

PHILIP ROBBINS

TO STRUCTURE, OR NOT TO STRUCTURE?

ABSTRACT. Some accounts of mental content represent the objects of belief as structured, using entities that formally resemble the sentences used to express and report attitudes in natural language; others adopt a relatively unstructured approach, typically using sets or functions. Currently popular variants of the latter include classical and neo-classical propositionalism, which represent belief contents as sets of possible worlds and sets of centered possible worlds, respectively; and property self-ascriptionism, which employs sets of possible individuals. I argue against their contemporary proponents that all three views are ineluctably plagued by generation gaps: they either overgenerate beliefs, undergenerate them, or both.

Belief is a relational affair. To be a bit more specific: To have a belief is to stand in an appropriate relation to some sort of semantic object. One major task for a theory of belief, then, is to give an account of these objects. Some philosophers recommend that we treat them as 'unstructured', typically using sets or functions. The idea here is not that unstructured doxastic objects lack structure altogether (whatever that might mean exactly) but rather that they lack anything like the constituent structure characteristic of the syntactic objects analyzed by linguists. Other theorists favor a 'structured' approach, using entities that formally resemble the sentences used to represent beliefs in natural language.

To structure, or not to structure, the objects of belief? That is the question. To address it, we will examine three varieties of unstructured semantics: classical propositionalism, property self-ascriptionism, and neo-classical propositionalism. Though these views don't exhaust the options for the unstructured camp, the failure of all three would constitute fair cause for alarm. My aim is to present new considerations which suggest that, despite persistent claims to the contrary, none of these accounts work: they either overgenerate beliefs, undergenerate them, or both. Interestingly, each generation gap shows up at a metaphysical extreme. Overgeneration affects an agent's beliefs about interworldly states of affairs, maximally objective or impersonal facts like the number of primes or the size of the continuum; undergeneration involves beliefs about intraworldly matters, subjective or personal facts like the time of day. Updated versions of the original unstructured scheme finesse the second problem by incorporating local entities into doxastic objects: substituting individuals, or worlds



Synthese 139: 55–80, 2004.

© 2004 Kluwer Academic Publishers. Printed in the Netherlands.

centered on individuals, for worlds. Yet sets of individuals or centered worlds fare no better than sets of worlds at discriminating among logically equivalent but distinct states of affairs. So even these newer schemes individuate beliefs too coarse-grainedly to capture distinctions that a good general account can't afford to ignore.

1. CLASSICAL PROPOSITIONALISM

Possible-worlds semantics for belief states, or classical propositionalism, has taken a heavy beating over the years, largely due to its propensity to both overgenerate and undergenerate the belief stock of ordinary agents. Nonetheless there are philosophers who continue to swear by it, arguing that the method can be salvaged by means of a technical device called diagonalization.¹ This latter claim is my main target in this area. Diagonalization, I will argue, cannot save classical propositionalism from either of the generation gaps which seem to plague it.

Here's the basic setup. On the object side of things classical propositionalism (CP) defines propositions as functions from possible worlds into truth-values. At a minimum, such a function maps each world in its domain in which the proposition holds to 1 (the True), and each world in which its negation holds to 0 (the False). Different sorts of proposition emerge when the formal character of these functions is further specified. For example, we can think of these functions as domain-invariant. In that case they all share a single, context-independently fixed domain: the universal set of worlds, W . Or we may think of propositions as functions definable on a plurality of domains, whose membership is determined by which worlds are cognitively salient. These two modes of specifying propositions-as-functions yield two variations on the unstructured theme: (V1) total functions on the universal domain of possible worlds, and (V2) total functions on context-dependent domains of worlds. V1 reduces to the streamlined propositions-as-sets conception with which classical propositionalism is ordinarily associated. V2 disallows this reduction of propositions to sets, since two propositions can agree on the worlds they map to 1 while their overall domains differ.

Classical propositionalism formally characterizes the belief relation in terms of set containment.² An agent is correctly reported as believing that P just in case P holds in every world which is doxastically accessible to her, where doxastic accessibility is defined as follows: given a pair of worlds (w_1, w_2) and an agent A at w_1 , w_2 is doxastically accessible to A from w_1 just in case for all A believes at w_1 , she might as well be at w_2 . So, for instance, if A has no beliefs whatsoever about panda bears and w_2 is

accessible, the panda situation at w_2 can be any which way, so long as it is compatible with what A believes about other matters. But if A 's beliefs are inconsistent – that is, if there is no world in which everything she believes is true – then neither w_2 nor any other world is accessible to A .

Classical propositionalism individuates doxastic objects and relations alike in a fairly coarse-grained way. Objects of belief are classified by sets of worlds, the belief relation by set containment. That ensures the stability, or inter- and intrapersonal shareability, of beliefs, since sets of worlds are robustly objective, mind-independent entities. But it also opens up generation gaps, in stereo.

2. OVERGENERATION

Overgeneration stems from the problem of logical omniscience. It's easy to see how this problem arises, given the classicalist account of belief. Let P be a proposition which A believes, and Q any of its deductive consequences. Since A believes that P , every world in her belief-set – the set of worlds doxastically accessible to her – is a P -world. And since P entails Q , every P -world is a Q -world. So every world in A 's belief-set is a Q -world. Hence A believes that Q . For instance, once you believe the axioms of Peano arithmetic, you believe all the theorems too. In fact, there's a sense in which you can't help but believe all necessary truths, arithmetical or otherwise, straight off the bat. For no matter what its composition, your belief-set is a subset of the universal set of worlds W , viz., the necessary proposition, \top . So you believe \top . Thus, if s is a sentence which expresses \top , every claim to the effect that you believe that s will come out true. Since all necessarily true sentences express \top , you believe all necessary truths.

A related difficulty arises with respect to agents whose beliefs form a logically inconsistent set (which is probably most of us). The belief-set of an inconsistent agent is empty, so it is contained by all sets of worlds. And an agent believes all the propositions represented by supersets of her belief-set. So an inconsistent agent believes absolutely everything. Doxastic logics which assume that agents are consistent in their beliefs go at best from fire to frying pan, since the assumption at issue is no more plausible than that of omniscience.

These are unlovely results. In response, friends of classical propositionalism have devised various remedies. Most of these rely in some form on one of four strategies:

1. *Non-Normalization*, which trades on the notion of an impossible world.

2. *Bifurcation*, which introduces a distinction between explicit and implicit belief.
3. *Idealization*, which defends classical propositionalism as a strongly idealizing theory of belief.
4. *Diagonalization*, which exploits the idea of metarepresentational content.

I will focus on the relatively little-known diagonalization strategy, after lightly sketching each of the other three.

2.1. *The Non-Normalization Reply*

This strategy, due to Hintikka (1975), drops the condition that, given an agent A at a world w , all worlds doxastically accessible to A from w are logically possible. To see how this works, assume A believes P . Then A 's belief-set is a subset of P , as before. But if we extend logical space to include logically impossible worlds, the fact that P entails Q no longer means that every P -world is a Q -world; it only means that every logically possible P -world is. So it doesn't follow that A 's belief-set is contained in Q . Hence it doesn't follow that A believes Q . So A 's beliefs are not closed under deduction.

