

Secondary Qualities and Self-Location¹

Andy Egan
Australian National University

Introduction

Colors aren't as real as shapes. Shapes are full-fledged qualities of things in themselves, independent of how they're perceived and by whom. Colors aren't. Colors are merely qualities of things as they are for us, and the colors of things depend on who is perceiving them. When we take the fully objective view of the world, things keep their shapes, but the colors fall away, revealed as the mere artifacts of our own subjective, parochial perspective on the world that they are.

Or so some have thought. And so it's often tempting to think, even for those of us who, at the end of the day, wind up thinking otherwise. It's even more tempting to think that, even if this is all crazy with respect to *colors*, there are *some* qualities of things that deserve the sort of treatment that colors are subjected to in the previous paragraph. Whatever we think about the case of color, there's definitely a pull to the idea of drawing some distinction between the fully real, objective, observer-independent qualities of things as they are in themselves, and the less-than-fully-real, subjective, observer-dependent qualities of things as they are for us.

Let's call qualities of the first kind *primary qualities*, those of the second kind *secondary qualities*, and the distinction between them the *primary/secondary quality distinction*. This terminology is partly stipulative—I won't retract anything in what follows if confronted with good textual evidence that, for example, Locke had nothing

¹ Thanks to Stephen Yablo, Robert Stalnaker, Alex Byrne, Ned Hall, Hud Hudson, Tyler Doggett, Alison Simmons, Richard Hall, Daniel Stoljar, Laurie Paul, and Karen Bennett for comments, suggestions, and discussion.

even remotely like this distinction in mind—but the choice of terminology isn't arbitrary, either. The (as yet intolerably vague and mushy) distinction I'm after is, I think, what underlies much of the interest of the historical distinction(s).

In any event, whatever we *call* the distinction, and whatever its relation to what people have actually said over the course of philosophical history, it's philosophically interesting because it's (a) often quite an attractive distinction to draw, and (b) incredibly hard to spell out in a satisfying and sensible way. I will attempt such a spelling-out in what follows, after first trying to pin down in more detail what we want from the primary/secondary quality distinction, and saying a little bit about why that's such a hard thing to get.

1. The Job Description

There is a lot of rhetoric that suggests the sort of distinction inadequately characterized above. I'll attempt to pin down the distinction a bit more carefully by looking at some of the rhetoric, and seeing what the primary/secondary quality distinction would have to be like in order to justify saying that kind of thing about the secondary qualities.

I take the task of justifying the rhetoric to be more important than making sure that the distinction classifies particular qualities as primary and secondary along the lines that philosophers have traditionally wanted to divide them. In particular, I take it to be of only secondary importance that the traditional paradigm cases of secondary qualities—sensory qualities like colors, tastes, smells, etc.—turn out to be secondary qualities, when the distinction between primary and secondary qualities is cashed out in the way that I

suggest. While I do think that, in fact, my distinction does carve pretty close to the traditional one, what I'm primarily concerned to do is to provide a distinction that justifies the rhetoric, not one that justifies the standard classifications of particular qualities as primary or secondary.

Following are some examples of the sorts of rhetoric that I'm concerned to justify. Again, what I want to focus on from these passages is just the characterization of what's supposed to be distinctive of the secondary qualities, and how they are supposed to be different from the primary qualities, rather than the claims about which particular qualities are primary and which are secondary.² I'll present fairly long string of examples all at once, and then draw out themes afterwards.

Democritus:

By convention, sweet; by convention, bitter; by convention, cold; by convention, color;
but in reality, atoms and void.³

St. Paul:

There is nothing unclean of itself: but to him that esteemeth any thing to be unclean, to him it is unclean...⁴

Locke:

What I have said concerning colors and smells, may be understood also of tastes and sounds, and other the like sensible qualities; which, whatever reality we, by mistake,

² I suspect that, for example, *colors* probably aren't secondary qualities, but the claims that are made about the colors below are still useful for characterizing the primary/secondary quality distinction.

³ From McKirahan (1994) (DK68B9). Thanks to Richard Hall for the reference.

⁴ Romans 14:14. Thanks to Tyler Doggett for the reference.

attribute to them, are in truth nothing in the objects themselves, but powers to produce various sensations in us, and depend on those primary qualities, viz. bulk, figure, and motion of parts; as I have said.⁵

The particular bulk, number, and motion of the parts of fire, or snow, are really in them, whether anyone's senses perceive them or no: and therefore they may be called real qualities, because they really exist in those bodies. But light, heat, whiteness, or coldness, are no more really in them, than sickness or pain is in manna. Take away the sensation of them; let not the eyes see light, or colors, nor the ears hear sounds; let the palate not taste, nor the nose smell, and all colors, tastes, odors, and sounds... vanish and cease...⁶

Bernard Williams:

In understanding, even sketchily, at a general and reflective level, why things appear variously coloured to various observers, we shall find that we have left behind any idea that, in some way which transcends those facts, they 'really' have one colour rather than another. In thinking of these explanations, we are in fact using a conception in which colour does not figure at all as a quality of the things.⁷

We can say, and indeed say truly, that grass before there was consciousness was green...

But it is, nonetheless, relative, relating to human tastes and interests.⁸

Thomas Nagel:

⁵ Locke (1690/1996: 51). (Book 2, chapter 8, section 14.)

⁶ *Ibid.*: 51-52. (Book 2, chapter 8, section 17.)

⁷ Williams (1978: 242).

⁸ *Ibid.*: 243. Williams is actually talking about amusingness, not green, in the second part of the quote, but nothing hangs on this.

The third step [in adopting an objective conception of the world] is to try to form a conception of that true nature [of the physical world] independent of its appearance either to us or to other types of observers. This means not only not thinking of the physical world from our own particular point of view, but not thinking of it from a more general human perceptual point of view either: not thinking of how it looks, feels, smells, tastes, or sounds. *These secondary qualities then drop out of our picture of the external world...*⁹

Colin McGinn:

Secondary qualities resemble properties like being poisonous or nourishing in this respect: plainly, these properties are relative to some implicit or explicit choice of creature as that with respect to which a substance is declared poisonous or nourishing. This relativity implies that there is no genuine disagreement between us and the Martians when they call an object green which we call red...¹⁰

I think it is an *a priori* truth that only the primary qualities correspond to how things are in themselves...¹¹

What the scientifically informed view denies is [not that objects are coloured, but] just that objects are *objectively* or *intrinsically* coloured, i.e. that objects have colour in the way that they have shape; it denies that possession of colour is an observer-independent condition.¹²

⁹ Nagel (1989: 14). My italics.

