know there was H₂O” or “I know there is an external world.” If the astronomical community agrees that radio observations warrant the assertion that there is a galaxy with a red-shift of nine, it will simply dismiss as frivolous objections to the effect that the *words* “there is a galaxy with a red-shift of nine” are somehow literally false. What the words themselves mean is a question for lexicographers, not astronomers. I fail to see why philosophy should differ from astronomy in this respect.

corresponding thought to the effect that q entails a difference between the objects of the thoughts that p and that q . This view is prevalent enough to be called *the standard view*.⁴

On the other hand, there is the neglected view that the possibility of thinking that p without thinking that q entails no difference between what you think in thinking that p and what you fail to think in failing to think that q . On this view, someone might know that there is water but not that there is H_2O even if there is no difference in content between the knowledge that there is water and the knowledge that there is H_2O . This is the multiple relation theory. The goal of this paper is to show that the multiple relation theory is at least as plausible as, and in some respects preferable to, the standard view. In the next section, I consider the main argument for, and the main versions of, the standard view. I then lay out my alternative. Finally, I argue that this alternative is superior to the standard view in several important respects.

2 The standard view

According to the standard view, if it is possible to believe that p without believing that q , then what one believes in believing that p differs from what one fails to believe in failing to believe that q . As Frank Jackson reasons,⁵

What we say about what the world is like using the sentence “There is water” is different from what we say about what the world is like using the sentence “There is H_2O .” Otherwise it would not have been a *discovery* that water is H_2O .

Generalizing this reasoning yields the basic argument for the standard view (replacing “belief” with “knowledge,” “desire,” “saying,” etc. will not affect the argument):

1. Belief is a relation.
2. If belief is a relation, then if it is possible to believe that p without believing that q , what one believes in believing that p is different from what one fails to believe in failing to believe that q .
3. One can believe that there is water without believing that there is H_2O .
4. Therefore, what one believes in believing that there is water differs from what one fails to believe if one fails to believe that there is H_2O .

Different versions of the standard view differ in what they take belief to have for its objects, such that the belief that there is water has a different object from the belief that there is H_2O (and likewise for other forms of thought). There are three main proposals on this head: the traditional theory, the structured proposition approach, and two-dimensional semantics.

⁴ Throughout, I use “content” and “object” interchangeably to mean that which is thought (believed, known, desired) in thinking (believing, knowing, desiring) that so-and-so. The words “thought,” “belief,” etc. are notoriously ambiguous, sometimes indicating an act or state (that of thinking or believing something), and sometimes indicating that which constitutes the truth-evaluable substance of such an act or state; this is the familiar “ing/ed” ambiguity. In the present essay I use these words only with their former, “ing” senses, and use “content” and “object” when I need to speak of what is cognized.

⁵ Jackson (2004, p. 262). See also Russell (1905, pp. 487–488) and Quine (1953a, p. 9).

2.1 The traditional theory

The traditional theory is that we can account for the Possibility in all its forms by individuating the objects of thought as sets of logically possible worlds, so that the object of the belief that p , for example, is the set of worlds in which it is the case that p . Call the totality of possible worlds in which there is water “the proposition that there is water,” and the totality of possible worlds in which there is H_2O “the proposition that there is H_2O .”⁶ What the traditional theory says is that the proposition that there is water is different from the proposition that there is H_2O . This way we can explain the Possibility by just letting propositions serve as the objects of thought. In believing that there is water, what you believe is a proposition that is true at any possible world that contains stuff with all the outward qualities of water (clarity, solvency, potability, etc.); in believing there is H_2O , what you believe is a proposition that is true at any possible world containing stuff with this chemical composition. Since there are possible worlds that contain something with all the outward qualities of water but nothing with this chemical composition, these are different propositions, whence the possibility of believing that there is water without believing that there is H_2O .⁷

The traditional theory explains the possibility of believing that Twain wrote *Huckleberry Finn* without believing that Clemens wrote it similarly, by maintaining that the proposition that Twain wrote it differs from the proposition that Clemens wrote it; maybe these propositions diverge in truth value at possible worlds in which Samuel Clemens never takes up writing, and another person very much like the actual Twain writes *Huckleberry Finn* using the pseudonym “Mark Twain.” Discovering that water is H_2O or that Twain is Clemens, according to the traditional theory, is a matter of learning that the relevant propositions have the same truth conditions at the actual world; for instance, that it is actually true that there is water iff there is H_2O , or, that it is actually true that Twain wrote *Huckleberry Finn* iff Clemens wrote it. In order to explain the possibility of mathematical discovery, the traditional theory must construe mathematical beliefs as having for their objects something besides the propositions expressed by mathematical statements, since there are only two of these (the proposition that is true at every possible world, and the one that is false at every possible world). One way to do this is to treat the objects of mathematical belief as contingently true propositions about mathematical expressions, such as the proposition that a given pair of mathematical sentences are logically equivalent.⁸

The traditional theory does not have as many proponents as it once did, because it now seems to many philosophers that the proposition that there is water is the *same* as the proposition that there is H_2O , and likewise for Twain and Clemens and other cases illustrating the Possibility. For instance, it seems that it is not just true, but *necessarily* true, that there is water iff there is H_2O . This is because it seems to be a necessary truth that water is H_2O . Everything is, after all, necessarily the same as

⁶ I use the term “proposition” purely as a term of art. Some authors use it to designate *whatever* serves as an object of thought, which they may or may not identify with what I am here calling a proposition. Propositions in my sense are equivalent to Stalnaker’s propositions (1987, p. 2) and Jackson’s C-intensions (1998a, pp. 46–52).

⁷ The clearest versions of the traditional theory are Russell (1905) and Wittgenstein (1961/1921, nos. 5.541–5.542).

⁸ See, for example, Stalnaker (1987, pp. 72–77).

itself, and nothing is the same as anything besides itself. Thus, wherever the relation of identity obtains, it obtains necessarily. So, given that the relation of identity obtains between water and H_2O , any possible world in which water exists is a possible world in which H_2O exists, and vice versa. Hence, by the definition of “proposition,” the proposition that there is water is the same as the proposition that there is H_2O .

Admittedly, not all identity statements state necessary truths. But that is only because not all identity statements are simple statements of identity. It is only contingently true that Ottawa is the capital of Canada, because to say that Ottawa is the capital of Canada is to say that Ottawa uniquely satisfies a certain condition—that of being a capital of Canada—which it satisfies only contingently. We can easily imagine a situation in which Ottawa is not the capital. But how are we to imagine a situation in which water is not H_2O ? Here it will not do to entertain a world without H_2O that contains a substance in all outward respects just like water. Such a world contains something like water, and perhaps something that is even *called* “water,” but it does not contain water. It is as much in the nature of water to be H_2O as it is in the nature of cats to be mammals, or in my nature to be me.⁹

At least, this is the prevailing opinion. Perhaps it is mistaken; perhaps people do normally use “water” as shorthand for a description that water uniquely but contingently satisfies, as Ottawa uniquely but contingently satisfies the description “capital of Canada.” Still, the traditional theory does not yield a general explanation of the Possibility unless people *must* use names and natural kind terms as if they were such descriptions. But suppose that Locke used “water” as a rigid designator; say, as a synonym for “the substance with the hidden microstructure of what actually falls as rain.” (Surely it would have been possible for him to have used the word this way. It seems that Kant, for one, *did* use it as a rigid designator.¹⁰) In that case it remains false that Locke knew that there was H_2O , even though he did know there was water, and even though as he used the term “water,” the proposition that there is water is the same as the proposition that there is H_2O (since it is a necessary truth that the substance with the hidden microstructure of what actually falls as rain is H_2O).

The upshot for the traditional approach is that it cannot maintain *both* that the content of the knowledge that there is water is the proposition that there is water *and* that the content of the knowledge that there is H_2O is the proposition that there is H_2O , if it is to succeed in giving a truly general explanation of the Possibility. This is an awkward implication, for two reasons. First, it is difficult to say which proposition you *do* know in knowing that there is water, if not the proposition that there is water, and similarly for H_2O . A person might know that there is H_2O without knowing any more about water than someone who does not know that there is H_2O , other than that water is called “ H_2O .” But what one knows in knowing there is H_2O can hardly be that water is called “ H_2O ,” since it is neither necessary nor sufficient for the existence of H_2O that anything be called “ H_2O .”¹¹

⁹ Thus Kripke (1972, pp. 116–140).