This solves the problem of logical omniscience, but at considerable cost. Two points here. First, the notion of an impossible world is less clear and intuitive than that of a possible world.³ Second, as Vardi (1986) points out, even if such a nonstandard theory could be devised, that would only displace the problem of logical omniscience, not solve it. For non-normalization leaves an agent's beliefs closed under deductive consequence. The difference is that, instead of assuming agents are omniscient with respect to classical logic, we now assume they are omniscient with respect to some deviant alternative, such as relevance logic. Relevance logic helps especially with the subproblem of inconsistency, since it drops the principle of *ex falso quodlibet* (from a contradiction anything follows); relevance doxastic logic, in turn, drops the assumption that agents whose beliefs are inconsistent believe everything. That advances the cause somewhat, to be sure. But the deeper problem of omniscience isn't that ordinary mortals fail to reason perfectly in classical logic. The deeper problem is that they're not very good at deductive reasoning of *any* kind, classical or otherwise.

2.2. *The Bifurcation Reply*

The present strategy turns on a proposed distinction between *explicit* and *implicit* belief.⁴ On this view, an agent's explicit beliefs constitute a rela-

tively small finite set; her implicit beliefs, infinitely numerous, are all and only the logical consequences of her explicit ones. Believing a proposition explicitly entails believing it implicitly, but not vice versa. Implicit belief is closed under deduction, whereas explicit belief is not. So with respect to explicit belief, overgeneration does not occur.

The main difficulty with this take on the omniscience problem is its lack of psychological realism. There doesn't seem to be any serviceable notion of belief according to which ordinary agents believe all the truths of arithmetic, or any other noncontingent subject-matter. The reason for this is not far to seek. As various authors have emphasized, it is of the essence of belief that beliefs are representational states which play a role in the guidance of action.⁵ In Ramsey's celebrated phrase, beliefs are "maps by means of which we steer" through our environment. Having set action-guiding potential as a necessary condition on beliefhood, it follows that Levesque-style implicit beliefs are beliefs as decoy ducks are ducks – that is, in name only. For assuming that an agent implicitly believes all the truths of arithmetic, the vast majority of those beliefs will never show up in behavior. That gives us a good reason to be skeptical of their existence.⁶

2.3. *The Idealization Reply*

The idealization strategy, developed in Stalnaker (1991), grants that classical propositionalism is not a descriptively adequate theory of the attitudes – at least not in any straightforward sense. It then seeks to deflect the impression of inadequacy by arguing that classicalism must be understood as a strongly idealizing theory, and can be fairly evaluated only in terms of that tendency. An analogy to theorizing in another domain will help here. Just as classical mechanics abstracts away from complications due to friction and air resistance, classical propositionalism abstracts away from the "cognitive friction" which spoils deductive closure. In both cases, simplifying assumptions are made to smooth theory construction. The result is a more synoptic, though grainier, picture of the phenomena under study.

This is a natural strategy for coping with the omniscience problem. But given a plausible constraint on the sort of idealizing appropriate for psychology, an idealization that floats as high above doxastic reality as the classicalist one will be hard to justify.

Theory building typically involves a trade-off between tractability and applicability. The trick is to find the right balance between these desiderata, not sacrificing one too much in favor of the other. No one doubts that classical propositionalism yields a formally elegant theory of belief. The worry is that it thereby gives up too much on the side of applicability.

That's against the background of the following candidate constraint on idealization:

Minimal Applicability Constraint. An appropriate idealization must offer broad explanatory and predictive coverage of non-breakdown cases in its proper domain.

Putting the point negatively, an idealization which fails to apply in at least a broad range of intuitively normal cases cannot be considered theoretically adequate, for it sacrifices too many phenomena. The Minimal Applicability Constraint is backed up by the fact that, were it not to be met, the theory in question might be vacuously confirmed, and would have no explanatory or predictive power to speak of.

Classical propositionalism, by idealizing to deductive closure and consistency, does not meet this constraint. After all, the classical theory applies only to agents for whom the deductive sciences are entirely trivial, namely, no one.⁷ Surely such an extreme idealization puts a wealth of psychological phenomena out of epistemic reach. For example, it makes it impossible to explain or predict at the intentional level what seem like undeniable differences in behavior between agents who know something about topology and agents who don't. Ignorance of mathematics may be unfortunate, but it's scarcely a sign of cognitive breakdown.

2.4. *The Diagonalization Reply*

Finally, we can try to 'diagonalize' the problem away, as proposed in Stalnaker (1984). To see how this works, consider the sentence

- (1) Lewis Carroll is Charles Dodgson.

Tokens of this sentence, it is suggested, are ambiguous in content. Sometimes they express the proposition that a certain individual, call it x , is identical with itself. This proposition, call it P , is the set of worlds in which x is self-identical. Every individual is necessarily self-identical, so $P = \top$. That's a proposition we can't doxastically refuse.

But sometimes tokens of (1) express a different proposition. The alternate content P' of (1) is the metalinguistic proposition given by

- (2) The proposition expressed by (1) is true.

We call P' the *diagonal proposition*, because it represents the diagonal of a i -by- j matrix of truth-values whose i th row, j th column entry is the truth-value of the proposition expressed by the sentence at the i th context

of tokening evaluated at the j th context of evaluation. The matrix as a whole represents the *propositional concept* of (1). Its rows correspond to the *horizontal propositions* expressible by (1), which vary according to what world the sentence is tokened at.

Sentences like (1) can express different horizontal contents in different contexts of tokening because linguistic items such as names and predicates do not have their semantic properties essentially. So there are worlds in which either the names ‘Lewis Carroll’ and ‘Charles Dodgson’ do not co-refer or ‘is’ does not mean identity.⁸ This metaphysical fact articulates the horizontal/diagonal divide. Compare: the horizontal content of (1) as uttered at the actual world is just the set of worlds in which x is self-identical, namely, \top ; the diagonal content is the set of worlds in which the referents of ‘Lewis Carroll’ and ‘Charles Dodgson’, respectively, jointly instantiate the dyadic relation expressed by ‘identity’. While the sentence (1) determines a large class of horizontal contents, some of which are necessary, the diagonal content it determines is unique. The latter content is also contingent, for while actual tokens of (1) mean that Lewis Carroll and Charles Dodgson are the same person, it isn’t necessary that the sentence mean that. The same string of symbols could have meant that Lewis Carroll and Charles Dodgson are different people, or that Nixon was not a crook. (Note that the interpretation of the identity predicate is subject to transworld variation just as it is with any other sentential element.) That is, it’s possible that (1) says something that’s actually false.