¹⁰ McGinn (1983: 10).

¹¹ *Ibid* p114.

¹² *Ibid* p118. His italics.

Some themes emerge from these passages, which provide us with three desiderata for an account of what the secondary qualities are:

Observer-Dependence: Secondary qualities are supposed to depend, in some non-trivial way, on the existence or the peculiarities of observers in a way that primary qualities do not.¹³

Relativity: Secondary qualities are supposed to be unlike primary qualities in that an object can have a secondary quality relative to one observer that it lacks relative to another.¹⁴

Less-than-full Reality: The secondary qualities are supposed to be somehow metaphysically second-class. They display a sort of *unreality*—though we’re correct in attributing secondary qualities to things, there’s something *less than fully real* about them.¹⁵

This last desideratum looks particularly problematic. It’s quite hard to say just what this could *possibly* amount to, prior to giving a bit more of an account of the distinction. Still, it does seem pretty clear that this is part of the idea.

It’s worth pointing out, on the topic of less-than-full reality, that the sort of unreality that’s wanted is *not* the sort that we’d get by saying that the secondary qualities are never instantiated, or that they’re only instantiated by mental entities and we’re

¹³ See for example the Locke, Nagel, and McGinn passages above.

¹⁴ McGinn, St. Paul, and possibly Williams are good exemplars of this kind of talk.

¹⁵ This strand of the rhetoric is particularly strong in the passages from Democritus and Williams.

mistaken when we attribute them to things outside the mind. Not that no one has ever said such things about the secondary qualities—of course many people have. But it's not part of the initial, intuitively appealing picture of them. The appealing distinction isn't a distinction between qualities that we correctly attribute to things and qualities that we mistakenly attribute to them. It's between two kinds of qualities, both of which are had by the things we attribute them to, but which differ in some metaphysically important way that privileges qualities of one kind over those of the other in terms of their capacity to genuinely characterize the way the world really is in itself. Because it's so hard to provide a sensible way of understanding "genuinely", "really" and "in itself" in sentences like the preceding, it's tempting to move to some other notion of less-than-full-reality, on which we're just mistaken in attributing the secondary qualities to the things that we in fact attribute them to. But it seems clear that this is not the sort of distinction in metaphysical status that's involved in the initially compelling picture of the distinction between primary and secondary qualities.

The project of identifying the primary/secondary quality distinction, as I'm construing it here, is the project of finding some distinction that satisfies the three desiderata above. That is, it's the project of finding a distinction that gives us an important difference between the fully real, observer-independent, non-relative qualities of things, and the metaphysically second-class, observer-dependent, relative qualities of things.

2. Trouble

One distinction that would at least get us the more real/less real part of the primary/secondary rhetoric is the distinction between properties that things actually have and properties that they don't. But this would be an extremely unsatisfying way to draw the distinction. For one thing, it would probably involve us in attributing an awful lot of systematic error. For another, it doesn't really justify the rhetoric. Any distinction between primary and secondary qualities that justifies the traditional rhetoric will have to run a lot deeper. To justify, for example, the talk about relativity and observer-dependence, it will have to turn out that, among the qualities that we *correctly* attribute to things, some of them are more real, more objective, more a part of the absolute conception of the world, than others.¹⁶

But it might be that some sort of projectivist error theory is the best that we can do. Gideon Rosen (1994) discusses the urge to draw a metaphysically substantive distinction between two kinds of *facts*, such that facts of one kind characterize the world as it is in itself, are objective, fully real, etc., while facts of the other kind characterize the world as it is for us, are subjective, less than fully real, etc. He makes a persuasive case that this can't be done.¹⁷

If Rosen is right that there's no interesting metaphysical distinction in point of objectivity and subjectivity between *facts*, then it's hard to see how there could be one at the level of *qualities*. If we had qualities that were, in some metaphysically interesting

¹⁶ Not everyone says this. One could be a projectivist, and think that the secondary qualities are qualities of our experiences that we mistakenly attribute to things outside the mind. Then the more real/less real distinction would be the distinction between properties that are correctly/mistakenly attributed to the things we attribute them to. (See for example Boghossian and Velleman (1989).) This requires accepting an error theory, though, instead of just a theory about the nature of the qualities that we correctly attribute to things. And a lot of the primary/secondary rhetoric seems to suggest a difference in the kind of quality being correctly attributed (to things outside the mind), rather than a difference in the accuracy of the attribution.

¹⁷ Alex Byrne (2001, 2002) makes some very similar points (influenced by Rosen), in response to Stroud (2000).

sense, subjective, less than fully real, etc., then we could get metaphysically second-class facts by correctly attributing those qualities to things. So if there aren't any interestingly subjective facts, there must not be any interestingly subjective qualities, either.

Rosen argues pretty convincingly that the usual ways in which people have tried to draw the distinction won't work. I'll briefly discuss just two examples here.

That some fact, quality, or entity is *mind-dependent* doesn't seem to impugn its full, first-class reality. Certainly a fact's being *causally* mind-dependent—in that its obtaining was brought about by some mental activity—doesn't make the fact metaphysically second-class. Facts about the existence of artifacts, for example, are as metaphysically respectable as facts get, and they're causally mind-dependent—the existence of my kitchen table was brought about, at least in part, by the thoughts, plans, and intentions of some carpenter. Other sorts of mind-dependence don't seem to fare any better.

Response-dependence¹⁸ also seems not to do the trick. That a thing is disposed to cause some response R in a subject S in circumstances C is a perfectly objective fact about the thing. Consider for example Locke's tertiary qualities, such as being disposed to melt wax in ordinary circumstances, or Rosen's example of being disposed to annoy fox terriers in ordinary circumstances. These are perfectly objective features of whatever has them—the fact that something is disposed to melt wax, or to annoy fox terriers, isn't in any way metaphysically second-class, subjective, or less than fully real. And if these aren't metaphysically second-class, then neither is, for example, being disposed to cause sensations of kind K in humans.

¹⁸ For general discussion of response-dependence, see for example Johnston (1989, 1993, 1998), and Pettit (1991).

Our disproportionate *interest* in facts of this last kind—facts about which things are disposed to cause which responses in humans—is subjective, parochial, and so forth, but the *facts themselves* are perfectly objective. Facts about what's disposed to cause certain kinds of sensations in us are just as much a part of the world as it is in itself as the facts about which things are disposed to melt wax or annoy fox terriers. So response-dependence doesn't seem to get us any metaphysically interesting distinction between the objective, genuine facts and those that are subjective and therefore somehow second-class.