¹⁰ See Kant (1965/1781, A728/B756). In the same place, Kant argues that our use of natural kind terms is non-descriptive, anticipating the “new” theory of reference by about 200 years.

¹¹ One might equate knowing that there is H_2O with knowing that there exists that which is *actually* called “ H_2O .” However, this would lead away from the traditional explanation of the Possibility, since the possible worlds in which there exists what is actually called “ H_2O ” are the same as the possible worlds in which there exists water.

Second, it is difficult to accept the implication that “there is water” and “there is H₂O” express the same proposition standing on their own, but different propositions when they occur in the scope of an epistemic verb. For if sentences expressed different propositions depending on whether they occurred in direct or indirect discourse, then an inference like:

My broker tells me that American Necktie is a bargain at \$17.

Everything my broker tells me is true.

American Necktie is a bargain at \$17.

would be invalid, despite following a pattern that we rely on all the time in drawing conclusions about the world on the basis of people’s linguistic and epistemic performance.¹²

2.2 Structured propositions

More recent versions of the standard view steer clear of such difficulties. According to these, the objects of thought are not propositions at all, or at any rate not just propositions. There are two main strategies of this sort. One can begin with propositions (individuated as sets of logically possible worlds) and then specify the objects of thought as combinations of propositions with other entities, the most popular candidates being linguistic entities (such as expression types), psychological entities (such as mental states), and abstract constructions out of sets.¹³ According to this “structured proposition” (or “guised proposition”) theory, the beliefs that there is water and that there is H₂O have the same propositional content (corresponding to the same set of possible worlds), but they have this content only by virtue of having for their objects distinct, complex entities that somehow package or incorporate this proposition. To take a crude example, the object of the belief that there is water might be a structure consisting of the proposition that there is water combined with the neural states that characterize the brain of someone who is prepared to assert, “There is water,” while the object of the belief that there is H₂O might be a distinct structure consisting of the same proposition combined with certain other neural states (those characteristic of someone who is prepared to assert “There is H₂O”). Since these are different neural states, the structured proposition theory succeeds in distinguishing between the objects of the two beliefs, thus explaining the possibility of having one without the other. The structured proposition theory handles other cases similarly, such as by positing differently structured propositions as the objects of the beliefs that Twain wrote *Huckleberry Finn* and that Clemens wrote *Huckleberry Finn*, or for the act of wondering whether $\sum_{n=0}^{\infty} 2^{-n} = 2$ and the act of wondering whether $2 = 2$.

¹² This is what Barwise and Perry (1981) call a violation of “semantic innocence”; see also Davidson (1975, pp. 172–173).

¹³ For the first approach, see Richard (1990), Larson and Ludlow (1993), and Larson and Segal (1995, pp. 437–464). For the second, see Kaplan (1969, 225f), Forbes (1990), and Loar (1990). For the third, see Carnap (1956, pp. 53–64), Lewis (1970), Cresswell (1985, pp. 85–92), and Zalta (1988).

On a related approach, the objects of thought are structures that involve functions from contexts of utterance to propositions.¹⁴ According to this theory, the objects of thought are propositions under sentential “characters” that intuitively correspond to the linguistic meanings of declarative sentences. For example, the object of the belief that *she* is a punk rocker may differ from that of the belief that Sheena is a punk rocker, even if she is Sheena, since “She is a punk rocker” expresses different propositions from “Sheena is a punk rocker” in various contexts (and therefore differs in character from it) even if they express the same proposition in this particular context. One advantage of individuating the objects of belief this way is that it helps explain the difference between egocentric and non-egocentric attitudes of identical propositional content (“I want you to give MWP shock treatment” versus “I want you to give me shock treatment”); more on this in our discussion of *de se* cognition. However, there are problems with extending this approach to more standard illustrations of the Possibility, since it is hard to argue that “There is water” differs in character from “There is H₂O,” or that “ $\sum_{n=0}^{\infty} 2^{-n} = 2$ ” differs in character from “ $2 = 2$.” It also seems possible for someone to believe that Sheena is a punk rocker on one occasion under the character of “She is a punk rocker” (seeing her in torn jeans and spiky hairdo) while doubting the same proposition under the same character on some other occasion (seeing her in business tweeds, hair in a tight bun), despite not having changed his mind in the interim.¹⁵

2.3 Two-dimensional semantics

As an alternative to the structured proposition theory, one can individuate objects of thought as sets of elements drawn from a domain that includes things other than (or in addition to) logically possible worlds. The most popular version of this “two-dimensional” approach treats objects of thought as sets of *epistemic* possibilities. More precisely (since which situations are epistemically possible can vary from one person to the next), it treats the objects of a given individual’s thoughts at a given point in time as sets of worlds epistemically possible for that individual, at that time. This theory turns on the fact that a logically impossible world might be epistemically possible for someone, and might go from epistemically possible to epistemically impossible for him as he acquires more information.

Suppose we identify the worlds epistemically possible for an individual at a given time as the ones that are logically possible, as far as he can tell at that time. Then we can say that the object of a person’s thought that *p* is the totality of worlds epistemically possible for him each of whose actuality would, as far as he can tell given the information at his disposal, be compatible with the claim that *p*. By this account, *N*’s thought that *p* has for its object the set of all worlds *w* epistemically possible for *N*

¹⁴ See Kaplan (1990) and Perry (1990).

¹⁵ As I have described it, the structured proposition approach treats objects of thought as complex proposition-involving structures that serve as relata of binary epistemic relations. Alternatively, one might treat belief, knowledge, etc. as triadic relations among believers, propositions, and propositional guises or modes of presentation; the hidden indexical theory of Crimmins and Perry (1989), Crimmins (1992), and Schiffer (1992) takes this approach. One can also try to get by with the “guises” alone, as in sententialism. Each of these proposals seeks to explain the Possibility along standard lines, and, for the purposes of this essay, we may regard them all as variants of the structured proposition approach.

such that N cannot, with only the information at his disposal, validly infer that if w is actual, then not- p .¹⁶

For example, the object of Locke's knowledge that there is water will be the set containing each world epistemically possible for Locke whose actuality he cannot, on the basis of the information at his disposal, infer to be incompatible with the existence of water. Now, as far as Locke can tell, water is a basic element. Therefore, some worlds in which the only water-like substance has no chemical structure are such that Locke cannot infer that their actuality would entail that there was no water. Thus, the object of his knowledge that there is water is a set of epistemically possible worlds that contains some worlds devoid of H_2O .

This does not mean that in knowing there is water, Locke knows something that is true at some logically possible worlds that lack H_2O . The propositional value of his knowledge that there is water is given by the set of logically possible worlds included among the epistemically possible worlds that constitute the complete object of his knowledge that there is water (namely, the set of logically possible worlds in which there is water). This, according to current mainstream thought, is the same as the propositional value of knowledge that there is H_2O . Still, it does not follow that Locke knows that there is H_2O . For all Locke knows, " H_2O " might be the formula for gold. Therefore, he cannot infer from the information he has that the actuality of a world containing nothing but gold would entail that there was no H_2O . The set of worlds epistemically possible for him whose actuality would, as far as he can tell, be consistent with the existence of H_2O therefore includes worlds (such as all-gold worlds) that the object of his knowledge that there is water does not include.

This account explains the possibility of *discovering* that water is H_2O by reference to the possibility for the set of water-containing worlds that are epistemically possible for a given person to come to coincide with the set of worlds epistemically possible for him in which there is H_2O . What makes it possible for someone to wonder whether $\sum_{n=0}^{\infty} 2^{-n} = 2$ without wondering whether $2 = 2$ is the fact that $2 = 2$ in worlds epistemically possible for him in which $\sum_{n=0}^{\infty} 2^{-n} = \infty$.¹⁷

3 The alternative

Earlier (Sect. 2) we considered the basic argument for the standard view; only if that argument is unsound can there be any alternative to this view. The alternative that I

¹⁶ According to the two-dimensional theory sketched here, the objects of thought are roughly equivalent to Jackson's A-intensions (2004) and Chalmers' primary or epistemic intensions (2004a).