Mathematical cases can be handled in a similar fashion. The horizontal content of actual tokens of

(3) For $n > 2$, $x^n + y^n = z^n$ has no integral solutions

is, again, the universal set of worlds, \top . We believed this proposition all along, since Fermat’s Last Theorem is true, mathematical truths are necessary, and necessary truths are doxastically irresistible. But until recently most of us didn’t believe the diagonal content of (3), that is, the contingent proposition that

(4) The proposition expressed by (3) is true.

As before, (4) isn’t necessary because there are possible worlds in which (3) means something false, like the negation of the theorem. The upshot is that, by diagonalizing, classical propositionalists seem able to deflect the criticism that agents don’t believe all of mathematics or any other noncontingent subject-matter.

Alas, this solution to the problem of logical omniscience isn’t much better off than the other three. Here are two reasons to reject it.

First, it's intuitively weak. Only a theorist wedded to her theory would claim that what we came to believe after Wiles's proof was not the mathematical fact that $x^n + y^n = z^n$ lacks integral solutions for $n > 2$, but the metalinguistic fact that the sentence " $x^n + y^n = z^n$ lacks integral solutions for $n > 2$ " expresses a truth. Common sense suggests the contrary, namely, that we came to believe a bit of pure mathematics. This intuition seems fairly robust: though it might be discounted or explained away, it should not be dismissed out of hand. Likewise, *mutatis mutandis*, for identities between proper names and other cases of noncontingent truth. Call this the *Topicality Objection*.

Second, diagonalization undermines natural strands of doxastic continuity across linguistic communities. Here's why. The strategy dictates that the principal objects of belief and doubt in a noncontingent domain like mathematics are propositions about the relation between sentences of mathematical language, familiarly known as 'mathematese', and what those sentences express, rather than propositions about mathematical structures and their properties. But mathematese is not a universal canonical notation. Nor is it a fragment of any one natural language, since the mathematese written and spoken in Berlin, say, differs from Chicago mathematese. This means that a number theorist in Berlin trying to determine the truth-value of Goldbach's Conjecture is actually working on a different problem than her counterpart in Chicago. The Chicagoan is considering whether or not the sentence

- (5) Every even number ≥ 4 can be expressed as the sum of two primes

expresses a truth. The Berliner is asking this about a different sentence, namely

- (6) Jede gerade Zahl ≥ 4 lässt sich als Summe zweier Primzahlen darstellen.

These are different questions, since presumably there are worlds in which (5) and (6) express different propositions; at least, it's not clear why there shouldn't be such worlds. So if a genuine proof of the conjecture were found and published, the Chicagoan would come to believe one thing and the Berliner another. But this is implausible. We've no reason to fear that the mathematical community suffers from the sort of doxastic discontinuity that diagonalization implies. Call this the *Continuity Objection*.

This pair of objections to the Diagonalization Reply should give the classicalist pause. But it would be unfair to close a discussion of this approach to the omniscience problem before considering how an advocate of diagonalization might respond. So let us do just that, taking up each objection in turn.

The Topicality Objection. This objection charges the diagonalizer with mischaracterizing the topic, or subject matter, of our mathematical (and other noncontingent) beliefs. More specifically, the objection says that diagonalization has the effect of saddling such beliefs with an essential metalinguistic component which seems alien to what they are rightly about, viz., noncontingent states of affairs. To block this move, Stalnaker (1984) appeals to the general motivation behind modelling belief contents using sets of possible worlds, what Stalnaker calls the “pragmatic picture” of mental representation. According to this (broadly functionalist) picture, the analysis of attitudes is governed first and foremost by the role such states play in the explanation of rational action, with possible worlds representing the alternative possible outcomes of actions among which agents choose to actualize by acting. Transposed to this theoretical setting, the claim that mathematical thought concerns mathematase expressions and their properties gains plausibility, for

[i]f one looks at the kind of actions that might be explained by mathematical beliefs, and at the abilities that constitute mathematical knowledge, one finds that they are actions and abilities that essentially involve operating with some kind of notation – for example, calculating and constructing proofs. So it is appropriate, according to the pragmatic picture, which in the general case motivates a separation of propositions from linguistic expressions and semantic structure, to treat mathematical propositions as essentially involving expressions and semantic structure by having them as [their] subject matter. (Stalnaker 1984, 74–75)

Seen from a pragmatic perspective, the metalinguistic slant on mathematical belief is, if not natural, then at least “not entirely artificial.”⁹

By way of counterreply, let me briefly make three points.

First, the pragmatic defense fails to meet the Topicality Objection head-on, and is itself vulnerable on pragmatic grounds. To see this, consider the question: What is mathematics about? A plausible answer along realist/structuralist lines goes as follows:

Structure Principle. Mathematics is the study of mathematical structures: things like number systems, lattices, groups, and graphs.

Though I haven’t space to defend this claim here, philosophical reflection on mathematical practice gives fairly robust *prima facie* support for

the Structure Principle. One of the most striking features of mathematics, for example, is its applicability to the physical world – a fact which the principle goes some distance toward explicating, given that mathematical structures can be physically instantiated. This principle in turn supports the idea that mathematical knowledge consists in knowledge of the properties of such structures, knowledge which is typically acquired by manipulating symbols.

Contrast this idea with the Stalnakerian view. According to diagonalization, mathematical knowledge is not knowledge of mathematical structures, as the Structure Principle implies, but rather knowledge of mathematese expressions and their semantic properties. The sense in which knowledge of the latter sort counts as knowledge of mathematics seems attenuated at best. Elevating it to primary status on the basis of pragmatic considerations will be no small feat, given the predominantly object-oriented spirit in which mathematics is conducted.

Second, while pragmatic considerations indicate the importance of language to mathematical thought, they also point to the need to distinguish between knowing what a mathematical sentence means and knowing whether it's true. Fermat's Last Theorem illustrates this phenomenon nicely, since grasping what (3) says is (seems?) so easy but settling its truth-value is so hard. Diagonalization makes no room for this distinction: from the diagonal point of view, if you don't know a theorem, it can only be because you don't understand what the theorem says, viz., which proposition it expresses at your world. In short, ignorance of mathematics is, at root, ignorance of mathematical meaning. This is surely an odd result, but the pragmatic defense of diagonalization sheds no light on the issue.

Third, as Stalnaker concedes, the pragmatic defense does not obviously generalize beyond the mathematical case. For example, it fails to cover noncontingent beliefs about the identity of individuals (Lewis Carroll and Charles Dodgson) or natural kinds (gnus and wildebeests), since the actions and abilities connected with such beliefs needn't have a constitutive connection with uses of public language. In these cases, the proponent of the Diagonalization Reply has to bite the bullet and insist, arguably against appearances, on the metalinguistic character of such beliefs.