The moral of Rosen's story seems to be: if an object has a property, then it's a perfectly objective matter of fact that it has that property. We can find properties such that our reasons for being *interested* in which things have them are subjective and parochial, but that doesn't make them metaphysically second-class—it doesn't make the fact or the property subjective in any metaphysically interesting sense.

This is trouble. It looks as if we won't be able to find any class of properties about which *any* of the secondary quality rhetoric is justified—there aren't any properties the having of which is either relative or observer-dependent in any interesting way, and the only sense we can make of a property's being metaphysically second-class is it's being uninstantiated.¹⁹ So, we won't be able to find a distinction between properties that can play the role that we wanted the primary/secondary quality distinction to play.

I actually think that this is exactly right, but it still doesn't doom the project of finding a sensible primary/secondary quality distinction. That's because I think that the distinction between primary and secondary qualities is not a distinction between two

¹⁹ One might still think there's some hope for the distinction between *natural* and *unnatural* properties as place to locate the primary/secondary quality distinction. This *is* a metaphysically interesting distinction between kinds of properties that we correctly attribute to things. But it's probably not metaphysically interesting in the right way—it doesn't, e.g., license the sorts of relativity and observer-dependence rhetoric that surrounds the secondary qualities. (Thanks to Robert Stalnaker for discussion here.)

kinds of properties; rather, it's a distinction between properties and another type of property-like thing. Explaining exactly what these other property-like things are, and why they're good candidates to be the secondary qualities, will require a detour into discussions of self-locating belief.

3. *Centering Features*

I am partial to a picture of mental and linguistic content according to which the role of mental states and linguistic representations is to distinguish between possibilities. My beliefs distinguish between the possibilities that I take to be candidates for actuality and the ones that I rule out, my desires distinguish between the possibilities that I hope for and those that I dread, and my assertions distinguish between (roughly) those possibilities that I'm asking you to rule out and those that you're free to leave open.²⁰

If we like this possibility-sorting picture of content, then it's very natural to represent contents as sets of possible worlds. The content of a belief, desire, or assertion is the set of worlds where things are as they're believed, desired, or asserted to be.

Typically (always?) the relevant sets of worlds are picked out by attributing properties to things—I believe that Fido is furry, desire that the cheese is melted, and assert that France is hexagonal. The content of my belief, desire, or assertion is the set of worlds in which the relevant things have the relevant properties.

There's a distinctive role that the property plays in determining the possible-worlds content of my belief, desire, or assertion: given an object, the property determines a set of worlds—the worlds in which the object has the property. So given Fido, the property, *being furry* determines a set of worlds: the worlds in which Fido is furry.

²⁰ See for example Stalnaker (1984).

Equivalently, given a world w , the property determines an extension—the set of things that have the property in w . Given a world w , *being furry* determines an extension—all of the things that are furry in w . Fido is in the extension of *being furry* in w iff w is a member of the proposition expressed by “Fido is furry”.²¹

So we can think of properties as functions from worlds to extensions (or, equivalently, as functions from objects to possible-worlds propositions).²² This illustrates the distinctive role that properties play in characterizing, and maybe in determining, the contents of representations—we can characterize the possible-worlds content of a representation by saying which properties it attributes to which things.²³

Take “Kermit is a frog”. This sentence picks out a set of worlds—the ones in which Kermit’s got the property, *being a frog*. (Maybe this is all of the worlds in which Kermit exists, maybe not. The question of whether frogginess is essential to Kermit doesn’t matter for present purposes.) A belief or assertion whose content is that set of worlds is one that represents Kermit as having a certain feature—namely, *being a frog*.²⁴

²¹ There are worries about *identifying* properties with such functions. For one thing, maybe part of what properties are supposed to do is explain the causal powers of things, and the similarities and differences between them, and maybe functions aren’t well suited to do that kind of work. For another thing, this proposal won’t allow us to distinguish between necessarily coextensive properties, and maybe that leads to trouble. But never mind—the identity claim isn’t important for present purposes. Whatever properties are, they at least *determine* such functions, and that’s all that they need to do in order to play the role in determining and describing contents that I’ll be talking about here.

²² Note that we need to have merely possible objects available to be elements of the extensions in the first case, and values of the functions in the second, for this to work. Also, this isn’t quite my official story—I actually think that properties are functions from <world, time> pairs to extensions, but nothing hangs on the difference here. (See Egan 2004.)

²³ Maybe it’s also true that the way the content gets determined is by bits of the thing doing the representing that are associated with the properties getting put together in the right way with bits that are associated with the objects. (That *is* how it works with natural language—at least, for the simple, subject/predicate bits of it – things are a bit more complicated for, e.g., universal and existential claims, where the properties aren’t being attributed to any particular object, but not, I think, in a way that makes any trouble. Fodor & co. think that’s how it works with mental representation, too, but that’s more controversial. So I’m going to be adamant about the description/characterization claim, tentative about the determination claim.)

²⁴ I should point out that I’m working with a pretty undemanding notion of what it takes to correctly attribute a property to something. I correctly attribute a property F to an object a iff I believe some proposition p such that (i) p is the proposition that you get by applying F to a , and (ii) p is true. (In order to

Not-so-bold claim: this is a really central, indispensable way of characterizing the possible-worlds contents of mental states, assertions, etc.

There's good reason to think that some contentful mental states—beliefs and desires, for example—don't have (merely) possible-worlds content, but have *centered-worlds* content instead (or as well). (This is motivated by arguments and examples from Perry (1979) and Lewis (1979).)

A *possible world* is a way things might have been. For present purposes, we can either follow Lewis and take possible worlds to be big concrete universes—more things of the same kind as us and all of our surroundings—or we can follow Stalnaker, and take them to be maximally specific properties that the whole universe might have instantiated.²⁵ A *centered world* is to a possible world what a map with a “you are here” arrow added is to an arrowless map. Centered worlds single out not just a way for the world to be, but also a location within the world. They're probably best thought of as ordered pairs of a world and a center. There are different ways of picking out a center, and different advocates of centered worlds make different choices, but it will be convenient for present purposes to take centers to be <individual, time> pairs.²⁶

If we think of representation in terms of selection among possibilities, then we can think of representations with *possible-worlds contents* as selecting among possibilities *for the world*, and those with centered-worlds as selecting among

just attribute F to *a*, all I have to do is satisfy (i).) This is a pretty minimal standard for correct attribution—it doesn't make any claims about the *structure* of the relevant mental representation, for example.