¹⁷ Mathematical cases actually pose more of a challenge to two-dimensionalism than this gloss suggests. If the object of a person's belief that $\sum_{n=0}^{\infty} 2^{-n} = 2$ is the set of worlds each of whose actuality would, as far as he can tell given the information at his disposal, be compatible with the claim that $\sum_{n=0}^{\infty} 2^{-n} = 2$, then whether someone can know that $2 = 2$ without knowing that $\sum_{n=0}^{\infty} 2^{-n} = 2$ depends on what information we count as being "at his disposal." No doubt Zeno could have told that the actuality of any and every world in which $2 = 2$ would comport with the claim that $\sum_{n=0}^{\infty} 2^{-n} = 2$, given suitable knowledge of *a priori* truths; thus, for the two-dimensional theory to work in this case, it must make it clear in what sense these *a priori* truths were not at Zeno's disposal. More generally, it must offer an account of information availability that explains how the information that p can be available to someone to whom the information that q is unavailable, in a case where p iff q . Depending on how it handles these complications, the two-dimensional theory may simply revert to a more conventional structured proposition account.

favor rejects its first premise, which says that belief is a relation. I maintain that belief is not a relation, but a family of relations, and likewise for doubt, desire, knowledge, and other familiar forms of cognition. This is the multiple relation theory.¹⁸

The basic idea behind the multiple relation theory is simple. Where the standard view sees the possibility of knowing that there is water without knowing that there is H₂O as analogous to that of marrying Jane without marrying her sister Julia, the multiple relation theory sees it as analogous to the possibility of dating Jane without marrying her. Proponents of the traditional theory explain the possibility of believing that Twain wrote *Huckleberry Finn* without believing that Clemens wrote *Huckleberry Finn* as similar to that of believing there is soda without believing there is beer; I explain it as similar to the possibility of hoping there is beer without believing that there is. In general, the standard view sees the Possibility as arising in cases where a person bears a psychological relation to one object of thought while failing to bear the same relation to a different, propositionally equivalent object of thought, whereas the multiple relation theory sees it as arising in cases where a person bears one psychological relation but not another, similar relation to a single object of thought (a single set of possible worlds, as the case may be).

To say that belief is a family of relations is to say something about the practice of ascribing beliefs. It is to say that there is no relation such that to ascribe a belief just is to put it forth that this relation holds between a believer and something that can be believed; no relation we can point to and declare: “The whole point of engaging in the practice of belief attribution is to indicate the obtainment of this relation between believers and what they believe.” For everyday purposes, it may be neither necessary nor practical to have a separate label for each way of bearing a proposition in mind, any more than for each discernible shade of blue. On the multiple relation theory, we do not use the verb “believe” to designate the same relation in every use, but to designate any one out of a range of epistemic relations, depending on changing circumstances of use. Other epistemic verbs correspond to other ranges of relations, so that, like the buttons on an old car radio, using a particular one of them places the speaker’s meaning within a specific interval of the overall spectrum of epistemic relations, but leaves it to contextual factors to tune in to a specific relation in that range.

The multiple relation theory therefore treats verbs of propositional attitude ascription as indexical predicates, relying for their representational content in any given context on certain features of that context. But which features are these? They are not, or at least not primarily, the contextual features most commonly associated with indexicality, such as the identity and physical location of the speaker, or the time at which the speech act occurs. The content of a verb like “to believe” is not

¹⁸ The theory I propose bears a superficial resemblance to Russell’s multiple relation theory of judgement, according to which belief is a series of relations of increasing adicity (1985/1918, pp. 79–93); unlike Russell’s theory, mine does not treat belief as involving relations of more than two places. The multiple relation theory defended here has a closer affinity to epistemic contextualism, which construes “knows” as having the capacity to designate different dyadic relations in different contexts of utterance; see Lewis (1979b), Unger (1984), and DeRose (1992). Within the literature on opacity and attitude ascription, only Richard regards attitude verbs as indexical; see (Richard, 1990, p. 107). However, the indexicality thesis plays only a secondary role in his overall account, which treats propositional attitudes as triadic relations among agents, propositions, and “mappings of Russellian annotated matrices,” and holds that any clause occurring in the *scope* of an attitude verb is thereby rendered indexical: see Richard (1990, pp. 140–147).

generally sensitive to who uses it, or when or where he does so. But these are not the only features of a context on which an indexical might rely for its content. In addition to these *situational* features (as I shall call them), a context exhibits various discursive features. These are features of the discourse (conversation, monologue, or train of thought) in which an utterance takes place.

The discursive features of a context of utterance roughly comprise (a) what is being discussed in that context; (b) what is being said about (a); and, (c) the history of the conversation or line of thought (if any) which has culminated in (b) being said about (a). Without pretending to give an exhaustive account of discursive contextual features, we can say, slightly more precisely, that given an occasion of word use *u*, the *topic* of the context in which *u* occurs is a function of the following factors:¹⁹

- (a) the *theme* of the context: what is under discussion or consideration in the context in which *u* occurs, up to *u*'s occurrence;
- (b) the *conceptual background* of the context: what concepts have been used in the conversation or inquiry of which *u* is a part, up to and including the point at which *u* occurs; and,
- (c) the *discursive history* of the context: what conversational or dialectical moves have been made up to *u*—what questions raised, discussed, or dismissed; what assertions made or challenged; what evidence brought forward; etc.

These three factors loosely add up to what we normally think of as *topic of conversation*, and we may call context-dependent terms that depend primarily on such factors for their semantic content *topical indexicals*, in contrast to *situational indexicals*, such as “I” and “today,” which vary in content primarily in response to situational contextual factors. Many indexicals are mixed in nature, depending on a variety of topical and situational factors: “here” and “soon” are good examples. This dual-dependence is sometimes obscured by idealized or oversimplified semantic treatments, according to which, for example, the extension of “here” in a given use is simply a matter of the location at which that use occurs. This is fine as a rough characterization, but its illicit appeal to “the” location at which the indexical gets used obscures the fact that topical factors often play a major role in determining a meaning for this word. Similarly, “soon” may designate a point in time more or less distant from the time of its utterance, depending on the topic: it designates a more distant point in “Soon the human population will exceed 7 billion” than in “The baby will fall asleep soon.”

Of course, indexicals like these also depend quite heavily on situational factors. Except under unusual circumstances, a person speaking in New York in the year 2006 cannot use “here” to refer to Singapore, or “soon” to refer to some time a billion years hence. An exception might be a case in which the person points to a map or timeline while making his utterance; but normally, the significance of a mixed indexical like this is heavily constrained by situational factors, such as location and time of utterance. Still, such terms' sensitivity to topical factors is an important part of their use. In general, we use indexicals to capitalize on the circumstances surrounding a given act of speech, by directing the hearer's attention to certain aspects

¹⁹ Formally, the multiple relation theory would treat epistemic verbs as binary predicates of a double-indexed semantics—in effect, an extension of Kaplan's logic of demonstratives, with verbs like “to believe” indexed to discursive contextual factors, and contexts of utterance enriched correspondingly; see Kaplan (1989, pp. 541–546). I make no attempt to formalize the theory in this paper.

of these circumstances in a way that is likely to bring about mutual understanding. Indexicals like “I” and “today” primarily exploit situational aspects of these circumstances, while “here” and “soon” exploit both situational and topical aspects. This makes the use of a mixed indexical more complicated, more subject to misunderstanding, but also richer and more flexible than that of a situational indexical.

As a conversation (or story, or solitary meditation) progresses, discursive factors accumulate: new considerations are brought to the table, new evidence is gathered, and a history of conversational moves (assertions, challenges, etc.) builds up as the conversation unfolds. We may think of these accumulated discursive factors as so much semantic potential energy. Mixed indexicals like “here” and “soon” often tap into this potential, by allowing discursive contextual factors to bear on their semantic values. According to the multiple relation theory, epistemic terms tap this potential in a similar but more aggressive way: similar, since here too there is sensitivity to topic; but more aggressive, since unlike a mixed indexical, a topical indexical is sensitive to little besides discursive factors. Whereas mixed indexicals are tuned to two contextual frequencies (discursive and situational), topical indexicals receive at just one frequency (the discursive).

