



WILEY

Our Entitlement to Self-Knowledge: I. Tyler Burge

Author(s): Tyler Burge and Christopher Peacocke

Source: *Proceedings of the Aristotelian Society*, New Series, Vol. 96 (1996), pp. 91-116

Published by: Wiley on behalf of The Aristotelian Society

Stable URL: <http://www.jstor.org/stable/4545230>

Accessed: 11-04-2017 02:15 UTC

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at <http://about.jstor.org/terms>



The Aristotelian Society, Wiley are collaborating with JSTOR to digitize, preserve and extend access to *Proceedings of the Aristotelian Society*

OUR ENTITLEMENT TO SELF-KNOWLEDGE

Tyler Burge and Christopher Peacocke

I—Tyler Burge

I want to understand our epistemic warrant for a certain range of judgments about our own thoughts and attitudes. I am guided by two hypotheses. One is that there are certain sorts of self-knowledge that are epistemically special. The other is that the epistemic right or warrant we have to these sorts of self-knowledge is, in a sense, environmentally neutral. I want to understand this specialness and this environmental neutrality.

The hypothesis of epistemic specialness will be argued for in this paper. The hypothesis of environmental neutrality is relevant to a project that deals with scepticism and the nature and functions of reason. I will not develop this latter hypothesis in depth here, but I will comment on it for the sake of orientation.

Most of our empirical thoughts and our thoughts about our empirical thoughts depend for their individuation conditions on relations that we bear to a particular environment. But, on my guiding hypothesis, our epistemic warrant for our judgments about our thoughts does not depend on particular relations to a particular environment. It is common to any environment and derives from the nature of the thinker as a critical reasoner. This point is relevant to showing that certain claims to self-knowledge which are among the premises in a certain anti-sceptical argument do not beg the question by depending on presumptions about the environment that the sceptic calls into question. In this paper I will not discuss scepticism. But I begin with the sort of *cogito*-like judgments that figured in traditional anti-sceptical arguments. I believe these judgments relevant not only to scepticism, but to the epistemic specialness of some self-knowledge. Although some striking features of *cogito*-like judgments are not shared by all members of the wider range of judgments about one's thoughts whose epistemic status interests me, *cogito*-like judgments provide a useful paradigm for reflection.

So I begin with some remarks about a judgment that:

- (1) I am thinking that there are physical entities.

This judgment is an instance of *cogito*-like thoughts, an elaboration of Descartes' *I am thinking*. Let us construe 'thinking' in (1) minimally—as *engaging in thought* or *having a thought*, regardless of whether it is merely entertaining a thought, making a judgment, or whatever. In this sense, one 'thinks' all propositional components of any thought one thinks (including negated ones, antecedents of conditionals, and so on). (1) is the content of my judgment. I accept it as true. To be true, (1) requires only that I am engaging in some thought whose content is that there are physical entities.

We do not rest this judgment upon any observation or perception such as was traditionally called 'inner sense'. The judgment is direct, based on nothing else. Making the judgment requires sufficient understanding to think (1). But once one makes the judgment, or indeed just engages in the thought, one makes it true. The thought is *contextually self-verifying*. One cannot err if one does not think it, and if one does think it one cannot err. In this sense, such thinkings are *infallible*.

I do not claim that judgments like (1) are indubitable. The scope for human perversity is very wide. One could be so far gone as to think to oneself: 'I do not know whether I am now thinking or not; maybe I am dead or unconscious; my mantra may have finally made me blissfully free of thought'. Such mistaken doubt would evince cognitive pathology, but I think it possible. It is an error, however, that most people would avoid without swerving.

Key features of (1) are shared by judgments of

- (2) I judge, herewith, that there are physical entities.

When *judge* in (2) is used to execute not merely describe a judgment, judgments of (2) are contextually self-verifying. (2) is not made true by the mere thinking of it, nor does it have quite the same quasi-logical self-evident status that (1) does. These are subtleties that I will have to discuss on another occasion.

(1) and (2) are not mere philosophical curiosities. I think that they represent the form of many ordinary self-aware judgments (at least when (1) is taken to have the 'herewith' reflexivity of (2).) When one makes a judgment and is conceptually aware of one's so doing,

whether or not one spells out this conceptual awareness, one's judgments have a reflexive form like that of (2). Such conceptual self-awareness goes beyond simply consciously thinking a thought, but it is not an unusual phenomenon among people with normal second-order abilities. Thus I believe that *cogito*-like judgments constitute a significant segment of our everyday mental activity.