²⁵ See Lewis (1973, 1986), and Stalnaker (1984, 1996).

²⁶ This means that there's a potentially confusing point of disanalogy between maps and worlds. While the map/map-with-arrow distinction is a class/subclass distinction, the world/centered world distinction is not. While a map with a “you are here” arrow is a kind of map, a centered world is not a kind of world—it's a set-theoretic construction with a world as one of its constituents. This is a point of disanalogy, but it's not one that makes any trouble for the uses that I'm going to put the analogy to. (Thanks to Tyler Doggett for pointing this out.)

possibilities *for the agent*. My beliefs, desires, etc. with possible-worlds content draw distinctions between ways the world might be, while my beliefs, desires, etc. with centered-worlds content draw distinctions between situations that I might be in. The latter are more fine-grained than the former, since (at least in most cases), each world is going to contain a number of positions for some agent to occupy.

Let us say that, when a particular centered world is the one that corresponds to my current position, I am *correctly located by* – in the sense that my location is fixed by – that centered world. No one is ever correctly located by more than one centered world at a time, and no one is ever correctly located by any one centered world for very long. I am, as I write this, correctly located by the centered world <@, <Egan, 10:07pm>>. (Adopting the <individual, time> method of picking out centers.) You are, as you read this, correctly located by <@, <you, whenever you’re reading>>. And as you read *this*, you’re correctly located by some other centered world, whose time element is very slightly later.

We ought to say that some mental representations have centered-worlds content. When I have beliefs about my location, or the time, my belief is not well represented as an attitude toward a possible-worlds proposition. Possible-worlds propositions don’t cut finely enough—knowledge of, and belief about, possible-worlds propositions can pin down which *world* I’m in, but cannot pin down my *location* (either spatial or temporal) within that world.²⁷

This much is more or less standard – all of the elements of the picture described above are likely to be familiar to readers of Lewis, Perry, and the ensuing literature.

²⁷ See Lewis (1979) and Perry (1979) for the classic examples. Stalnaker (1981) also discusses a number of nice examples, though he offers an alternative account of the phenomenon there.

Before moving on to the new stuff, we should pause for a moment to address two complications.

First, anything that determines a possible-worlds proposition determines a centered-worlds proposition. It just determines (to introduce a technical term) a *boring* centered-worlds proposition: one that includes, for each world w , either all of the positions in w or none of them. So the claim is really that many contentful mental states have contents that determine *interesting* (i.e., non-boring) centered worlds propositions.

Second, Perry draws an importantly different conclusion from the cases. He wants to locate the extra structure not in the contents of belief, but in the belief-states. So, when each of us believes that our own pants are on fire, there's no common *content* to our beliefs, but we're in the same kind of belief-state. If this kind of account is right, then we don't need centered-worlds contents for propositional attitudes. I'm going to assume without argument that this kind of account *isn't* right, and that propositional attitudes do, at least sometimes, have centered-worlds contents. (And that since centered-worlds contents are the contents of propositional attitudes, it's safe to call them propositions.) It's worth noticing that if I'm wrong about the need for centered-worlds content, pretty much everything I say from now on will be false. I feel no pangs of conscience as I move on without further discussion because: (a) I'm pretty convinced that, at the end of the day, we will want to locate the extra structure at the level of content. (b) One thing that will be relevant to *deciding* where to locate the structure is a comparison of the costs and benefits of locating it in different places. Since the ability to provide a satisfying primary/secondary quality distinction would be a significant benefit, it's worth

exploring whether or not the assumption that some beliefs, etc. have self-locating content will allow us to do that.

We describe the *possible-worlds content* of a representation (largely) in terms of the *properties* that things are represented as having. By representing Kermit as having the property, *being a frog*, we pick out a class of worlds—all and only the worlds in which Kermit is a frog. When I believe that the world is a certain way, I represent some things as being frogs and others as being bears, some things as being slimy and others as being furry, and so on. In this way I narrow down the range of worlds that I take to be candidates for actuality. (We can tell the same kind of story of other kinds of representational states and entities—for example desires, fears, and natural language sentences.)

When I have beliefs not just about what the world is like, but also about my location within it—when I have *self-locating* beliefs—something very similar is going on. I represent some things as being *nearby* and others as *far away*, some events as *present* and others as *past* or *future*, and some objects as being *to my left* and others as *on my nose*. In this way I narrow down the range of possible *predicaments*—possible locations within worlds—that I take to be candidates to be the one that I am in. (Again, the same goes for other representations with this kind of content.)²⁸

The ways in which we describe centered-worlds contents are very similar to the ways in which we describe possible-worlds contents. In the possible-worlds case, one very common way to single out a set of worlds is by attributing some property like *being*

²⁸ In talking this way, I don't intend to commit myself to any particular view about the semantics of 'nearby' in English. (And more importantly, I don't think I *do* so commit myself, any more that someone like Lewis commits himself to a view about the semantics of indexical pronouns in English when discussing the self-locating belief that *my pants are on fire*.)

a frog to some object like Kermit. In the centered-worlds case, a very common way to single out a set of centered worlds is by attributing some “property” like *being nearby* to some object like Kermit. Just as *being a frog* is a function that, when we plug Kermit into it, delivers a set of worlds (the ones in which Kermit is a frog), *being nearby* is a function that, when we plug Kermit into it, delivers a set of centered worlds—the ones where Kermit is near the center. (That is: it delivers the set of $\langle w, \langle i, t \rangle \rangle$ pairs such that Kermit is near i at t in w .) At least, this looks like a tempting thing to say. But maybe we shouldn’t be so quick to say that there are any such “properties” as *being nearby* or *being on my foot* for us to attribute to things.

It does seem right that if there *are* such “properties”, they deserve the scare quotes—they probably don’t deserve to be called *properties*. As I’ve been using the term, properties are (or at least determine) functions from worlds to extensions. A “property” like *being nearby* won’t do that.²⁹ Which things are nearby—which things are in the extension of the “property”, *being nearby*—depends not just on which world is actual, but also on *where you are* within the world.³⁰ You and I are worldmates, but (very probably) there are lots of things that are near me and far away from you. So *being nearby* doesn’t determine a function from worlds to extensions, and so it’s not a property.