So, if epistemic representations are indexical, they differ from textbook indexicals by capitalizing more exclusively on discursive contextual features. Particularly important to the multiple relation theory are the reasons available in a given context for accepting a given propositional attitude ascription as true, or rejecting it as false. Obviously, these reasons can take many forms, and on different occasions one might have different reasons for attributing belief of the same proposition to the same individual. Typically the reasons have to do with the behavior of the ascriber. Normally, although not always, we are prepared to defend an assertion to the effect that so-and-so believes that such-and-such by reference to so-and-so’s overt linguistic behavior, such as his having assertively uttered some sentence to the effect that such-and-such. Non-linguistic reasons to suppose that someone believes a given proposition include the proposition’s being obvious to most people in that person’s position, or an obvious consequence of such a proposition. But such reasons are generally secondary. Even when we accept an attitude ascription on the strength of second-hand information, this information itself normally has to do with the reported linguistic propensities of the ascriber.

Sometimes it is unclear what evidence, if any, the person who ascribes a belief takes himself to have for the ascription (or, intends for others to take him to have). But often the supposed reasons for accepting or rejecting a given attitude ascription are clear enough to form part of the common ground of the discourse in which the ascription is made. More important for present purposes, the Possibility can arise only in discursive contexts that *do* in one way or another supply reasons to accept one attitude ascription and reject a corresponding ascription. If there were no clear justification for claiming that Locke knew that there was water, or no clear justification for denying that he knew that there was H₂O, we would have no clear reason to accept (1) and reject (2) in the first place. And if we had no clear reasons to accept (1) and reject (2), and likewise for similar cases, we would have no reason to acknowledge the Possibility at all.

This consideration forestalls an otherwise natural objection to the multiple relation approach, which is that it cannot account for the compatibility of (1) with the negation of (2) in contexts wherein the verb “to know” represents the same cognitive relation in both. It is true that the multiple relation theory cannot account for this, but only because what it is being asked to account for is a fiction. Again: *any* case in which we are justified in accepting (1) and rejecting (2) is a case in which we have some basis for doing so. It cannot be just a brute fact that (1) is true and (2) false. But whatever basis exists for accepting (1) we can treat as part of what we mean by “knew” as it occurs therein, and likewise we may factor into what we mean by “know” in (2) our basis, whatever it may be, for rejecting (2). A case that afforded no such opportunity to distinguish between the relations expressed by the verb in the two utterances would be a case in which we had no basis to think that the utterances diverged in truth value. In such cases, the Possibility simply does not arise to be accounted for.

However, it is important not to overstate this reply. The Possibility may well arise *vis à vis* a true statement such as:

$$(3) \quad (\exists\phi)(\exists\psi)(\Box(\phi \leftrightarrow \psi) \ \& \ \text{John believes that } \phi \ \& \ \text{John does not believe that } \psi).$$

Now, if we are supposing the statement true, we are supposing that it has a true instance. But what if we do not have any specific instance in mind? Then, on my view, what we accept in accepting (3) is:

$$(4) \quad (\exists\phi)(\exists R_1)(\exists R_2)(R_1 \text{ is a belief relation} \ \& \ R_2 \text{ is a belief relation} \ \& \ R_1(\text{John}, \phi) \ \& \ \neg R_2(\text{John}, \phi)).$$

Here, the existence of the belief relations required to make (4) come out true is guaranteed by the fact (if it is a fact) that we have good reason to accept (3). For we have good reason to accept (3) only if we have good reason to think that for *some* pair of propositionally equivalent sentences, S and S', John can be got to assent to S but not S', or to manifest some other, similar complex of dispositions.

This connects with the multiple relation theory's treatment of *de re* attitude ascriptions. According to this, to say that John believes of Clemens that he wrote *Huckleberry Finn* is to say, roughly, that

$$(5) \quad (\exists R)(R \text{ belongs to the class of belief relations} \ \& \ R(\text{John}, \text{the set of possible worlds in which Clemens writes } \textit{Huckleberry Finn})).$$

This is only rough, because it fails to distinguish between, e.g., “John believes, of Clemens, that he wrote *Huckleberry Finn*” and “John believes, of *Huckleberry Finn*, that Clemens wrote it.” What (5) really gives is the multiple relation analysis of the “pure” *de re* ascription: “John believes, of Clemens, writing, and *Huckleberry Finn*, that the first did the second to the third.” To capture the difference between the two impure *de re* attributions, we must give separate, modified versions of (5), each including a further clause to the effect that R is a relation that someone bears to a proposition only if he has the sort of dispositions appropriate to having one of these *de re* beliefs versus the other.

3.1 Twain/Clemens

According to the multiple relation theory, the inference from:

- (6) Eva knows that Mark Twain wrote *Huckleberry Finn*.

to:

- (7) Eva knows that Samuel Clemens wrote *Huckleberry Finn*.

is invalid for basically the same reason we cannot validly infer from:

- (8) Eva is looking for a husband.

to:

- (9) Eva is looking for a married man.

Since (8) might be true even in a case where (9) is false, a proponent of the standard view might want to conclude that there is some difference between the content of “husband” and that of “married man.” But this conclusion is not forced upon us, since we may account for the invalidity of the inference from (8) to (9) as due rather to a variation in the significance of the verb phrase “looking for.” In (8), “looking for” stands for a condition that Eva is in as long as she seeks to make something her own, while in (9), it stands for a condition that she is in as long as she seeks to ascertain the whereabouts of something. What accounts for this difference is the occurrence of “married man” in place of “husband” in (9). The occurrence of “husband” in (8) fixes the meaning of “looking for” as one that involves a certain combination of intentions and behavior; the occurrence of “married man” in (9) has a different effect, fixing the meaning of “looking for” as one that involves a different combination of intentions and behavior. Since Eva can exhibit the former combination without exhibiting the latter, the inference from (8) to (9) is invalid, despite the fact that “husband” and “married man” do not differ in content (extension or intension).²⁰

Similarly, according to the multiple relation theory, the word “knows” in (6) expresses a relation in which Eva stands to the proposition that Mark Twain wrote *Huckleberry Finn* only if she exhibits certain dispositions, such as the disposition to use the sentence “Mark Twain wrote *Huckleberry Finn*” to make an assertion. Call this relation, whatever it comes to in detail, “t-knowledge,” and the proposition that Twain wrote *Huckleberry Finn* (= the proposition that Clemens wrote *Huckleberry Finn*) “the proposition that q.” In (7), the verb “to know” signifies a different relation—one in which Eva stands to the proposition that q only if she is prepared to make an assertion with the sentence “Samuel Clemens wrote *Huckleberry Finn*,” or to engage in some relevantly similar form of behavior. Call this relation, whatever it

²⁰ We could try to analyze (8) and (9) differently from how I have proposed; for example, we might construe (8) as equivalent to “Eva is looking for an eligible bachelor,” and (9) as equivalent to “Eva is looking for a man who is married.” This has the disadvantage of erroneously implying that the conjunction of (8) and (9) is incompatible with the claim that Eva is looking for just one man.

may come to in detail, “c-knowledge.” As long as c-knowing that q is not necessary for t-knowing that q , the inference from (6) to (7) is invalid. And c-knowing that q will not be necessary for t-knowing that q in any case in which (6) and (7) exemplify the Possibility, since, as we noted earlier, if we had no consistent evidence for the truth of (6) and the falsity of (7), we would have no good reason to regard the inference from the one to the other as invalid in the first place. Thus we can always account for the compatibility of (6) with the negation of (7) without supposing there to be any difference in content between the belief that Twain wrote the book and the belief that Clemens wrote it.

I have focused on the role of background assumptions about the linguistic behavior of attitude ascribers in determining contents for epistemic concepts and representations. As previously observed, this is not the only kind of evidence we might have to accept or reject an attitude ascription. If we know that Eva’s knowledge of Twain is limited to what she acquired from a teacher notoriously ignorant of the fact that Twain wrote under an assumed name, we might infer on that basis that she knows that Twain wrote *Huckleberry Finn* but not that Samuel Clemens did so. Our exact reasons for accepting (6) but not (7) are unimportant. As long as we have *some* reason to believe Eva knows that Twain wrote *Huckleberry Finn* and some reason to believe she does not know that Clemens wrote it, the presence of these reasons in the discourse context can be brought to bear on the content of the verb “to know” in the way the multiple relation theory suggests.