The Continuity Objection. The complaint here is that assigning diagonal contents makes it unlikely that continuity of mathematical belief will be preserved across linguistic divisions. This follows from the fact that mathematese sentences in distinct languages which express the same proposition at our world could express distinct propositions at some other world. Our example featured two number theorists, one American and the other German: the American being assigned the diagonal content of an

English-language formulation of Goldbach's Conjecture, and the German, the diagonal of its German-language counterpart (sentences (5) and (6), respectively). It was then suggested that, since the diagonal proposition expressed by (5) need not coincide with the diagonal proposition expressed by (6), the American and the German would be working on different problems – a counterintuitive result.

A natural rejoinder is to insist that intertranslatable sentences belonging to distinct sectors of mathematese have the same diagonal content. The trick is finding some way to justify this insistence. One proposal involves cutting down the space of possible worlds by appealing to facts about cognitive salience, where a possible world is salient at a context if compatible with the presuppositions of – hence, doxastically accessible to – the agent(s) of the context. This means shifting to a different version of classical propositionalism than the one under discussion until now. Recall (from Section 1) how we distinguished two versions of this doctrine, corresponding to two ways of modelling propositions as functions from worlds into truth-values. So far we've been working with V1, on which those functions are total and defined on a single, context-invariant, universal domain of worlds. The present effort to meet the Continuity Objection shifts the action to V2. On this refinement of classicalism, propositions are total functions on domains which vary from context to context according to which worlds are salient. In other words, V2-propositions are total functions from context sets into truth-values. The defense then claims that only worlds in which (5) and (6) agree on the diagonal will be salient to competent speakers of mathematese. By restricting propositional construction to salient parts of logical space, as per V2, we exclude worlds in which (5) and (6) determine different horizontal contents. This strategy should ensure continuity of mathematical belief across the usual linguistic lines.

There are, however, two serious problems with this response to the Continuity Objection. The first relates to the specific use made here of context-dependent propositional domains. The second is broader in scope, as well as deeper in philosophical significance. It bears on the general prospects for versions of classicalism which exploit such domains, such as V2. I will elaborate each of these problems in turn.

The first difficulty is obvious. Let's say that Fred (in St. Louis) and Friedrich (in Berlin) are monolingual speakers of English and German, respectively. Both are number theorists working on Goldbach's Conjecture, intuitively expressible in English via (5) and in German via (6). The current defense of diagonalization grants that there may be worlds in which (5) and (6) express different horizontal propositions of the sort favored by V1, so that the V1-diagonals of these sentences need not coincide. What

it denies is that the context sets for either Fred or Friedrich contain any such worlds, and hence that the V2-diagonals of (5) and (6) could diverge. But why should this be so, if Fred and Friedrich are monolingual? Surely a monolingual speaker of English or German will presuppose little or nothing about the status, synonymy-wise, of this (or any other) pair of English and German sentences.

The second point concerns the compatibility of V2 with the intuitive – and arguably non-negotiable – demand that thought contents be stable. Since agents invariably diverge in at least some of their estimates of live optionhood, context sets will vary from agent to agent. Only in cases where all presuppositions are shared will all the same worlds be salient, and the context sets coincide. Such cases will be rare, if indeed they ever arise. The moral is clear. Propositional domains can be relativized to context only at the cost of violating the stability constraint on content. And that may be a cost that no adequate theory of content can bear.

After the dust settles, then, it seems that diagonalization offers only marginally more help with the omniscience problem than non-normalization, bifurcation, or idealization. Despite all these efforts, the overgeneration gap persists. This is bad news for classical propositionalism. And there's more bad news to come.

3. UNDERGENERATION

According to one formulation of classical propositionalism, due to Lewis (1979), belief amounts to a kind of cognitive self-location, namely, self-location at the global level, the level of worlds. Worlds are points in logical space, propositions regions of that space. As a first-pass approximation, to believe that P is to partition logical space into two regions, the region where P holds and its complement, and to assign one's world to the first region.

This is only approximate because a defining feature of the Lewisian analysis is that all beliefs, regardless of their propositional content, involve a first-person or self-conceptual component. The reason for this is that self-location is an intensional notion, not an extensional one. It requires more of an individual x than thinking about x in one way or another; rather, x must be thinking about herself *as herself*, or 'self-consciously' – whatever that comes to.¹⁰ To highlight this aspect of the analysis, I shall adopt the pronominal '*'-notation introduced in Castañeda (1966), which indicates that the referent of a superscripted term is thought of by the agent in a self-conscious manner. Lewis's general line on classical propositionalism, then, amounts to this: An agent A believes that P just in case A locates herself*

at a *P*-world, viz., in the *P*-region of logical space. So, for instance, if *A* believes that panda bears are herbivorous, then she locates herself* at a world in which pandas are herbivorous.

But agents also occupy positions *within* worlds. As a subject of attitudes at a world *w*, *A* is a particular individual *x* at a particular place and time. And she can think of herself as such. For example, she can locate herself* at the position in the space of individuals which is occupied by Ingrid Bergman – a belief she might express by uttering the sentence

(7) I am Ingrid Bergman.

Similarly, she can locate herself* in ordinary space, with beliefs such as that expressed by

(8) I am in St. Louis.¹¹

In such cases the type of self-location involved is local, not global: the agent is fixing her* intraworldly locus, rather than her* location in logical space.

We can think about the difference between these two kinds of self-locational content in terms of a difference in representational granularity, along the lines of Hobbs (1985). Granularity is like degree of resolution, or level of detail. For example, a map of the Midwest which uses a network of dots and lines to indicate St. Louis and other urban areas, and the roads connecting them, pictures the geography of the region in a coarser-grained way than a map which displays the grid of streets in Clayton and adjoining districts. The first map may be all that's needed to get to St. Louis, but only a representation of finer granularity will help us navigate within city limits. Similarly, a coarse-grained cognitive map of logical space which displays worlds as structureless points will enable global but not local self-location.

The idea of a built-in granularity limit to CP underlies the problem of self-location as presented in Lewis (1979), and the undergeneration gap to which this problem gives rise. The problem is that omniscience (true belief) with respect to propositions, conceived of as sets of possible worlds, seems compatible with local self-locational ignorance. If so, classical propositionalism is in trouble. Here is the argument:

- (P1) If CP is true, then omniscience with respect to classical propositions is equivalent to omniscience *simpliciter*.
- (P2) It is possible for an agent to know, more or less exactly, both her* global and local location.

- (P3) Any agent who knows her* exact global location, i.e., knows exactly which world she* inhabits, knows all propositions true at that location.
- (P4) It is possible for an agent to know her* exact global location yet not know her* approximate local location.
- (P5) It is possible for an agent to know all true propositions yet not know her* approximate local location. [by P3 and P4]

Hence,

- (C) CP is false. [by P1 and P5]

The argument is valid. *Prima facie*, the more interesting premises are P3 and P4.

P3 is fairly intuitive, but were an argument wanted, the following *reductio* should serve. First, assume that *A* knows at *w* exactly which world is hers*. Assume also that there is a proposition *P*, true at *w*, which *A* doesn't know. Then *A* could not discriminate her world *w*, in which *P* holds, from an otherwise indiscernible (hence, accessible) world *w'* in which *P* fails. But recall that an agent *A* knows (believes) that *P* at *w* only if there is no $\neg P$ -world *w'* epistemically (doxastically) indiscernible to *A* from *w*.¹² So this contradicts the hypothesis that *A* has exact global self-locational knowledge. Hence P3 is true.