Why am I so sure that they’re not properties? Well, this is partly stipulative. Nothing bad happens if I allow that they’re properties, but insist that there’s an important distinction between two kinds of properties. One reason for insisting on the name, though, is that *properties* probably ought to be the sort of thing about which a

²⁹ Again, my official view is that properties are functions from $\langle \text{world}, \text{time} \rangle$ pairs to extensions, but that doesn’t matter for present purposes; *being nearby* doesn’t determine one of those, either.

³⁰ It also depends on the currently active standards of *nearness*, but let’s ignore that for now and pretend that there’s only one standard of nearness that applies in all contexts.

nominalist/universalist/trope theorist debate wouldn't just be *crazy*. There's no plausibility at all to the idea that centering features are some peculiar kind of universal, or sets of peculiar kinds of tropes.

If *being nearby* and its ilk *were* properties, they wouldn't be fit to play the role that they do in determining centered-world contents, because we wouldn't be able to use them to distinguish between different positions within a given world. Instead, they are the *analogues* of properties for centered-worlds contents. That is, they're functions (or things that determine functions) from centered worlds to extensions, in the same way that a property is (or determines) a function from *worlds* to extensions. So given a world and a center, *being nearby* will give us the set of things that are nearby if we're at that location in that world. And given an object, *being nearby* will give us the set of centered worlds in which that object is nearby (that is, near the center). We can pick out a set of centered worlds by saying that they're the ones in which Kermit is nearby—the ones that, when plugged into the *being nearby* function, deliver an extension that includes Kermit.

We need a name for these things. I'll call them *centering features*. 'Features' to indicate that they're property-analogues, not properties, and 'centering' because their descriptive role is to select not just a world, but a center (attributing a property to Kermit selects between worlds that might be the one that I inhabit, while attributing a centering feature to Kermit selects between predicaments that might be the one that I am in).³¹

³¹ Thanks to Alex Byrne for the name. It's important to notice the difference between centering features and predicates with hidden indexicals. When I attribute a property to something using a predicate with a hidden indexical, I'm still attributing a *property*, and so I'm still expressing a possible-worlds proposition. It's just that which property I attribute to things with a use of the predicate varies from context to context. If we had a predicate that expressed a centering feature, sentences in which it occurred (in the usual way) would express *centered-worlds* propositions. (One potential difficulty with the primary/secondary quality distinction I'm trying to draw here is that it's not obvious that there *are* any sentences that express centered-worlds propositions. I argue elsewhere (in Egan, Hawthorne, and Weatherson (forthcoming) and Egan (MS)) that there are, and that some apparent difficulties with them are merely apparent.

In the same way that there are *interesting* and *boring* centered-worlds propositions, there are interesting and boring centering features—the boring ones are the ones that, when combined with objects, always determine boring centered-worlds propositions, and the interesting ones are the ones that sometimes determine interesting centered-worlds propositions.

It's worth drawing attention to the fact that not all self-locating contents are *geographically* self-locating. Consider the centered-worlds propositions, *that my pants are on fire*, *that everyone is out to get me*, *that Kermit is out to get me*, *that Kermit is out to set my pants on fire*, and *that Kermit is disposed to cause greenish sensations in me in standard viewing conditions*. (Which are, of course, distinct from the corresponding possible-worlds propositions, *that Egan's pants are on fire*, *that everyone is out to get Egan*, *that Kermit is out to get Egan*, *that Kermit is out to set Egan's pants on fire*, and *that Kermit is disposed to cause greenish sensations in Egan in standard viewing conditions*.) These sorts of self-locating contents have nothing special to do with *geographical* self-location, and the centering features that they attribute to things (e.g., *being my pants*, *being out to get me*, *being disposed to cause greenish sensations in me in standard viewing conditions*) are importantly unlike the ones (such as *being nearby*) that I've discussed so far in that they're not features of having a certain geographical location relative to the believer.

So just as self-locating belief isn't restricted to *geographical* self-location, centering features aren't restricted to relative-position features like *being nearby*. Most importantly for our purposes, they include features that are tied up with the effects that things have on our sensory apparatus (*being disposed to cause F sensations in me in*

normal viewing/hearing/tasting/etc, conditions, for example). This is particularly relevant because these sorts of centering features look much more like the traditional secondary qualities than centering features like *being nearby*.

I'll continue to frame the discussion of centering features primarily in terms of *being nearby* and *being far away*, but this is simply because it keeps the presentation cleaner, and avoids unnecessary complications, not because centering features are one and all concerned with the geographical locations of things.

By now, two things about centering features should be pretty clear: First, they're useful for describing centered-worlds contents. In the same way that it's useful to talk about the properties that people attribute to things in describing their beliefs about which world is actual (try specifying the content of my belief that Kermit is green without appealing to properties), it's useful to talk about the centering features that people attribute to things in describing their *self-locating* beliefs (beliefs about which *predicament* they're in). Second, there's nothing mysterious about them. Properties are one kind of set theoretical object–functions from worlds to extensions—and centering features are set-theoretical objects of another, analogous kind–functions from centered worlds to extensions.³²

Let me introduce some stipulative terminology:

³² Or at least they stand in some very intimate relation to such set-theoretic objects. The identity claims are more controversial. But even if the identity claims are false, centering features don't seem to be any *more* mysterious than properties. (Two further, somewhat sketchy points: The set-theoretic identity claims seem, if anything, more plausible for centering features than for properties, since centering features don't have, as part of their job description, the objective-similarity-grounding and relevance-to-causal-powers roles that properties are often taken to have in theirs. Second, if there's a split within the property role, between, say, the properties that play a role in describing and determining content and the ones that explain the similarities between and causal powers of objects, then the set-theoretic identity claims look more plausible for the content-related properties, and centering features will be of the same kind as these sorts of properties.)

A *quality* or *feature* is something that determines a function from objects to contents, where the contents might be either sets of possible worlds or sets of centered worlds.

A *property* is a kind of quality—the kind that determines a possible-worlds content.