To remark that (1) and (2) are *contextually self-verifying* is to remark on their truth conditions, not on our justification or epistemic warrant in thinking them. It does seem that understanding (1) suffices for knowing that it is true. And the relevant understanding requires no great perspicacity. But noting that it is self-evidently self-verifying (supposing that this needed no more comment—which of course it would) would not capture fully what is involved in its epistemic status. For I think that *cogito*-like judgments share an interesting epistemic status with a number of types of self-knowledge that are not contextually self-verifying or infallible, and that lack the quasi-logical status of (1). I have in mind a wider class of judgments about states, not just reflexive occurrences—judgments about what one believes, wants, intends.

When we make judgments about many of our mental states and events, our judgments commonly constitute knowledge. I know very well that I believe that there are physical entities—if I judge that I do. Such judgments do not merely evince an inner state in the way that a yelp evinces a pain; nor are they avowals or conventional practices without cognitive value. What is the epistemic status of such judgments? What epistemic warrant do we have to make them?

I take the notion of epistemic warrant to be broader than the ordinary notion of justification. An individual's epistemic warrant may consist in a justification that the individual has for a belief or other epistemic act or state. But it may also be an *entitlement* that consists in a status of operating in an appropriate way in accord with norms of reason, even when these norms cannot be articulated by the individual who has that status. We have an entitlement to certain perceptual beliefs or to certain logical inferences even though we may lack reasons or justifications for them. The entitlement could in principle presumably—though often only with extreme philosophical difficulty—be articulated by someone. But this articulation

need not be part of the repertoire of the individual that has the entitlement.

Our epistemic warrant to much of our self-knowledge is of this sort. Most of us have no justifying argument or evidence backing the relevant judgments. The judgments are immediate, non-inferential. Although *cogito*-like judgments may count as self-evident or self-verifying, most judgments that interest me do not. Wherein are we being reasonable—in the sense of operating under norms sanctioned by reason—in making judgments about our own minds?

As I have intimated, the remarks about self-verification suggest an initial analogy between *cogito*-like judgments and knowledge of simple logical truths. The truth of judgments of (1) and (2) *is*, in a broad sense, present in the form and logic of the thought. There is something of the same self-evident and obvious features here as there are in simple logical truths. The main differences are that *cogito*-like judgments are dependent on being thought for being true, and are in their specially direct way self-verifying.

Another analogy to knowledge of simple logical truths is this: The key to the epistemic status of *cogito*-like judgments seems to reside in ordinary understanding, not in some mechanism connecting the knower with a sensed object.

This point will be one of the key elements in my account of the environmental neutrality and specialness of self-knowledge. Perceptual experiences particular to a given environment inevitably figure in the acquisition of understanding of almost any given content. But one's epistemic warrant for believing the content may not incorporate the perceptual experiences or beliefs that go into understanding it. This is the traditional view of knowledge of logical or mathematical truths. One may need perceptual experience to come to understand simple logical or arithmetical notions and truths. (This is surely the case with such logical truths as 'nothing is both a dog and not a dog'.) But on the traditional view such experience is not a constituent of one's justification or entitlement in believing simple logical or arithmetical truths.

I am not arguing for the traditional view—just recalling it. The element in it relevant to our purposes is the following. The account of epistemic justification or entitlement may presuppose understanding, which may be dependent on particular perceptual

relations to a given environment. But the account need not include perceptual beliefs or experiences as *constituents* in the individual's justification or entitlement. The account can allow attribution of concepts to the individual which could be acquired only in a limited range of possible environments, while itself taking a form that is applicable to any critical reasoner, regardless of the particular environmentally dependent contents of his or her thought.

I want to illustrate the relevance of this idea to our discussion by reconsidering the scenario of one's being switched between different environments unawares—a scenario I discussed in a paper some years back.¹ Let us assume for the sake of argument that my thinking that there are physical entities (hence my thinking that I am thinking that there are physical entities) is the thought that it is because of relevant causal relations I bear to actual physical objects in my environment. Let us also assume that an individual with a chemically identical body could have been brought up in a situation in which such relations were lacking—and in which the concept of physical object could not be acquired—but in which different, counterpart thoughts occurred. (I doubt that *physical object* is a concept universal to all possible critical reasoners; but if one did not doubt, another concept could be chosen.) Finally, let us grant that if at any time one were switched unawares from one's actual situation into such a counterpart situation, one would have no resources that would tip one off to the difference.