3.2 Water/H₂O

The multiple relation theory also accounts for the possibility of discovering that water is H₂O. The discovery amounted to the discoverer’s coming to bear a new cognitive relation to the proposition that water is water, a relation that requires for its obtainment between him and this proposition that he be prepared to express it with the words “water is H₂O,” or to engage in some other relevant form of behavior, linguistic or otherwise. If someone doesn’t know that water is H₂O, he cannot infer that there is H₂O in the test tube from the fact that there is water in it. Nor can he translate a qualitative description of soap-making into a chemical equation for the hydrolysis of fatty acids. And in general, not knowing that water is H₂O means not being able to use the expression “H₂O” in certain ways in which someone who does know that water is H₂O *can* use it. Perhaps there are also non-linguistic ways of knowing that water is H₂O. Maybe there are even ways of knowing it that do not give rise to any overt behavior at all. But the *capacity* for such behavior is certainly a requirement for knowledge that there is H₂O, and this behavior is paradigmatically language-involving.

None of this implies that knowing that water is H₂O is knowing some linguistic fact. What we know in knowing that water is H₂O has nothing whatever to do with language, since water is H₂O in every possible world, including possible worlds that contain no language. In order to know that water is H₂O, we may well require certain linguistic capacities, but this does not entail that *what* we know in knowing that water is H₂O has something to do with language.

On this view, the difference in cognitive significance between the thought that water is water and the thought that water is H_2O does not lie in what in the world it takes for these thoughts to be true. Rather, it lies in the fact that knowing that water is H_2O involves having certain (paradigmatically linguistic) abilities, whereas knowing that water is water does not involve these abilities. There is therefore no need to distinguish between the content of the thought that water is water and the content of the thought that water is H_2O .

$$3.3 \quad 1 + \frac{1}{2} + \frac{1}{4} + \frac{1}{8} + \dots / 2$$

In addition to standard water/ H_2O cases, the multiple relation theory explains the possibility of mathematical discovery. The difference between these cases is really only one of degree. Given the multiple relation theory, it can be the same thing you know in knowing that $\sum_{n=0}^{\infty} 2^{-n} = 2$ as you know in knowing that $2 = 2$, that π is irrational, that azure is a shade of blue, or any other necessary truth; “ $\sum_{n=0}^{\infty} 2^{-n} = 2$ ” and “ $2 = 2$ ” need have no truth-assessable content besides the propositional content they share. This idea is not as eccentric as it sounds. It is obvious that necessary truths do not rule out any logical possibilities; as far as the fact that $\sum_{n=0}^{\infty} 2^{-n} = 2$ is concerned, any possible world might logically be our own. Modern versions of the standard view recognize this by arranging for all statements of necessary truth to have the same propositional content, and cooking up an appropriate array of complex entities to serve as the objects of distinct mathematical beliefs. The multiple relation theory regards correct mathematical beliefs as identical in their object, distinguishing them according to the behavior and capacities of mathematical believers.

Proponents of the standard view might argue that the proposition that $2 = 2$ is not a proposition about infinite series, and that this shows that we have to distinguish between this proposition and the one whose verification solved Zeno’s most enduring paradoxes. But what is a proposition “about”? If it is about whatever must exist for it to be true, then the proposition that $\sum_{n=0}^{\infty} 2^{-n} = 2$ and the proposition that $2 = 2$ are about the same thing: since both of them are true in every possible world, neither depends for its truth on the existence of anything but what necessarily exists. If the proposition that p depends for what it is about on what we do with the statement “ p ,” then the claim that there is a difference between what the proposition that $2 = 2$ is about and what the proposition that $\sum_{n=0}^{\infty} 2^{-n} = 2$ is about boils down to the claim that we put “ $2 = 2$ ” and “ $\sum_{n=0}^{\infty} 2^{-n} = 2$ ” to different uses. Of course we do put them to different uses despite their logical equivalence, and the existence of this difference is all the multiple relation theory needs to do its job.²¹

3.4 Explanatory power

The task of explaining the Possibility is that of explaining the difference between believing that there is water and believing that there is H_2O in a way that leaves open the possibility for someone to believe there is water without believing there is H_2O , and likewise for other, similar examples. The standard view sees the difference

²¹ Most important, these statements engage differently with the rules that govern the language of arithmetic, with the result that not everything derivable by these rules from the one statement is derivable by them from the other.

between water-believers and H₂O-believers as a difference between what they believe; the multiple relation theory sees it as a difference in how they behave or are disposed to behave, but not in what they believe. This can make it seem as if the standard view explains more than the multiple relation theory. Since the multiple relation theory uses behavioral differences to explain the Possibility, it can hardly constitute an explanation of those behavioral differences. By contrast, the standard view appears able to explain both the Possibility and the behavioral differences between water-believers and H₂O-believers by reference to distinct objects of thought.

This appearance is misleading. It is one thing to ask why some people behave in ways that justify us in describing them as knowing that there is water but not that there is H₂O. Why people exhibit this variety of behavior is an empirical question that neither the standard view nor the multiple relation theory properly addresses. What explains the fact that a person engages in the sort of behavior that prompts us to describe him as knowing that there is H₂O is, ultimately, that he has a certain kind of body that interacts in certain ways with a certain kind of environment. A different question is: what does the fact that we are justified in describing some people as knowing that there is water but not that there is H₂O tell us about the nature of knowledge and its objects? This a conceptual question, and it is in the answers they give to this question alone that the standard view and the multiple relation theory differ. When it comes to explanatory power, therefore, neither view has any special advantage over the other.

4 Multiple relations versus multiple objects

In this final part of the paper, I argue that the multiple relation theory has various advantages over the standard view, in addition to those already discussed in relation to the issue of compositionality. The first three advantages discussed below concern matters of semantic detail that any complete discussion of propositional attitudes must address. The last advantage is of a more metaphysical nature, and concerns certain broad questions in the philosophy of mind.

4.1 Quantifying in

From the fact that

- (10) Homer believes that Mark Twain wrote *Huckleberry Finn*.

it would appear to follow that

- (11) There is someone who Homer believes wrote *Huckleberry Finn*.

The standard view has considerable difficulty accounting for this. It has trouble dealing with any inference that proceeds by quantifying into the scope of an epistemic operator, as in (11), which has the form:

$$\exists x(\text{Homer believes that } x \text{ wrote } \textit{Huckleberry Finn}).$$

On the standard view, any term we use to complete the open sentence
Homer believes that x wrote *Huckleberry Finn*.

must play a different semantic role there from what it normally plays. Normally, “Mark Twain” and “Samuel Clemens” make the same semantic contribution to the sentences in which they occur. But according to the standard view, they make different contributions to the semantic value of a sentence like (10); this is how the standard view explains the fact that we are not free to substitute these names for one another in (10). The standard view therefore sees occurrence within the scope of an epistemic operator as akin to occurrence in quotes (“‘Mark Twain’ contains nine letters”) or in logophoric contexts (“Mark Twain is so-called after a sounding of two fathoms”). On the strength of this analogy, Quine goes so far as to suggest that (11) is as senseless as

Someone is so-called after a sounding of two fathoms.²²

Other proponents of the standard view tread more lightly, arguing that the inference from (10) to (11) is valid, albeit not as a direct application of existential introduction like the one that takes us from:

Mark Twain wrote *Huckleberry Finn*.

to:

Someone wrote *Huckleberry Finn*.

The key, on these views, is to pry the reference to Twain in (10) out of the scope of the verb “believes,” and only then apply existential introduction. There is a variety of techniques for accomplishing this, but the basic idea behind all of them is to interpret (10) as having the following general form:

Mark Twain is such that Homer thinks of him in a certain way W , and W is a component of what Homer believes in believing that p .

where “ p ” designates whatever the standard view identifies as the object of Homer’s belief that Twain wrote *Huckleberry Finn*.²³

The details of these accounts are complex, and it is controversial how far they justify quantifying in. This is in stark contrast to the multiple relation theory, on which the move from (10) to (11) is licensed by a perfectly ordinary application of existential introduction. On the multiple relation theory, “Mark Twain” occurs in (10) in a perfectly ordinary way; all that prevents us from replacing it with the semantically equivalent “Samuel Clemens” are the different effects these terms have on the content of the accompanying “believes.” But the fact that this verb can

²² See Quine (1953b, 1956).