A *centering feature* is a quality that determines a centered-worlds content.³³

A representation *attributes a quality* to an object iff the representation has as its content the proposition that's determined by applying the quality to the object.³⁴ (So representations that attribute properties to things have possible-worlds content, and representations that attribute centering features to things have centered-worlds content.)³⁵

4. Secondary Qualities as Centering Features

As we saw in section 2, here don't seem to be any *properties* that display the relativity, observer-dependence, or less-than-full reality that is supposed to be characteristic of secondary qualities. In this section I will argue that centering features *do* display all of these characteristics, and so that they are well suited to play the role of secondary qualities, and that the property/centering feature distinction is well suited to play the role of the primary/secondary quality distinction.³⁶

³³ This way of setting things up is very different from Locke's terminology. While Locke takes the qualities to be a subset of the properties, I'm taking the properties to be a subset of the qualities (the ones that combine with objects to determine possible-worlds propositions). Thanks to Hud Hudson for pointing this out.

³⁴ Better, if its content is something that entails the proposition that you get by applying the quality to the object. (This is to get it to turn out that 'Kermit is green and Big Bird is yellow' attributes *being green* to Kermit.)

³⁵ Note that these needn't be exclusive—there might be representations with both kinds of content.

³⁶ Let's distinguish two questions: (i) Can we make sense of there being *any* features of things represented in experience that behave in the way secondary qualities are supposed to? (ii) Is it plausible that the *particular* features that people have taken to be secondary qualities (e.g. colors) behave that way?

A clarification and a qualification: On my preferred view, both properties and centering features are set-theoretic constructions of a certain kind. The view I'm proposing is *not* that one of these kinds of set-theoretic construction is more real than the other. That would be crazy. Every set theoretic construction is metaphysically on all fours with every other. The proposal is, instead, that the having of a property by a thing is, in senses to be explained below, more real, more objective, etc., than the having of a centering feature by a thing.

Now the qualification: it's actually only *interesting* centering features (in the sense defined in the previous section) that are good candidates to be the secondary qualities, because only interesting centering features display the right sort of secondary-quality-ish behavior. Boring centering features act, in the relevant respects, just like properties.

I said earlier that centered worlds are to possible worlds what maps with “you are here” arrows are to maps without such arrows. A helpful extension of that metaphor is to think of possible-worlds contents as distinguishing between classes of *maps*, and centered-worlds contents as distinguishing between classes of maps-with-arrows. This picture of the difference between centered-world and possible-world contents makes it easier to see how the distinction between properties and centering features parallels the

The claim I'm mostly going to be concerned with is that the answer to the *first* question is “yes, and they're centering features”. I'm not so sure about the second, though I suspect that, at least for very many of the qualities traditionally classified as secondary, the answer is ‘yes’. The main thing that I want to do here is open up a space for a primary/secondary quality distinction, or something very like it—to show that there can be a metaphysically significant distinction in point of objectivity, genuine reality, etc. between two classes of qualities, both of which we correctly attribute to things outside the mind. I'm less worried (at least for now) about saying which side of the line particular qualities like colors, smells, etc. fall on.

The second question is interesting, but it's only interesting once we've got a serviceable answer to the first question—a serviceable characterization of what the distinction *is*. Before we can sort out whether or not, for example, colors are secondary qualities, we need to know just what (or at least *roughly* what) being a secondary quality amounts to.

distinction between primary and secondary qualities. (How to restate the same points without the metaphor should be clear enough.)

Less-than-full reality

This understanding of the primary/secondary quality distinction gives us a satisfying way of cashing out perhaps the most puzzling strand in the secondary quality rhetoric—the strand that paints the secondary qualities as somehow metaphysically second-class, less than fully real.

One way that the less-than-full-reality talk is sometimes put is that the secondary qualities go missing from the “absolute conception” of the world. They’re features of things “as they are for us” rather than “as they are in themselves”. Bernard Williams (1978) characterizes the distinction between primary and secondary qualities as the distinction between those properties that do and those that do not feature in the absolute conception of the world, where this is something like the conception that all rational inquirers, regardless of how their particular sensory equipment is set up, would have to agree on in order to be maximally well-informed.

As a distinction between *properties*, this falls victim to something very like Rosen’s attack on response-dependence as a kind of subjectivity. Which properties a thing has is just a straightforward fact about that thing. If some object has a certain property, then *everybody* has to attribute that property to the thing in order to completely and correctly characterize it—*all* of the properties of a thing will appear in the maximally complete conception of the world to which all maximally well-informed inquirers must agree. The conception of the world on which maximally well-informed Australians and

North Americans must agree includes the fact that loud drumming is disposed to annoy fox terriers. It also includes the facts that Vegemite is disposed to cause unpleasant sensations in North Americans, and that Vegemite is disposed to cause pleasant sensations in Australians. If the Australians leave out the first, or the North Americans the second, then they're not maximally well-informed. No *property* will fail to appear in the (maximally complete) conception of the world to which all maximally well-informed rational inquirers must agree.

However, when we characterize the distinction between primary and secondary qualities as the distinction between properties and centering features, it's quite plausible that secondary qualities don't appear in the absolute conception of the world. The absolute conception of the world is supposed to be the one that *everybody* has got to accept, regardless of what their perceptual apparatus, etc. is like. That looks like the conception that picks out which world is actual, and no more. Any conception of the world that has a content that's more fine-grained than a possible-worlds proposition is going to be non-mandatory—it needn't be shared by all of the maximally well-informed inhabitants of a given world.

So the features that appear in the absolute conception will be the ones that determine possible-worlds contents—the properties. The features that determine centered-worlds contents—the centering features—won't show up in the absolute conception. They'll only show up in our particular, located, parochial conception. Maximally well-informed Australians and North Americans must agree that Vegemite tastes great to Australians. They do *not* need to agree that Vegemite tastes great. That is: both Australians and North Americans must attribute the property, *tasting great to*

Australians to Vegemite, if they are to be maximally well-informed. If they fail to do so, they will be leaving something out or getting something wrong. But Australians and North Americans need *not*, in order to be maximally well-informed, both attribute the centering feature, *tasting great*, to Vegemite. Supposing (as is near enough to true) that Vegemite really does taste great to all Australians, and awful to all North Americans, a North American who attributed *tasting great* to Vegemite would be making a mistake, and so would *fail* to be maximally well-informed. (Since no North American is correctly located by the self-locating proposition *that Vegemite tastes great*.) Similarly, an Australian who failed to attribute *tasting great* to Vegemite would be making a mistake, and so would fail to be maximally well informed. (Since all Australians *are* correctly located by the self-locating proposition *that Vegemite tastes great*.) In the same way, if I am in Canberra and you are in Boston, and we are both to be maximally well-informed, we must both agree that Sydney is near Egan. We ought not to agree about whether Sydney is *nearby*.