Unless memory and learning connections to the original environment were broken, it is hard to describe a switch of actual situations that would produce a new twin set of the concepts, with no residue from the past experiences. So in the case I am imagining one's thoughts do not switch to twin thoughts. Because of a switch one's thoughts might, however, change content, broadening their extensions without one's being aware of their doing so.

I take it that this observation is sufficient to prompt the following question. Given that we are insensitive to such alleged possible changes in content, how can we know what we are thinking?

1 'Individualism and Self-Knowledge' *Journal of Philosophy* 85 (1988), pp. 649–663. Reprinted in Quassim Cassam ed. *Self-Knowledge* (Oxford, Oxford University Press, 1994).

I will not try to deal with this question in all its ramifications here. But as I noted some years back, some of the negative force of the question can be shown to be illusory by this consideration: There is no way for one to make a mistake about the content of one's present-tensed thought in the relevant cases.

Suppose that I think that I am engaging in a thought that there are physical objects. In thinking this, I have to engage in the very thought I am referring to and ascribing to myself. The reference to the content—expressed in the that-clause—cannot be carried out unless I actually engage in the thought. The intentional content mentioned in the that-clause is not merely an object of reference or cognition; it is part of the cognition itself. It is thought and thought about in the same act. If background conditions are different enough so that I am thinking different thoughts, they will be different enough so that the objects of reference and self-ascription will also be different. So no matter how my thoughts are affected, no matter how I am switched around, I will be correct in self-ascriptions of content that are correctly expressed in *cogito*-that-clause form.

It would be a mistake to reply that because one's correct reference does not give one any grasp of what one is referring to, this reference is empty. For to self-ascribe thoughts in the way expressed by that-clauses, one has to understand the thoughts one is referring to well enough to think them. One need not have any more explicatory understanding of one's thoughts than is necessary to think them. One need not master anti-individualism, much less have an empirical mastery of the conditions that have established the identity of the thoughts one thinks. Such mastery is emphatically not guaranteed by mastery of *cogito*-self-ascriptions. But one *is* guaranteed that one ascribes something of which one has the ordinary understanding involved in using concepts and thinking thoughts.

This understanding presupposes the causal-perceptual relations to a particular environment that help determine what content is available for being understood. What one *can* think is partly dependent on relations to one's environment. And one's second-order self-ascriptions inherit both the content and the background environmental content-determining conditions from one's first-order understanding.

I have granted that one need not be sensitive to actual or counterfactual changes in what one understands under transportations into

environments where the content of one's understanding changes or would be different. One need not be capable of detecting such changes. But in any situation in which a person can think the relevant *cogito*-like judgments, the person would think them with understanding—and to all appearances, knowledgeably. At any rate, there is no obvious reason why knowledge in such judgments would be prevented by such changes, much less such possible changes.² In any such twin situation, the person would understand the self-ascribed contents and would self-ascribe them with a justice and reliability that is equal to that in any ordinary situation. Some entitlement attaching to understanding seems to be what the self-knowledge depends upon, not on some knowledge of what the understanding consists in, or whether it differs from understanding that is past or possible.

The person's epistemic entitlement to the self-ascriptions presupposes understanding. Understanding is, as I have noted, dependent on and local to causal-perceptual relations to a given environment. But the entitlement that underlies knowledgeable *cogito*-like thoughts and other self-ascriptions does not seem local and seems to survive such switches. It seems to be carried somehow by the fact that we correctly self-ascribe any content at all with understanding. Where does the entitlement derive from? And what makes it capable of surviving such environmental switches?

- 2 Are there switching situations in which one would have reasonable ground for doubting what contents one is thinking, so that a *cogito*-type judgment would not constitute knowledge? This is very complex, but I will make a few remarks here. The self-ascription in the that-clause way cannot involve a mistake about the intentional content. So the possibility of a switch does not threaten a mistake. I think therefore that such possibilities pose no relevant alternative threat to one's entitlement to one's judgment about the that-clause content of one's thoughts. I believe that the relevant minimal understanding suffices for knowledge in *cogito*-like judgments. Even in non-*cogito*-like judgments, switches in content cannot, for the same reason, undermine knowledgeability of the content of self-ascriptions. Cf. my 'Individualism and Self-Knowledge' *op. cit.*, p. 659. A fuller story has to be told about the propositional-attitude concepts in non-*cogito*-like judgments. I think the possibility of switching, or of errors of incomplete understanding, do not by themselves undermine knowledge; but I will have to discuss these matters further elsewhere.