²³ See Kaplan (1969) and Forbes (1996).

express different relations in different contexts no more prevents us from quantifying into its scope than the ambiguity of “lies” prevents us from quantifying into:

Homer lies on the couch.

to get:

There is something on which Homer lies.

Whatever may be the prospects for standard accounts of quantifying in, it is a distinct advantage of the multiple relation theory that it affords such a simple and well-integrated justification of this practice.²⁴

4.2 John Smith a.k.a. John Smith

Suppose George Jones knows a man named John Smith who mows his lawn once a week every summer for 40 dollars. If we were to ask Jones if he thought that Smith was a physician, he would sincerely reply in the negative. Thus it appears that

(12) Jones does not believe that John Smith is a physician.

Unbeknownst to Jones, however, Smith is a retired surgeon who picks up yard work on the side as an eccentric form of relaxation. Late one winter, Jones takes his wife on a Caribbean cruise. Aboard the ship they make various new acquaintances, including that of a Dr. and Mrs. Smith. Although the latter is the man who mows his lawn, Jones does not recognize him, since Smith has recently lost a lot of weight, and in any event looks very different in a dinner jacket from how he looks in coveralls, sunglasses, and a floppy straw hat. Mrs. Jones learns from Mrs. Smith that Dr. Smith is a medical doctor, which information she passes on to her husband, who takes it at face value. Thus it appears that

(13) Jones believes that John Smith is a physician.²⁵

Considered in one light, cases like this practically *entail* a multiple relation account of epistemic concepts and idioms. How else could (12) be consistent with (13), other than by a difference in significance of a term that occurs in both statements? And what term do the members of every such pair of statements have in common, but the main epistemic verb? Even the standard view must have recourse to some doubling

²⁴ Of course, the practice is not justified when the position to be quantified is occupied by an empty term, but then the culprit is not the propositional attitude verb, but the empty term: we can no more go from “Santa does not exist” to “ $\exists x(x \text{ does not exist})$ ” than from “Tom believes that Santa lives at the North Pole” to “ $\exists x(\text{Tom believes that } x \text{ lives at the North Pole})$.”

²⁵ This scenario is reminiscent of Kripke (1988, pp. 130–131), but owes its direct inspiration to a conversation with Mitch Green.

of semantic role to explain this kind of case, assigning different semantic values to “John Smith is a physician” when tokened in relevantly different contexts.²⁶

This reveals an unexpected advantage of the multiple relation theory over the standard view. *Any* sentence can occur in the scope of an epistemic verb, and therefore any sentence might figure in a Smith/Smith type of scenario. Since the standard view handles such cases only by holding that the pertinent embedded sentence varies in semantic value, the standard view implies that every sentence has the capacity to vary in this way. This includes sentences that themselves contain epistemic idioms. We can construct a Smith/Smith scenario for:

- (14) Brown doesn’t believe that Jones believes that Smith is a physician.

and:

- (15) Brown believes that Jones believes that Smith is a physician.

no less than for (12) and (13). In effect, the standard view works only on the assumption that every linguistic expression varies in semantic value, where the multiple relation theory requires only that a small class of expressions (the propositional attitude verbs) vary in semantic value. To the extent that multiplying meanings is an issue, it surprisingly weighs for, rather than against, the multiple relation theory.²⁷

4.3 Attitudes *de se*

Suppose I notice someone reflected in a plate glass window and think to myself: “He has bad posture.” This thought does not prompt me to any particular course of action; I just go about my business. But suppose that unbeknownst to me, the reflection I see is my own. What I know in knowing that *he* has bad

²⁶ The hidden indexical theory posits an unrealized component of (12) and (13) to account for their compatibility. According to this theory, propositional attitude ascriptions normally contain a covert ellipsis. Such ellipsis is not totally unheard of; as Reimer (1996) points out, someone might go through life unaware that his utterances of statements like “I’ll be home by noon” can be construed as elliptical for statements that contain an explicit reference to a time zone (“I’ll be home by noon, Eastern Standard Time”). But if a person really has no conception of time zones, then he cannot make elliptical reference to them unless he relies on the broader linguistic community of which he is a member, where this community uses chronological language in accordance with a time zone convention. Similarly, the ellipsis posited by the hidden indexical theory would require some background consensus on the logic of belief, knowledge, etc., if only as much as what lies behind our everyday use of clock vocabulary. But there is no such background consensus in this case, as the long-standing debate about the logic of such ascriptions attests. More crucially for the central concerns of this paper, the hidden indexical theory follows the standard view in positing a difference of cognitive objects to account for the compatibility of (12) with (13) (or of (1) with the negation of (2)); it is just that the difference occurs at the level of a third, indexically determined relatum of the (supposedly unique) relation of belief.

²⁷ The multiple relation theory has no trouble with iterated attitude ascriptions such as (14) and (15). Take the Mates-style ascription: “Brown doubts that everyone who believes that physicians are physicians believes that physicians are doctors”—See Mates (1952). According to the multiple relation theory, this has the following form: “Brown doubts that everyone who believes₁ that *p* believes₂ that *p*,” where “*p*” designates the set of possible worlds in which physicians are doctors (= the set of worlds in which physicians are physicians), and the relational contents of “believes₁” and “believes₂” depend partly on discursive elements of the context in which we attribute this doubt to Brown.

posture then has the same propositional content as what I would know if I knew that *I* had bad posture. Yet, if I knew that *I* had bad posture, I would stop to correct it.²⁸

This example shows that there can be a difference between my knowing that *he* has bad posture and my knowing that *I* have bad posture, even when I am *he*. Characteristically, the standard view sees the task of explaining this difference as reducing to “the problem of finding appropriate objects for such knowledge.”²⁹ But here the standard theorist faces special challenges. No matter what we substitute for the phrase “I have bad posture” in:

(16) I know that I have bad posture.

it will fail to yield an equivalent knowledge attribution, unless it contains some term of personal self-reference. For example, we might try to paraphrase (16) with the statement:

(17) I know that the person who is speaking has bad posture.

But this does not work. From (16) and the fact that I am disposed to maintain good posture, we may infer that I will correct my posture. But we may not infer this from the same fact and (17). Only given (17) and the further assumption that I believe that *I* am the person who is speaking may we infer that I will correct my posture.

The traditional theory is ill-equipped to deal with *de se* cognition. It must maintain that in knowing that I have bad posture, I know a proposition distinct from the proposition that MWP has bad posture, distinct from the proposition that *he* has bad posture (where “he” refers to me), and distinct from any proposition expressible without the aid of the first person pronoun. What could this proposition be? At which possible worlds does it diverge in truth value from the *prima facie* identical proposition that MWP has bad posture? To this the traditional theory has no good answer.³⁰

Modern variants of the standard view scarcely do better. John Perry entertains a view according to which the objects of *de se* thought are “relativized propositions”—abstract objects corresponding to sentences containing indexicals.”³¹ These are a form of structured proposition, with sentential characters (distributions of propositions over contexts of utterance) providing the structure. Thus, in order for me to have *de se* knowledge that my posture is bad, I must bear the relation of knowledge to a structure that combines the character of the sentence “my posture is bad” with the proposition this sentence expresses in contexts in which I am the speaker or thinker. I have only ordinary, non-egocentric knowledge that my posture

²⁸ For other examples (on which this one is based), see Castañeda (1966), Perry (1979), and Richard (1988, pp. 183–188).

²⁹ Lewis (1979a, p. 522).

³⁰ Frege, who, on the strength of a comment in (1966/1892), is widely associated with the traditional theory, seems to advocate a guised proposition account of *de se* knowledge: see Frege (1997/1918, pp. 332–333).

³¹ Perry (1979, p. 97). Perry’s own views on the matter are not entirely clear, but they appear to have some affinity with the multiple relation approach.

is bad if the object of my thought includes the same proposition, but fails to incorporate this character (perhaps it incorporates instead the character of the statement “that person’s posture is bad”).