Here are two alternative statements of the same idea: (i) If you fail to believe some true possible-worlds proposition, then you've failed to completely characterize the world. Failure to believe all of the true *centered-worlds* contents, though, is compatible with having completely characterized *the world*, though not your place within it. (ii) You don't need centering features in order to draw the map right—all you need for that is properties. You only need centering features in order to put the "you are here" arrow in the right spot. And while all of the maximally well-informed inhabitants of a world have got to agree on what the map looks like, they don't have to agree on where the "you are here" arrow points.

This difference between properties and centering features—that all of our maximally well-informed worldmates need to agree on which properties things have, though they needn't agree on which centering features they have—promises to justify a great deal of the rhetoric of less-than-full-reality (as well as the rhetoric of relativity) that surrounds the secondary qualities.

We've mentioned several times that secondary qualities are supposed to be, “not part of the world as it is in itself, but of the world as it is for us”. We can see how this is a natural way to describe centering features, too. Representations with possible-worlds contents describe the world as it is in itself—they tell us (if they're accurate) what the *world* is like. Representations with centered-worlds content describe what the world is like *for us*—they tell us (if they're accurate) about our own individual *situation* in the world; our own individual *predicament*. We can change how things are represented as being *for us* without changing how the world is represented as being *in itself*, because we can represent ourselves as being in a different *predicament* without representing ourselves as being in a different *world*. (This is one of the main points of the examples in Lewis (1979).)

All of this—the failure of centering features to appear in a Williams-type absolute conception of the world, the fact that we can pick out a unique world as actual without recourse to centering features, and the natural sense in which properties are qualities of things “as they are in themselves”, while centering features are qualities of things “as they are for us”—provides us with, I think, a quite satisfactory way of understanding the rhetoric of less-than-full reality. This is a big deal. It's easy to see the less-than-full-reality talk as hopelessly obscure, and without any hope of being cashed out in any

remotely plausible way. It's a remarkable fact that, despite this, we are still drawn to talking this way, and that passages in which people do talk this way can still ring true. I count it as a substantial benefit of my account of the primary/secondary quality distinction that it gives us a reasonably precise and plausible way of understanding the attributions of less-than-full reality to the secondary qualities.

Relativity

We also have a nice account of the *relativity* of secondary qualities. Colin McGinn (1983) says that colors are secondary qualities, and wants it to be possible that (i) the Martians attribute *being green* to the things we attribute *being red* to, and vice versa, and (ii) we're both right. If colors are centering features, we can get things being green for us and red for Martians, and vice versa. In general, we can get incompatible features F and G such that one observer represents some object as F, another represents it as G, and they're both right. We can also get the sort of change over time that McGinn wants—if we were all taste-permuted in the right way, things that used to be sweet would start being bitter. We'd still attribute the same features to things when we called them 'sweet' or 'bitter', but different things would have the features, because of the changes in us (assuming *being sweet*, for example, is something like the centering feature, *being disposed to cause S sensations in me now*).

It's far from clear that McGinn is right about how colors or tastes behave, but that's not really the point. The point is to make sense of the possibility of there being some features or other that act the way McGinn wants colors and tastes to act.³⁷

³⁷ McGinn (1983) points out a number of suggestive analogies between indexical terms and secondary qualities. One of the things I want to do here is explain the similarities—there are all these close analogies

Observer-Dependence

By taking secondary qualities to be centering features, we can also justify the rhetoric of observer-dependence that we encountered in the opening section. The first pass through this will be fairly metaphorical.

Take an ordinary possible world, without a center. Nothing has any centering features there, because there's no center for them to bear the relevant relations to. Pick a center—add the “you are here” arrow—and things suddenly take on a number of new qualities; qualities that they get not (or not entirely) in virtue of what the world is like in itself, but in virtue of where the arrow points. It's quite natural to think of these features as being *added* to the world by the selection of a center. The selection of a center—the adoption of some subjective position within the world—provides the world with all of these features which aren't present in the world considered on its own, without any center.

In the same way, centering features (unlike properties) disappear when you take away the “you are here” arrow. Nothing's *nearby*, *in my ear*, or *disposed to cause greenish sensations in creatures like me* until you specify a center. This looks like a satisfying sort of observer-dependence.

Less metaphorically: models that only contain centerless worlds won't have the resources to provide extensions for centering features. To assign an extension to a centering feature, you need to have a model that picks out a privileged center. It's much like what happens with sentences involving “actually”—you need a model that picks out a particular world as actual before you can assign them truth-values. Taking the

between secondary qualities and indexicals because secondary qualities are centering features, which are close cousins of indexical terms.

‘absolute’ view of logical space, where we don’t single out a particular world as actual, we don’t get truth-values for anything that’s got a significant occurrence of any rigidifying machinery. (See e.g. Davies and Humberstone (1980).) In the same way, taking the ‘absolute’ view of the world, where we don’t single out a particular location as *mine*, we don’t get extensions for any centering features. An *objective* model of the world—one that doesn’t privilege a particular center—doesn’t have the resources to deliver extensions for centering features. Add the center, though, and you get extensions for the centering features.

Note that there is a certain sort of observer-dependence that we *don’t* get. We don’t, even on the assumption that colors are centering features, get the truth of, “if there had been no observers, nothing would’ve been colored”. Lots of merely possible things have dispositions to cause various responses in me, even the ones in worlds where I don’t exist (or don’t have any counterparts).³⁸ The fact that the centering features don’t appear until we select a center seems like enough, though, to justify (at least most of) the kind of observer-dependence and less-than-full-reality talk that many philosophers go in for when discussing the secondary qualities. The secondary qualities “fall away” when you take the objective, observer-independent view—stop thinking from the perspective of some observer, and all of the secondary qualities disappear. (See the preceding paragraph if you’re worried about all of the metaphors here.)

³⁸ Does attribution of dispositional centering features get us involved in the potentially unpleasant business of trying to evaluate counterfactuals with centered-worlds propositions as antecedents and/or consequents? No. To see why not, it’s easiest to think about centered-worlds beliefs along Lewisian lines, as self-ascriptions of properties. When I attribute a centering feature like, *being disposed to cause G experiences in me in C* to Kermit, I’m self ascribing the property, *being one of the things in which Kermit is disposed to cause G experiences in circumstances C*. That is, I’m taking myself to be one of the things of which some open sentence along the lines of, “if x was in C, Kermit would cause G in x” is true. The only counterfactual that needs to get evaluated here is one with possible-worlds propositions as both antecedent and consequent.