P4 is more controversial. Lewis motivates it with a pair of fictional cases, my favorite of which is his "Tale of Two Gods" (TTG, for short). Of these hapless immortals Lewis writes:

They inhabit a certain possible world, and they know exactly which world it is. Therefore they know every proposition that is true at their world. Insofar as knowledge is a propositional attitude, they are omniscient. Still I can imagine them to suffer ignorance: neither one knows which of the two he is. They are not exactly alike. One lives on top of the tallest mountain and throws down manna; the other lives on top of the coldest mountain and throws down thunderbolts. Neither one knows whether he lives on the tallest mountain or on the coldest mountain; nor whether he throws manna or thunderbolts. (Lewis 1979, 520–521)

This scenario surely looks coherent. It also rallies some strong intuitions in favor of P4. What is a classical propositionalist to do?

I propose to examine four possible solutions to the problem of self-location. These are (1) the Haecceity Reply, (2) the Diagonalization Reply, (3) the Introspection Reply, and (4) the Vacuity Reply. I'll argue that none of them pans out.

3.1. *The Haecceity Reply*

This proposal, anticipated in Lewis (1979) and pressed again in Stalnaker (1981), appeals to a distinction between two kinds of propositions about individuals. First there are *qualitative* propositions: propositions about what properties or qualities individuals have. Then there are *nonqualitative* propositions: propositions about which individuals are which, that is, propositions about the distribution of individual essences or ‘haecceities’ (literally, thisnesses). The distinction comes into play as follows:

Consider a possible world w' which is exactly like the actual world w except for the fact that the gods have traded places and properties. At w' the god on the tallest mountain (TM) is on the coldest mountain, with all the properties that the god on the coldest mountain (CM) has at w , and vice versa.

The worlds w and w' , though distinct, agree on all qualitative features. If TM is locally self-locationally ignorant, it can only be because he doesn’t know which of these qualitatively indiscernible worlds is actual. Hence he must be missing some nonqualitative propositions. Likewise for CM. But by hypothesis the gods are omniscient about all propositions. So the story doesn’t add up.

However, as Lewis (1979) points out, this reply underestimates the depth of the problem. For let TM – identifiable across worlds as the individual x – be omniscient with respect to all facts, qualitative or otherwise. Then TM knows that w , and not w' , is actual. Since it is a fact of w that TM is on the tallest mountain, TM knows that x is there. But it does not follow that TM knows that he* is perched there, unless he also knows that he* is x . And that extra bit of knowledge is something that the classicalist has yet to account for.

3.2. *The Diagonalization Reply*

The next move, due to Stalnaker (1981), seeks to undermine the coherence of TTG by appealing to divine knowledge of diagonals.

According to the story, TM knows all propositions but he doesn’t know whether he* is on the tallest or the coldest mountain. However, consider the diagonal content of TM’s (actual or possible) utterances of

- (9) I am on the tallest mountain.

Call this proposition *D*. If TM knows *D*, then he knows that he* is on the tallest mountain. And since TM is omniscient with respect to propositions, he knows *D*. Contradiction!

This objection can be readily met. Granted, since TM knows all propositions, he does indeed know *D*. But his knowing *D* does *not* entail his knowing that he* is on the tallest mountain. TM knows precisely which individual at *w* his utterances of (9) refer to (namely, *x*) and what properties that individual has. But as O'Brien (1994) points out, he needn't have a clue as to which utterances of (9) are his*. Unlike ordinary agents, TM is estranged from his occurrent linguistic productions, unable to recognize them as his* or himself* as their author. He's mired in local self-locational ignorance as a result.

Here's a less fanciful case which exposes the flaw in the Diagonalization Reply very clearly. It's summertime, and you're standing in the corridor of a warehouse just outside a room marked COLD STORAGE. The corridor is tropically hot and humid, and you're sweating bullets. As you raise your hand to mop your brow, you hear a still small voice from inside the room say

(10) Baby, it's cold in here.

Now, under the circumstances it makes sense for you to think that whatever the speaker said is true. And you do think this. So you believe the diagonal content of the speaker's token of (10). But believing this content doesn't mean believing it's cold where you* are standing. What you believe is that it's cold at the context of the speaker; that is, inside the storeroom, not in the corridor outside. You'd believe it was cold where you* were if you believed both the diagonal of the relevant token of (10) *plus*

(11) that you* were the agent of the context of (10)'s tokening.

But you don't believe (11).

Analogously, TM knows the diagonal of his tokens of (9), and he knows that *x* is the agent of the context of (9)'s tokening. But this does not entail his knowing that he* is that agent. So he need not know where he* is. The apparent contradiction vanishes.

I can imagine someone sympathetic with the Diagonalization Reply trying to rescue it as follows:

By hypothesis, TM has the capacity to utter (9). So TM has the capacity to wield sentences containing the first-person pronoun. But genuine

competence with the first-person pronoun rules out the sort of cognitive estrangement embodied by TM's (supposed) failure to recognize his occurrent utterances as his*. In short, only those who can identify their occurrent tokens of 'I' as their own count as competent with 'I'.¹³ Yet TM, *qua* utterer of (9), has this competence. So he cannot be ignorant in the manner suggested.

We can make short work of the suggestion above by noting that, even if competence with the first-person pronoun does necessitate the ability to recognize one's occurrent utterances as one's own (which is not beyond dispute), Lewis's story nowhere assumes that the gods are competent in this way. Nor can the objector simply assume this to be the case. What's needed is an argument to the effect that the capacity for global self-location necessitates this competence, and it's far from clear how such an argument might go – especially since the very idea of such a competence is *prima facie* at odds with the view that all knowledge is classical.

3.3. *The Introspection Reply*

Even after the weaknesses in the Diagonalization Reply are pointed out, one might well persist in thinking that there is a tension between global self-locational knowledge and local self-locational ignorance. One might, for instance, reason as follows:

Though the gods may not be able to employ the first-person pronoun, they are able to think. And being able to think necessitates being able to introspect one's occurrent thoughts. Further, if one can introspect one's occurrent thoughts, then one can identify those thoughts as one's own, that is, self-consciously. Putting these points together we get this result: If you can think, then you can recognize your occurrent thoughts as yours*.

So, suppose TM has an occurrent thought *t*. Since TM is a genuine thinker, he can introspect *t*; and since he can introspect *t*, he can recognize *t* as his*. Being omniscient with respect to classical propositions, TM also knows that *x* is the unique host of *t*, and that *x* is on the tallest mountain. So he can figure out that he* is on the tallest mountain.