David Lewis offers a different (though still standard) account of *de se* cognition. According to him, the objects of thought are sets of possible *entities* rather than simply sets of possible worlds. This way we can still construe many statements as having standard propositional contents corresponding to sets of possible worlds (possible worlds being possible entities), and we can differentiate *de se* knowledge from ordinary knowledge by supposing the former but not the latter to involve a cognitive relation to a set containing some possible individuals that are *not* worlds, such as possible versions of myself (as opposed to near counterparts of me) perceiving my reflection in the window.³²

Equivalently, we can identify the objects of thought as sets of “centered” possible worlds. A centered world is a possible world with an emphasis on one location in it, like a map marked with an X. I have *de se* knowledge that my posture is bad only if what I know is the set of worlds where my posture is bad and the X marks me. If the object of my knowledge contains all worlds in which I have bad posture, including those *not* centered on me, then I know *of myself* that my posture is bad, but lack the *de se* knowledge *that I have bad posture*. Proponents of two-dimensional semantics have co-opted this approach (modified in some cases to include centered *impossible* worlds) to explain the difference between *de se* beliefs and their non-egocentric propositional equivalents.³³

All of these accounts are ontologically burdensome. Relativized propositions, individual possibilities, and centered worlds have no apparent claim to reality, beyond the fact that positing them allows us to distinguish between the objects of ordinary and *de se* thought. By contrast, on the multiple relation theory, the content of a *de se* belief is just a proposition, and the same proposition that is the object of an ordinary *de dicto* (or even *de re*) belief. Instead of construing the difference between egocentric and non-egocentric knowledge as a difference between knowing an egocentric versus a non-egocentric thing, the multiple relation theory explains it as a difference between knowing the same thing egocentrically versus non-egocentrically. For me to know that *my* posture is bad is for me to relate to the proposition that MWP’s posture is bad via a relation that a person bears to this proposition only if he acts on this knowledge by correcting his posture if he is disposed to maintain good posture. When we ask if I realized that it was *my* posture that was bad, we are not asking whether I possessed some further item of information than what has already been attributed to me, but whether I acted upon the information already attributed to me in certain ways (such as by straightening up). Thus we are spared the need to

³² Lewis (1979a, pp. 519–522).

³³ See Chalmers (1996), and Chalmers (2004a), and, for Lewis’s discussion of centered worlds, Lewis (1979a, pp. 531–532).

posit Lewisian possibilities or other exotic objects of cognition—welcome news to those who share Quine’s misgivings about such putative entities.³⁴

4.4 Multiple relations and the knowledge argument

If the multiple relation theory is a viable alternative to standard theories, this has an important bearing on the debate over traditional physicalist naturalism (the position that the physical facts about the world exhaust the empirical facts about it). If the standard view is correct, it lends considerable support to a main argument against physicalism. This argument concludes that physicalism is false on the grounds that a physically omniscient being could discover something about conscious experience.³⁵ In Jackson’s familiar thought-experiment, Mary has all the physical information in the world, but up until now she has always lived in a black and white prison. When she is finally released to view colorful objects for the first time, it is hard to deny that the experiences she has are enlightening, involving some element of discovery.³⁶ But according to the standard view, making a discovery involves coming to know something (whether a proposition, a structured proposition, or an epistemic intension) previously unknown to the discoverer. Since Mary already knew everything about the physical nature of the world, it follows that there is more to be known about the world than its physical nature, in which case physicalism is false.³⁷

Some physicalists argue that we can allow that Mary acquires factual knowledge without conceding that she gains knowledge of any new fact. On this view, we can explain Mary’s learning as a matter of acquiring a new item of information (structured proposition, epistemic intension, or what have you) that is factually (and propositionally) equivalent to some item of information that she already possessed. In this way, modern variants of the standard view are supposed to allow us to concede that not all information is physical information without committing ourselves to the physicalistically unacceptable claim that not all facts are physical facts.³⁸

In reality, they have the opposite effect. After all, Mary is supposed to know all the physical facts even before she has her first experience of color. But if physicalism is correct, this means that she already knows all the facts, period. But *all* the facts includes all the facts about all the information, all the facts about all the structured propositions, all the facts about all the sentences, guises, epistemic intensions, etc. Only if Mary could know all these facts without having all the information could we

³⁴ Quine, of course, rejects even logically possible worlds. (Quine (1960, pp. 200–211)) Although the view developed here does not dispense with these, a more aggressive version of the multiple relation theory might. So far we have eschewed structured propositions in favor of unstructured propositions, by allowing discourse contexts to influence the contents of psychological verbs in a way that permits us to regard the sentences occurring in their scope as representing sets of possible worlds. Going a step farther, we might forego propositions in favor of mere truth values, by allowing discourse contexts to select the contents of embedding psychological verbs from a suitably enlarged menu of relations. There is no space in the present essay to pursue this more ambitious proposal.

³⁵ See Broad (1925, p. 71), Nagel (1979), and Jackson (1982).

³⁶ Hard, but not impossible; see Dennett (1991, pp. 398–406) and Jackson (1998b).

³⁷ Wittgenstein (building on an argument from Moore) similarly argues that complete knowledge of mental and physical reality does not guarantee moral knowledge; he concludes that moral facts are supernatural, if they exist at all. See Wittgenstein (1965) and Moore (1903, pp. 10–21).

³⁸ Thus Horgan (1984, pp. 149–152), Loar (1990, pp. 87–90), and Lycan (1990, pp. 113–114, 117–212).

consistently maintain that she acquired new information upon her release. But how could someone know all the facts about all the information without having all the information? Surely she could not, since for every piece of information it is a fact that it is precisely that piece of information. As long as we follow the standard view in assuming that gaining new propositional knowledge requires gaining knowledge of some new cognitive object, even if only in the form of an old fact under a new guise, we are compelled to deny that Mary, who knows all the facts, can gain new propositional knowledge. The standard view therefore places severe constraints on how we can reply to the knowledge argument, and for this reason should not be adopted lightly by those who want to keep their options open with regard to the mind-body problem.³⁹

5 Conclusion

The idea that belief, knowledge and other familiar forms of thought are relations seems innocent enough, and has only been reinforced by the success with which philosophers have engineered logical constructions to serve as distinct objects of distinct thoughts. There is also considerable variety among different developments of the standard view, and this can give it the appearance of exhausting the theoretical possibilities. At a minimum, I hope to have shown that this appearance is misleading, and that the standard view is neither the only view available, nor obviously the most sensible one.

I have not argued that the multiple relation theory yields a faithful picture of how people actually use epistemic verbs. Given that people could do everything they ordinarily do with such verbs by using them as the multiple relation theory proposes, I do not know how to decide the empirical question of whether people do in fact use them this way, or in accordance with some alternative reconstruction of actual linguistic practice. But regardless of how it fares as a linguistic hypothesis, the multiple relation account can at least succeed as a prescription for how to use psychological terms in philosophical, logical, and psychological (or for that matter ordinary) discourse. As long as there is nothing we can say using epistemic verbs in accordance with the standard view that we cannot say using them in accordance with the multiple relation theory, there is no good theoretical reason not to use them as the multiple relation theory recommends. This by itself is enough to temper enthusiasm for philosophical arguments that presuppose the standard view, and that might seem inescapable but for some reasonable alternative to it.⁴⁰

³⁹ Among existing discussions of the knowledge argument, only that of Bigelow and Pargetter (1990) suggests a response of the kind made possible by the multiple relation theory. Most discussions tacitly or explicitly rule out such a response (although never with argument); see, for instance, Stoljar and Nagasawa (2004, p. 9), Hellie (2004, p. 336), and Chalmers (2004, pp. 272–273).

⁴⁰ Walter Ott, Todd Stewart, Tan Yoo Guan, Mark Textor, John Williams, and the editor of this journal provided valuable feedback on previous drafts of this paper.