An example by way of summary: We need centered-worlds contents for belief because we have beliefs not just about what the world is like, but also about our position within it. When Sarah the transatlantic sailor consults her charts, sextant, and GPS, she's not (primarily) deliberating and gathering evidence about how the world is—her main concern is with her position within a world that she already knows most of the relevant facts about.

If we have beliefs with centered-worlds contents, then we attribute centering features to things. When Sarah believes that Gibraltar is nearby, she believes the centered-worlds proposition that you get by applying the centering feature, *being nearby*, to Gibraltar. If Sarah's belief is accurate, then she's *correctly* attributed a centering feature to Gibraltar.

The centering feature that she attributes to Gibraltar is metaphysically second class compared to *properties* of Gibraltar like *being built around a big rock*. The property, unlike the centering feature, figures in the absolute conception of the world—the conception that has to be shared by all of Sarah's maximally well-informed worldmates. (Sarah's friends back in Cambridge ought also to believe that Gibraltar is built around a big rock, but they ought *not* to believe that Gibraltar is nearby.) Since Sarah correctly attributes both *being nearby* and *being built around a big rock* to Gibraltar, there's an interesting, metaphysically significant distinction in point of objectivity between two kinds of features that Sarah correctly attributes to things in her environment.

One might not be happy that it turns out that relative-location features like *being nearby* turn out to be secondary qualities. One might, for example, think that in order to

be a secondary quality, you need to have something to do with how things sensorily appear. We could accommodate this by insisting on a fourth desideratum, in addition to relativity, observer-dependence, and less-than-full reality. In that case, only a proper subset of the interesting centering features would be secondary qualities – the ones we might call the *sensory* centering features. (Features like *looking green to me*, *being disposed to cause unpleasant olfactory sensations in me in standard circumstances*, etc.) I have no strong view about whether to insist on this restriction or not, but it's certainly available for those who would like the primary/secondary quality distinction to cut closer to the traditional lines.

Conclusion

The sorts of secondary-quality rhetoric that we encountered in the first section can be extremely appealing. It turns out to be extremely hard, though, to find a primary/secondary quality distinction that justifies the rhetoric. Centering features display a lot of the behavior that's supposed to be distinctive of the secondary qualities—behavior that is difficult (probably impossible) to find in any kind of *property*. So the distinction between properties and centering features seems to be a good place to locate the primary/secondary quality distinction.

There is still an open question about just which qualities are primary and which are secondary. On which side of the line, for example, do the colors fall? And how do we tell, given some predicate, whether it expresses a property or a centering feature? The short answer is that the ones that express centering features will behave in ways that invite the rhetoric of relativity, observer-dependence, and less-than-full-reality. The long

answer is beyond the scope of this paper.³⁹

³⁹ Some potential criteria are discussed in Egan, Hawthorne, and Weatherson (forthcoming), in the discussion of expressions like “tastes great”.

References:

- Boghossian, P. and Velleman, J.D. (1989) Colour as a Secondary Quality, *Mind* **98**: 81-103.
- Byrne, A. (2001) Review of Barry Stroud, *The Quest for Reality*, *Philosophical Quarterly* **51**.
- Byrne, A. (2002) Yes Virginia, Lemons are Yellow, *Philosophical Studies* **110**:213-222.
- Comesaña, J. (2002) The Diagonal and the Demon, *Philosophical Studies* **110**: 249-266.
- Davies, M. and Humberstone, L. (1980) Two Notions of Necessity, *Philosophical Studies* **38**:1-30.
- Egan, A. (2004) Second-Order Predication and the Metaphysics of Properties, *Australasian Journal of Philosophy* **82**:48-66.
- Egan, A. (MS) Epistemic Modals, Relativism, and Assertion.
- Egan, A., Hawthorne, J. and Weatherson, B. (forthcoming) Epistemic Modals in Context, *Contextualism in Philosophy*, Preyer and Peters, eds., Oxford University Press, Oxford.
- Hume, D. (1751/1902) *Enquiry Concerning the Principles of Morals*, ed. L. A. Selby-Bigge, Clarendon, Oxford, (1902).
- Johnston, M. (1989) Dispositional Theories of Value, *Proceedings of the Aristotelian Society*, supp. **63**.
- Johnston, M. (1993) Objectivity Refigured: Pragmatism Without Verificationism, in *Reality, Representation, and Projection*, Haldane and Wright eds.
- Johnston, M. (1998) Are Manifest Qualities Response-Dependent?, *The Monist* **81**.
- Lewis, D. (1973) *Counterfactuals*, Blackwell, Oxford.
- Lewis, D. (1979) Attitudes *De Dicto* and *De Se*, *Philosophical Review* **88**:513-43.
- Lewis, D. (1986) *On the Plurality of Worlds*, Blackwell, Oxford.
- Locke, J. (1690/1996) *An Essay Concerning Human Understanding*
- McKirahan, R. (1994) *Philosophy Before Socrates*, Hackett, Indianapolis.
- McGinn, C. (1983) *The Subjective View*, Oxford University Press, Oxford.
- Nagel, T. (1989) *The View From Nowhere*, Oxford University Press, Oxford.
- Perry, J. (1979) The Problem of the Essential Indexical, *Nous* **13**:3-21.
- Pettit, P. (1991) Realism and Response-Dependence, *Mind* **100**: 587-626.
- Rosen, G. (1994) Objectivity and Modern Idealism: What is the Question?, in *Meaning in Mind*, ed. M. Michael and J. O'Leary-Hawthorne, Kluwer, Dordrecht, 277-319.
- Stalnaker, R. (1978) Assertion, *Syntax and Semantics* **9**: 315-332. Reprinted in Stalnaker (1999). Page references are to this reprinting.
- Stalnaker, R. (1981) Indexical Belief, *Synthese* **49**:129-151. Reprinted in Stalnaker (1999). Page references are to this reprinting.
- Stalnaker, R. (1984) *Inquiry*, MIT Press, Cambridge, Mass.
- Stalnaker, R. (1996) On What Possible Worlds Could Not Be, in *Benacerraf and his Critics*, ed. by A. Morton and S. Stich. Blackwell, Oxford (1996), 103-119.
- Stalnaker, R. (1999) *Context and Content*, Clarendon, Oxford.
- Stroud, B. (2000) *The Quest for Reality*, Oxford University Press, Oxford.
- Williams, B. (1978) *Descartes: The Project of Pure Inquiry*, Harvester Press.