Though this line has some face plausibility, it is vulnerable on much the same grounds as the Diagonalization Reply. The reasoning depends on two theses, namely:

- (T1) If *S* is a genuine subject of thought, then *S* can introspect her occurrent thoughts.

and

- (T2) If *S* can introspect her occurrent thoughts, then *S* can identify her occurrent thoughts as hers*.

From T1 and T2 we get

- (T3) If *S* is a genuine subject of thought, then *S* can identify her occurrent thoughts as hers*.

If T3 is true, then the story of the gods doesn't add up. So if we grant T1 and T2, that should take care of the problem.

But neither T1 nor T2 is self-evident. These are substantive claims, and they need to be argued for. What's more, they need to be argued for in a manner consistent with the principle that all knowledge is knowledge of classical propositions – including, say, introspective knowledge of the ownership of one's occurrent thoughts. In other words, it's not enough to show that T3 is true; one has to show both that T3 is true *and* that its truth is compatible with the truth of classical propositionalism. I don't imagine this will be easy to do. To see why, note that the closest classical approximation to knowing that a thought *t* belongs to you*, where you are the individual *y*, is knowing that you* inhabit a world in which *t* belongs to *y*, viz., knowing that *t* belongs to *y*. But what's to stop you from knowing that *t* belongs to *y* *without* knowing that *t* belongs to you*?

3.4. *The Vacuity Reply*

Here is one last attempt to bridge the undergeneration gap:

According to TTG, the gods know which world is theirs* (namely, *w*); so they must be self-conscious. But what does their self-consciousness consist in? In the normal human case, self-consciousness is linked to epistemic capacities like introspection and proprioception – capacities which reliably yield local self-locating knowledge. Thus, if the gods' self-consciousness were anything like ours, then each would know his* location within *w*. By assumption, the consequent of this conditional is false. But if the gods' self-consciousness is nothing like ours, then it's a mystery what it *is* like. So how can intuition be brought to bear in deciding whether the scenario described is a possible one?

There is something to this complaint. The notion of divine self-consciousness is obscure, and this obscurity makes it hard to get a grip on the

gods' situation. Quite right. But we don't need to get much of a grip here in order to appreciate the point of Lewis's story.

Let's concede that, if TM were self-conscious in more or less the way we are, he would know that he* was on the tallest mountain. This seems plausible enough, given a certain intuitive understanding of what ordinary self-consciousness involves. The point remains that TM *couldn't* know that he* was on the tallest mountain if all he knew was classical propositions. The most that classical omniscience affords TM is the knowledge that he* is at a world in which x is on the tallest mountain (or, more simply, the knowledge that x is on the tallest mountain). That's the intuition that counts, and the Vacuity Reply does nothing to undercut it.

4. PROPERTIES AND CENTERED WORLDS

A near neighbor to classical propositionalism, also in the unstructured camp, is property self-ascriptionism (PSA). The property view is developed in Lewis (1979, 1986). It diverges from all varieties of propositionalism by substituting non-truth-value-bearing abstract entities – namely, properties – for propositions as the primary objects of belief. But its distinctive features are best viewed against the background of the classical model which inspired it. As we shall see, shifting to properties eases the generational malaise somewhat, yet the gain is not enough to vindicate the unstructured approach to belief. The same goes for the (equivalent) proposal to model doxastic objects and relations using sets of centered worlds.¹⁴ Though safe from undergeneration, both accounts fatally overgenerate beliefs.

We begin with fundamentals, starting with a background check on the new objects. We've been working so far with a conception of propositions as total functions from the universal set of worlds W into the pair set of truth-values $\{0, 1\}$, or more simply as sets of worlds. Properties can be thought of along similar lines. A (first-order) property F can be represented as a total function from W into $\wp I$, the power set of the universal set of individuals, I . Equivalently, a property is a family of sets of individuals, namely, the collection of its extensions at each world. We can further simplify this conception by representing a property as the union of its extensions, that is, as the set of its instances. On this view a property reduces to a set of possible individuals, some of which may be actual – just as a classical proposition reduces to a set of possible worlds, some one of which may be actual.

The belief relation is also defined in a manner analogous to the classical. Where before we had the relation of doxastic accessibility on worlds, we

now have doxastic alternativity on individuals. Of an individual i and an agent A at w , we say that i is doxastically alternative to A at w just in case for all A believes at w , she might as well be i . Doxastic inconsistency arises exactly where the properties an agent self-ascribes are not jointly coinstantiable, in which case A has no doxastic alternatives.

It's not immediately obvious, however, how to cash out belief in terms of properties. Judging from the surface syntax of belief reports, belief looks to be a relation between an agent and a truth-value-bearing entity – hence the standard view of beliefs as attitudes toward propositions (of one sort or another). Yet the classical view of belief as an attitude toward propositions, taken as sets of worlds, naturally maps onto the view of belief as an attitude toward properties, taken as sets of individuals. To see how, it helps first to distinguish between two species of first-order property, the global and the local. These terms apply to properties in a manner roughly parallel to their earlier application to locations. A global property is one which is possessed by all the individuals at any world in which it has instances. More precisely, a property F is global just in case for any world w and any individual x at w , x has F only if all of x 's worldmates have F . A local property, on the other hand, may be possessed by some but not all of the individuals at a world. That is, a property F is local just in case there is some world w and some individual x at w such that x has F and some of x 's worldmates lack F .

We've already glimpsed the move from possible-worlds to possible-individuals theories of belief. Setting out the problem of self-location, we noted that to believe a classical proposition P is in effect to mentally situate oneself* in P -space, the region of *world-space* whose every point is a P -world. The same idea can be recast in property-theoretic terms as follows. To believe that P is to self-ascribe a certain global property F , the property of inhabiting a P -world. Self-ascribing F is, in turn, equivalent to mentally situating oneself* in F -space, the region of *individual-space* whose every point instantiates F . For instance, to believe that pandas are herbivorous is just to self-ascribe the property of inhabiting a herbivorous-panda world, or to locate oneself* in the subspace of individuals who inhabit such worlds.

To highlight the proximity of the property view to classical propositionalism, it's worth noting that the former is equivalent to a third view which represents a sophisticated variant of the latter. Neoclassical propositionalism (NCP) departs from the classical tradition by modeling doxastic objects with sets of centered worlds in lieu of ordinary worlds. A centered world, or c -world, is an unordered pair (w, i) consisting of a world w and a designated individual i at w . A centered proposition, or c -proposition, is a set of c -worlds. Centered world-space has c -worlds as

points and c -propositions as regions. If we index the density of a finite space S by the number of regions (propositions) S contains, centering stands to increase the density of S dramatically. This boost in density offers the advantage of greater representational granularity, since more propositional “pixels” per unit of logical space means enhanced resolution in the picture. NCP is otherwise exactly like CP, *modulo* the substitution of c -worlds for worlds in the definition of the belief relation.