Some worries about switching situations have focused on memory. I think that they tend to confuse preservative memory with memory of objects or with comparisons within memory, and to overrate the extent to which the content retrieved in memory is sensitive to immediate environmental context. For a discussion of preservative memory, see my 'Content Preservation' *The Philosophical Review* 102 (October 1993), pp. 457–488.

I think that the relevant entitlement derives not from the reliability of some causal-perceptual relation between cognition and its object. It has two other sources. One is the role of the relevant judgments in critical reasoning. The other is a constitutive relation between the judgments and their subject matter—or between the judgments about one's thoughts and the judgments' being true. Understanding and making such judgments is constitutively associated both with being reasonable and with getting them right.

Briefly drawn, my line of thought will be this. To be capable of critical reasoning, and to be subject to certain rational norms necessarily associated with such reasoning, some mental acts and states must be *knowledgeably* reviewable.³ The specific character of this knowledgeable reviewability requires that it be associated with an epistemic entitlement that is distinctive. The entitlement must be stronger than that involved in perceptual judgments. There must be a non-contingent, rational relation, of a sort to be explained, between relevant first-person judgments and their subject matter or truth.

All of us, even sceptics among us, recognize a practice of critical reasoning.⁴ Critical reasoning is reasoning that involves an ability to recognize and effectively employ reasonable criticism or support for reasons and reasoning. It is reasoning guided by an appreciation, use, and assessment of reasons and reasoning as such. As a critical reasoner, one not only reasons. One recognizes reasons as reasons. One evaluates, checks, weighs, criticizes, supplements one's reasons and reasoning. Clearly, this requires a second-order ability to think about thought contents or propositions, and rational relations among them.

3 I think that the following necessity also holds: To think the relevant first-person present tense thoughts about one's thoughts and attitudes, one must be capable of critical reasoning. Indeed, I think that to have a fully formed first-person concept or fully formed concepts of propositional attitudes, one must be capable of critical reasoning. To master concepts of propositional attitudes in a suitably rich sense, one must be capable of appreciating the force and relevance of reasons to attitudes as such, which amounts to being able to reason critically about reasons and reasoning. And to master a fully formed first-person concept, one must have concepts of propositional attitudes.

4 In actual practice, critical reasoning approximates what I call *reflective* reasoning. Reflective reasoning makes use of all the main concepts necessary to a full understanding of essential or fundamental elements in reasoning. Critical reasoning is simply reasoning that is sufficiently articulate to appreciate reasons as reasons and to employ articulated criticism of reasons and reasoning (as reasons and reasoning).

When one carries out a proof, one checks the steps of the reasoning, making sure that the inferences are valid. Any activity of proof requires some conception of validity, which requires an ability to think of the propositions in a proof as constituting reasons for what follows from them. Indeed, it is arguable that use of *therefore* in reasoning—deductive or otherwise—constitutes an exercise of this meta-cognitive ability. When one engages in practical deliberation, one articulates and weighs considerations on each side, goes over possible sources of bias, thinks through consequences. Essential to carrying out critical reasoning is using one's knowledge of what constitutes good reasons to guide one's actual first-order reasoning.

A non-critical reasoner reasons blind, without appreciating reasons as reasons. Animals and small children reason in this way. But reasoning under rational control of the reasoner is critical reasoning. Not all reasoning by critical reasoners is critical. Much of *our* reasoning is blind, poorly accessible, and unaware. We change attitudes in rational ways without having much sense of what we are doing. Often we are poor at saying what our reasoning is. Still, the ability to take rational control of one's reasoning is crucial in many enterprises—in giving a proof, in thinking through a plan, in constructing a theory, in engaging in debate. For reasoning to be critical, it must sometimes involve actual awareness and review of reasons; and such a reviewing standpoint must normally be available.⁵