References

- Barwise, J., & Perry, J. (1981). Semantic innocence and uncompromising situations. *Midwest Studies in Philosophy*, VI, 387–403.
- Bigelow, J., & Pargetter, R. (1990). Acquaintance with qualia. *Theoria*, 56, 129–147.
- Braun, D. (1998). Understanding belief reports. *Philosophical Review*, 107(4), 555–595.
- Broad, C.D. (1925). *The mind and its place in nature*. London: Kegan Paul, Trench, Trubner & Co.
- Carnap, R. (1956). *Meaning and necessity: A study in semantics and modal logic* (2nd ed.). Chicago: University of Chicago Press.
- Castañeda, H.-N. (1966). “He”: A study in the logic of self-consciousness. *Ratio*, 8, 130–157.
- Chalmers, D. J. (1996). *The conscious mind: In search of a fundamental theory*. Oxford: Oxford University Press.
- Chalmers, D. J. (2004a). Epistemic two-dimensional semantics. *Philosophical Studies*, 118(1–2), 153–226.
- Chalmers, D. J. (2004b). Phenomenal concepts and the knowledge argument. In P. Ludlow, Y. Nagasawa, & D. Stoljar, (Eds.), *There’s something about Mary: Essays on Frank Jackson’s knowledge argument against physicalism* (pp. 269–298). Cambridge: MIT Press.
- Cresswell, M. (1985). *Structured meanings: The semantics of propositional attitudes*. Cambridge: MIT Press.
- Crimmins, M. (1992). *Talk about belief*. MIT Press.
- Crimmins, M., & Perry, J. (1989). The prince and the phone booth: reporting puzzling beliefs. *Journal of Philosophy*, 86(12), 685–711.
- Davidson, D. (1975). On saying that. In D. Davidson, & J. Hintikka, (Eds.), *Words and objections: Essays on the work of W.V. Quine* (pp. 158–174). Dordrecht: D. Reidel.
- Dennett, D. (1991). *Consciousness explained*. Boston: Little Brown and Company.
- DeRose, K. (1992). Contextualism and knowledge attributions. *Philosophy and Phenomenological Research*, 52(4), 913–929.
- Forbes, G. (1990). The indispensability of Sinn. *Philosophical Review*, 99(4), 535–563.
- Forbes, G. (1996). Substitutivity and the coherence of quantifying in. *Philosophical Review*, 105(3), 337–372.
- Frege, G. (1966/1892). On sense and reference. In M. Black, & P. Geach (Eds.), *Translations from the philosophical writings of Gottlob Frege* (pp. 56–78). Oxford: Basil Blackwell.
- Frege, G. (1997/1918). Thought. In M. Beaney (Ed.), *The Frege reader* (pp. 325–345). Oxford: Blackwell.
- Green, M. S. (1996). Direct reference and implicature. *Philosophical Studies*, 41(1), 61–90.
- Hellie, B. (2004). Inexpressible truths and the allure of the knowledge argument. In P. Ludlow, Y. Nagasawa, & D. Stoljar (Eds.), *There’s something about Mary: Essays on Frank Jackson’s knowledge argument against physicalism*. Cambridge: MIT Press.
- Horgan, T. (1984). Jackson on physical information and qualia. *The Philosophical Quarterly*, 34(135), 147–152.
- Jackson, F. (1982). Epiphenomenal qualia. *The Philosophical Quarterly*, 32(127), 23–36.
- Jackson, F. (1998a). *From metaphysics to ethics: A defense of conceptual analysis*. New York: Oxford University Press.
- Jackson, F. (1998b). Postscript on qualia. In *Mind, method, and conditionals: Selected essays* (pp. 76–79). Routledge: London.
- Jackson, F. (2004). Why we need a-intensions. *Philosophical Studies*, 118(1–2), 257–277.
- Kant, I. (1965/1781). Critique of pure reason (Norman Kemp Smith, trans.). New York: St. Martin’s Press.
- Kaplan, D. (1969). Quantifying in. In D. Davidson, & J. Hintikka (Eds.), *Words and objections: Essays on the work of W.V. Quine* (pp. 206–242). Dordrecht: D. Reidel.
- Kaplan, D. (1989). Demonstratives: An essay on the semantics, logic, metaphysics, and epistemology of demonstratives and other indexicals. In J. Almog, J. Perry, & H. Wettstein (Eds.), *Themes from Kaplan* (pp. 481–566). Oxford: Oxford University Press.
- Kaplan, D. (1990). Thoughts on demonstratives. In P. Yourgrau (Ed.), *Demonstratives* (pp. 34–49). Oxford: Oxford University Press.
- Kripke, S. (1972). *Naming and necessity*. Cambridge: Harvard University Press.
- Kripke, S. (1988). A puzzle about belief. In N. Salmon, & S. Soames (Eds.), *Propositions and attitudes* (pp. 102–148). Oxford: Oxford University Press.
- Larson, R., & Ludlow, P. (1993). Interpreted logical forms. *Synthese*, 95(3), 305–355.
- Larson, R., & Segal, G. (1995). *Knowledge of meaning: Semantic value and logical form*. Cambridge: MIT Press.
- Lewis, D. (1970). General semantics. *Synthese*, 22(1–2), 18–67.

- Lewis, D. (1979a). Attitudes de dicto and de se. *Philosophical Review*, 88(4), 513–543.
- Lewis, D. (1979b). Scorekeeping in a language game. *Journal of Philosophical Logic*, 8(4), 339–359.
- Loar, B. (1990). Phenomenal states. *Philosophical Perspectives*, 4: *Action Theory and Philosophy of Mind*, 81–108.
- Lycan, W. G. (1990). What is the “subjectivity” of the mental? In J. Tomberlin (Ed.), *Philosophical perspectives IV: Action theory and philosophy of mind* (pp. 109–130). Ridgeview: Atascadero.
- Mates, B. (1952). Synonymity. In L. Linsky (Ed.), *Semantics and the philosophy of language: A collection of readings* (pp. 111–136). Urbana: University of Illinois Press.
- McKay, T. (1981). On proper names in belief ascriptions. *Philosophical Studies*, 39(3), 287–303.
- Moore, G. E. (1903). *Principia ethica*. New York: Cambridge University Press.
- Nagel, T. (1974). What is it like to be a bat? *Philosophical Review*, 83(4), 435–450.
- Perry, J. (1979). The problem of the essential indexical. In N. Salmon, & S. Soames (Eds.), *Propositions and attitudes* (pp. 83–101). Oxford: Oxford University Press.
- Perry, J. (1990). Frege on demonstratives. In P. Yourgrau (Ed.), *Demonstratives* (pp. 50–70). Oxford: Oxford University Press.
- Quine, W. V. O. (1953a). On what there is. In *From a logical point of view: 9 Logico-philosophical essays* (2nd ed., pp. 1–19). Cambridge: Harvard University Press.
- Quine, W. V. O. (1953b). Reference and modality. In *From a logical point of view: 9 Logico-Philosophical essays* (2nd ed., pp. 139–159). Cambridge: Harvard University Press.
- Quine, W. V. O. (1956). Quantifiers and propositional attitudes. *Journal of Philosophy*, 53(5), 177–187.
- Quine, W. V. O. (1960). *Word and object*. Cambridge: MIT Press.
- Recanati, F. (1993). *Direct reference: From language to thought*. Oxford: Blackwell.
- Reimer, M. (1996). What do belief ascribers really mean? A reply to Stephen Schiffer. *Pacific Philosophical Quarterly*, 77(4), 404–423.
- Richard, M. (1988). Direct reference and ascriptions of belief. In N. Salmon, & S. Soames (Eds.), *Propositions and attitudes* (pp. 169–196). Oxford: Oxford University Press.
- Richard, M. (1990). *Propositional attitudes: An essay on thoughts and how we ascribe them*. Cambridge: Cambridge University Press.
- Russell, B. (1905). On denoting. *Mind*, 14(56), 479–493.
- Russell, B. (1985/1918). *The philosophy of logical atomism*. La Salle: Open Court.
- Salmon, N. (1986). *Frege's puzzle*. Cambridge: MIT Press.
- Schiffer, S. (1992). Belief ascription. *Journal of Philosophy*, 89(10), 499–521.
- Soames, S. (1987). Direct reference, propositional attitudes, and semantic content. *Philosophical Topics*, 15(1), 47–87.
- Stalnaker, R. C. (1987). *Inquiry*. Cambridge: MIT Press.
- Stoljar, D., & Nagasawa, Y. (2004). Introduction. In P. Ludlow, Y. Nagasawa, & D. Stoljar (Eds.), *There's something about Mary: Essays on Frank Jackson's knowledge argument against physicalism* (pp. 1–36). Cambridge: MIT Press.
- Unger, P. K. (1984). *Philosophical relativity*. Oxford: Blackwell.
- Wittgenstein, L. (1961/1921). *Tractatus logico-philosophicus* (D.F. Pears and B.F. McGuinness, trans.). London: Routledge & Kegan Paul.
- Wittgenstein, L. (1965/1930). A lecture on ethics. *Philosophical Review*, 74(1), 3–12.
- Zalta, E. N. (1988). *Intensional logic and the metaphysics of intentionality*. Cambridge: MIT Press.