Since classical propositions correspond to global properties, and not all properties are global, it follows that there are properties to which no classical proposition corresponds. But due to the effect of centering just noted, the same cannot be said of neoclassical propositions. Let us say that an individual x centrally inhabits a c -world (w, i) just in case x is at w and $x = i$, and consider any property G , global or local. There is a unique c -proposition which corresponds to G , namely, the set of c -worlds centered on a G . So believing a centered proposition is equivalent to self-ascribing a property. This means that the fortunes of PSA and NCP rise and fall together.

5. GENERATION GAPS REVISITED

Like its unstructured propositionalist neighbors, PSA is in fine shape stability-wise, as properties are very much in the public domain. The outlook vis-à-vis generation gaps, however, is mixed. The good news is that PSA avoids undergenerating beliefs in the classical manner. The bad news is that it leaves the overgeneration gap wide open.

First the good news. Undergeneration stems from the problem of self-location. We can sum up this problem in a general way as follows. Assume that a theory of belief T says that doxastic objects are all and only X s. If omniscience with respect to X s is compatible with self-locational ignorance – whether ignorance about which world one* inhabits, or ignorance about who, where, or when one* is within that world – then T is false. The trick is to pick T so as to falsify the antecedent of this conditional. As we’ve seen, classical propositionalism fails this test. Property self-ascriptionism passes, for omniscience with respect to one’s properties precludes both global and local self-locational ignorance.

First, property omniscience precludes global ignorance. That’s because knowing all your* global properties is equivalent to knowing all the true classical propositions, and knowing the latter entails knowing exactly which world is yours*. Second, it precludes local ignorance. For just as knowing all true propositions (or all your* global properties) entails knowing your* exact location in the space of worlds, knowing all your*

properties entails knowing your* exact location in the space of individuals. That means knowing exactly which individual, and where and when, you* are. The local verdict should be intuitive, but here's a quick argument for it. Assume that you know all your* properties. Then there cannot be a distinct individual – call that person you₂ – who for all you know is you*. In order to be distinct from you, you₂ would have to have some property you lacked or lack some property you had; and if that were the case, you'd know about it – and know that you₂ wasn't you*.¹⁵ Similarly, if you know all your* properties, then in particular you know all your* local properties, positional and otherwise; in which case you must know everything there is to know about your* position in space and time, and be able to discriminate all other positions from your* own. So omniscience with respect to properties precludes self-locational ignorance of either the global or local variety.

As a corollary, note that omniscience with respect to centered propositions also precludes self-locational ignorance, since knowing a centered proposition is equivalent to self-ascribing a veridically ascribable property. If you know all *c*-propositions, then you know exactly which *c*-world is yours*; and if you know exactly which *c*-world is yours*, then you not only know which world is yours*, you also know which individual at that world is you*. So neoclassicalists needn't worry about undergeneration either.

Overgeneration is another matter, for the problem of logical omniscience again rears its ugly head. To see why, first recall that properties and propositions are twin notions. Both are regions of logical space, in one case of the space of individuals, in the other the space of worlds. We speak of the necessary propositions \top (necessary truth, the universal set of worlds) and \perp (necessary falsehood, the empty set of worlds). Analogously, we can speak of the extremal properties *Max* (maximality, the set of all individuals) and *Min* (minimality, the empty set of individuals). Propositions stand in entailment relations: a proposition *Q* is entailed by a proposition *P* just in case every *P*-world is a *Q*-world. Properties stand in a similar relation, what might be called *embedding*: a property *G* embeds a property *F* just in case every instance of *F* is an instance of *G*.

The omniscience problem takes much the same form as before. Let *F* be a property which an agent *A* self-ascribes, and *G* any property that embeds *F*. Each of *A*'s doxastic alternatives is an *F*, and every *F* is a *G*, so each of *A*'s doxastic alternatives is a *G*. Hence *A* self-ascribes *G*. In particular, *A* can't help but self-ascribe *Max*. And if *A* is, like most of us, self-ascriptively inconsistent – that is, if the properties she self-ascribes are not jointly coinstantiable – then *A* has no doxastic alternatives; hence she self-ascribes all properties. In other words, she believes everything.

To illustrate the problem, consider young Anne, who believes that wildebeests have short tails but has never heard of gnu. According to PSA, Anne self-ascribes the property of inhabiting a world in which wildebeests have short tails. Since ‘wildebeest’ and ‘gnu’ directly co-refer, every short-tailed-wildebeest world is a short-tailed-gnu world (and vice versa), so Anne self-ascribes the property of inhabiting a short-tailed-gnu world. Hence Anne believes that gnus have short tails. And like all her fellow individuals, Anne instantiates *Max*. We can describe *Max* as the property of inhabiting a world in which φ holds, where ‘ φ ’ is a meta-linguistic variable ranging over sentences which express necessary truths, such as true sentences of arithmetic or true identities flanked by natural-kind terms. Though Anne is six and has yet to encounter exponentiation, she self-ascribes the property of inhabiting a world in which Fermat’s Last Theorem is true. That is, Anne believes that $x^n + y^n = z^n$ has no integral solutions for $n > 2$. She also believes that wildebeests are gnus. Or so the property theorist would have *us* believe.

Neoclassicalists have no easier time of it. Let P_c be a centered proposition that *A* believes, and Q_c a deductive consequence of P_c . Then every *c*-world in *A*’s belief-set is a Q_c . Hence, *A* believes Q_c . In particular, *A* believes the necessary *c*-proposition \top_c . And if her beliefs are inconsistent, then she believes every *c*-proposition there is. Alas, NCP offers no resources for dealing with the problem of logical omniscience apart from those of its classical predecessor, and we’ve already seen that those resources are insufficient. So neoclassicalism inherits the overgeneration gap, as expected.

6. CLOSING MORAL

One might be tempted to think that persistent overgeneration would not justify our abandoning PSA and NCP, given that undergeneration has been taken care of.¹⁶ But this is a temptation we should resist.

Devising a semantics for belief states means devising a theory of belief *simpliciter*, not merely a theory of contingent belief. Like classical propositionalism, both property self-ascriptionism and neoclassical propositionalism purport to be general in scope, and we should judge their viability on that basis. Lewis himself rightly points out, in a different context, that causal accounts of knowledge stumble badly when they come to mathematics.¹⁷ What Lewis fails to appreciate is that unstructured accounts of belief like PSA, the account he favors, have the same problem. Such views are fine as far as they go, but they have their limits – limits which a minimally adequate account of belief must transcend. Among the

costs of transcendence may be the replacement of unstructured doxastic objects, like sets of worlds or sets of individuals, with structured ones, like Russellian propositions or Fregean thoughts. Whether accounts of the latter type fare any better overall, however, is a question that must be left to another day.

ACKNOWLEDGEMENTS

Thanks to Curtis Brown, Jonathan Cohen, Mark Moyer, Robert Stalnaker, Josef Stern, and Bill Tait for helpful discussion of the issues.

NOTES

¹ See Stalnaker (1981, 1984, 1995, 1999).

² See Hintikka (1969) and Lewis (1986).