Critical reasoning involves an ability not merely to assess truth, falsity, evidential support, entailment, and non-entailment among *propositions* or *thought contents*. It also involves an ability to assess the truth and reasonability of reasoning—hence *attitudes*. This is not to say that critical reasoning must *focus* on attitudes, as opposed to their subject matter. Normally we reason not about ourselves but about the world or about practical goods. But to be fully a critical reasoner, one must be able to—and sometimes actually—identify, distinguish, evaluate propositions as asserted, denied, hypothesized

5 I think Kant neglected distinctions between reasoning, critical reasoning, and reflective reasoning. But he clearly saw that it is the possibility of applications of 'I think' to our thoughts—not our being self-aware in this way all the time—that is basic to full reflective rationality. Of course, the form of 'I think' does not by itself make the relevant contribution to reflective rationality. One could dream *cogito*-thoughts. It is the ability to be conceptually aware of oneself as thinking with a certain control and agency that is crucial.

or merely considered.⁶ Such abilities and activities are central to argumentation. Similarly, in critical practical reasoning, one must be able to—and sometimes actually—evaluate propositions conceptualized as expressing pro-attitudes, to distinguish them explicitly from those that express beliefs, and to evaluate relations of reason among such propositions as so conceptualized. Such evaluation constitutes minimal evaluations of propositional attitudes.

To be a critical reasoner, one must also be able to, and sometimes actually, use one's knowledge of reasons to make, criticize, change, confirm *commitments* regarding propositions—to engage explicitly in reason-induced changes of mind. Critical reasoning here involves an ability to distinguish subjectivities from more objectively supportable commitments and to explicitly alter the former in favour of the latter. Its point is reasonably to confirm and correct attitudes and reasoning (not merely assess propositional connections), by reference to rational standards.

Critical reasoning must be exercised on itself. Any critical reasoning, even about abstract propositional relations or about the reasoning of others, involves commitments by the reasoner. And genuinely critical reasoning requires an application of rational standards to *those* commitments. A being that assessed good and bad reasoning in others or in the abstract, but had no inclination to apply such standards to the commitments involved in those very assessments, would not be a critical reasoner. To reason critically—to consider reasons bearing on the truth of some matter, to suspend belief or desire, to weigh values under a conception of the good—one must treat one's own commitments as matters to be considered and evaluated. Critical evaluation of one's own commitments is central to forming them and to rationally changing one's mind or standing fast.

So critical reasoning requires thinking about one's thoughts. But it further requires that that thinking be normally knowledgeable. To appreciate one's reasons as reasons—to check, weigh, criticize,

6 In effect, Frege's use of the assertion sign is an acknowledgement of a minimal use of these abilities. Without an ability to recognize that a proposition should be and is judged to be true, one cannot reason critically. Having a concept of judgment and using it in reasoning meets my requirement.

confirm one's reasons—one must know what one's reasons, thoughts, and reasoning are. One need not always be knowledgeable, or even right. But being knowledgeable must be the normal situation when one reflects on one's reasons in the course of carrying out reasonable inquiry or deliberation. The interest here is less in the requirement of normal knowledgeability—which is shared with other cognitive activities, such as perception. The interest lies in the ground of the requirement. Why must we be normally knowledgeable about our thoughts when we reflect upon them?

I will answer this question in three stages. First, I want to show that to evaluate reasons critically, one must have an *epistemic entitlement* to one's judgments about one's thoughts, reasons, and reasoning. Second, I want to support the stronger thesis that critical reasoning requires that one *know* one's thoughts, reasons, and reasoning. Third, I will try to show that this knowledge must take a distinctive, non-observational form.

So I begin with the matter of entitlement. The basic idea is simple. Put crudely: since one's beliefs or judgments about one's thoughts, reasons, and reasoning are an integral part of the overall procedures of critical reasoning, one must have an epistemic right to those beliefs or judgments. To be reasonable in the whole enterprise, one must be reasonable in that essential aspect of it.

Less crudely, consider the process of reasoning which involves the confirming and weighing of one's reasons. One must make judgments about one's attitudes and inferences. If one's judgments about one's attitudes or inferences were not reasonable—if one had no epistemic entitlement to them—one's reflection on one's attitudes and their interrelations could add no rational element to the reasonability of the whole process. But reflection does add a rational element to the reasonability of reasoning. It gives one some rational control over one's reasoning.