³ One might have thought that the latter concept was vexed enough. See Lewis (1986) and Stalnaker (1984) for rival views of the ontological status of possible worlds.

⁴ See Levesque (1984) for details.

⁵ See, for example, Dretske (1988) and Stalnaker (1984).

⁶ This is not to say, however, that *no* explicit/implicit distinction marks a joint in doxastic reality. It is only that we need to carve elsewhere than where Levesque suggests. For what seems to me the best alternative proposal currently available, see Crimmins (1992). On Crimmins's account, to believe something implicitly (or tacitly) is to believe it in a 'virtual' way. That means being cognitively disposed – equipped with tendencies of inference, emotion, and action – in a manner roughly isomorphic to the explicit case, *sans* explicitization. As a result, most implicit beliefs in Levesque's sense do not show up as implicit/tacit beliefs on Crimmins's account, since the requisite cognitive dispositions, such as the disposition to assent to a sentence expressing the relevant proposition, are typically absent.

⁷ This is a *leitmotif* of Cherniak (1986).

⁸ Compare Frege's treatment of identity in the *Begriffsschrift* as "a relation between names or signs of objects" (described as such in Frege (1984)). His later abandonment of this view led him to introduce senses, as a means of picking up the resulting slack.

⁹ Further, albeit rather indirect, support for diagonalization may be gleaned from 'combinatorial' or 'syntactical' accounts of mathematical truth, such as formalism and conventionalism (see Benacerraf (1973)). Such accounts, by defining truth in terms of derivability or theoremhood in a suitable formal system, also yield a sense in which mathematics is "about language," since they reduce mathematical facts to facts about syntax.

Notice, however, that the formalist/conventionalist sense in which mathematics is "about language" differs fundamentally from the Stalnakerian one. For a key feature of combinatorial accounts is the way they supplant semantic notions, such as reference and satisfaction, with syntactic notions, such as substitution and derivation. Diagonalization effects no such replacement; it leaves the notion of truth as an unreduced semantical primitive. In short, the analogy between combinatorial accounts and Stalnaker's view does not go very deep.

¹⁰ For qualms about the open-endedness of this characterization, see Section 3.4.

¹¹ If y is the individual Ingrid Bergman and z the location of St. Louis, then the standard propositional contents of A 's tokens of (7) and (8) are the set of worlds in which x is identical to y and the set of worlds in which x is at z , respectively. But the matter may not be that simple; see Section 3.2 for a complication.

¹² It is implicit here that worlds are 'total' in the sense of Perry (1986), since any world in which P fails must be a world in which $\neg P$ holds.

¹³ See Anscombe (1975), O'Brien (1994), and Bermúdez (1998).

¹⁴ The *locus classicus* is Quine (1969); for a more recent incarnation of the same idea, see Chalmers (1996).

¹⁵ The argument here relies on the Law of Identity of Indiscernibles, but in a version which should be weak enough to avoid controversy. It says that if two individuals have all their properties in common, including all spatio-temporal (positional) properties, then they're identical. See Armstrong (1978) for discussion.

¹⁶ Lewis suggests as much when he says that the problem of "ignorance of noncontingent matters . . . cuts across the issues I want to discuss, so I shall ignore it" (Lewis 1979, 515).

¹⁷ See Lewis (1986, 109).

REFERENCES

- Anscombe, G. E. M.: 1975, 'The First Person', in S. Guttenplan (ed.), *Mind and Language*, Oxford University Press, Oxford.
- Armstrong, D. M.: 1978, *Universals and Scientific Realism*, Vol. I, Cambridge University Press, Cambridge.
- Benacerraf, P.: 1973, 'Mathematical Truth', *Journal of Philosophy* **70**, 661–680.
- Bermúdez, J. L.: 1998, *The Paradox of Self-Consciousness*, MIT Press, Boston, MA.
- Castañeda, H.-N.: 1966, '"He": A Study in the Logic of Self-Consciousness', *Ratio* **8**, 130–157.
- Chalmers, D.: 1996, *The Conscious Mind*, Oxford University Press, Oxford.
- Cherniak, C.: 1986, *Minimal Rationality*, MIT Press, Boston, MA.
- Crimmins, M.: 1992, 'Tacitness and Virtual Beliefs', *Mind and Language* **7**, 240–263.
- Dretske, F.: 1988, *Explaining Behavior: Reasons in a World of Causes*, MIT Press, Boston, MA.
- Frege, G.: 1984, 'On Sense and Meaning', in B. F. McGuinness (ed.), *Collected Papers*, Blackwell, London.
- Hintikka, J.: 1969, 'Semantics for Propositional Attitudes', in J. W. Davis et al. (eds.), *Philosophical Logic*, D. Reidel, Dordrecht.
- Hintikka, J.: 1975, 'Impossible Possible Worlds Vindicated', *Journal of Philosophical Logic* **4**, 475–484.
- Hobbs, J. R.: 1985, 'Granularity', in A. Joshi (ed.), *IJCAI '85: Proceedings of the Ninth International Joint Conference on Artificial Intelligence*, Morgan Kaufmann, Los Altos, CA.
- Levesque, H. J.: 1984, 'A Logic of Implicit and Explicit Belief', in *AAAI-84: Proceedings of the National Conference on Artificial Intelligence*, Morgan Kaufmann, Los Altos, CA.
- Lewis, D.: 1979, 'Attitudes *de dicto* and *de se*', *Philosophical Review* **88**, 513–543.
- Lewis, D.: 1986, *On the Plurality of Worlds*, Blackwell, London.
- O'Brien, L. F.: 1994, 'Anscombe and the Self-Reference Rule', *Analysis* **54**, 277–281.

- Perry, J.: 1986, 'From Worlds to Situations', *Journal of Philosophical Logic* **15**, 83–107.
- Quine, W. V.: 1969, 'Propositional Objects', in *Ontological Relativity and Other Essays*, Columbia University Press, New York.
- Stalnaker, R. C.: 1981, 'Indexical Belief', *Synthese* **49**, 129–151.
- Stalnaker, R. C.: 1984, *Inquiry*, MIT Press, Boston, MA.
- Stalnaker, R. C.: 1991, 'The Problem of Logical Omniscience, I', *Synthese* **89**, 425–440.
- Stalnaker, R. C.: 1995, 'Stalnaker, Robert', in S. Guttenplan (ed.), *A Companion to the Philosophy of Mind*, Blackwell, London.
- Stalnaker, R. C.: 1999, 'Introduction', in *Context and Content*, Oxford University Press, Oxford.
- Vardi, M. Y.: 1986, 'On Epistemic Logic and Logical Omniscience', in J. Y. Halpern (ed.), *Theoretical Aspects of Reasoning about Knowledge: Proceedings of the 1986 Conference*, Morgan Kaufmann, Los Altos, CA.

Washington University in St. Louis
Department of Philosophy
Campus Box 1073
1 Brookings Drive
St. Louis, MO 63130-4899
U.S.A.