To put the point somewhat more fully: if one lacked entitlement to judgments about one's attitudes, there could be no norms of reason governing how one ought check, weigh, overturn, confirm reasons or reasoning. For if one lacked entitlement to judgments about one's attitudes, one could not be subject to rational norms governing how one ought to alter those attitudes given that one had reflected on them. If reflection provided no reason-endorsed judgments about the attitudes, the rational connection between the

attitudes reflected upon and the reflection would be broken. So reasons could not apply to how the attitudes should be changed, suspended, or confirmed *on the basis of* reasoning depending on such reflection. But critical reasoning just is reasoning in which norms of reason apply to how attitudes should be affected partly on the basis of reasoning that derives from judgments about one's attitudes. So one must have an epistemic entitlement to one's judgments about one's attitudes.

I turn now to the stronger thesis. One might imagine some gap between epistemic entitlement and knowledge. Might one have an epistemic entitlement but be systematically mistaken? Or might failure of some third Gettier-type condition (beyond truth and epistemic entitlement) undermine knowledge?

It is possible in given cases for reflection to be disconnected in these ways from the attitudes purportedly reflected upon. But both possibilities if generalized are incompatible with our having the sort of entitlement to the reflection just argued for. That entitlement rested on the assumption that reflection added a rational element to the reasonability of the whole process of critical reasoning—a process whereby object-level attitudes are guided by reflection on their reasonability. If reflective judgments were not normally true, reflection could not add to the rational coherence or add a rational component to the reasonability of the whole process. It could not rationally control and guide the attitudes being reflected upon (even though one could imagine situations in which such disconnected reflection would be mechanically or instrumentally beneficial in forming true or rational beliefs). So reflection would not add in the relevant way to the reasonability of the process, and therefore would not have the source of entitlement just argued for.

The same point applies to the possible failure of some Gettier-type condition. Again, if reflection were connected to the truth of our judgments about our thoughts in an accidental or non-knowledge-yielding way, the reason-guiding and rational-coherence-making functions of rational review would be broken. Since part of our entitlement to reflective judgments about our attitudes derives from their functions in critical reasoning, the entitlement itself would be undermined.

Not only the relevant entitlement to reflective judgments that derives from their functions within critical reasoning, but critical

reasoning itself is constitutively dependent on the truth- and Gettier conditions being met. If a being had an epistemic entitlement to its judgments about its attitudes but were systematically mistaken about them—never got them right—it would not be a critical reasoner. Or if our entitlement were always connected to the truth of our judgments about our thoughts in an accidental or non-knowledge-yielding way, critical reasoning would not be possible. For critical reason requires rational integration of one's higher-order evaluations with one's first-order, object-oriented reasoning. The former must be reason-guided and reason-guiding. And they must cement the rational coherence between the two levels. If the two came radically apart, or were only accidentally connected, critical reasoning would not occur.

So if we failed normally to know our thoughts and attitudes, in ordinary reasoning about reasons, either through systematic falsity of our judgments or through systematic mismatch between our entitlement and truth, critical reasoning would not occur among us. Indeed, the entitlement to reflective judgments that derives from those judgments' place in critical reasoning would lapse. But critical reasoning does occur among us; and we are entitled to reflective judgments by virtue of their contribution to the reasonability of critical reasoning. So as critical reasoners we must know our thoughts and attitudes.

Symptomatic of the connection I have noted between the rationality of reflection in critical reasoning and the truth of reflective judgments is the fact there are severe limits on brute errors in judgments about one's present ordinary, accessible propositional attitudes. A brute error is an error that indicates no rational failure and no malfunction in the mistaken individual.⁷ Brute perceptual errors commonly result from misleading natural conditions or look-alike substitutes. One can be perceptually wrong without there

7 I introduced the notion of brute error in 'Individualism and Self-Knowledge' *op. cit.*, p. 657. I intend rational failures to include any failure of entitlement or justification, not just ones that are epistemically culpable. I intend malfunctions to cover not only mechanical or biological failures in, say, the individual's perceptual apparatus, but also failures of normal understanding—as for example when an individual believes arthritis can occur in the thigh. The idea is that a brute error would have occurred even if the individual's epistemic warrants were in order and the individual's perception and ordinary understanding were functioning optimally.

being anything wrong with one. Such brute perceptual errors are unremarkable. But errors about what one's thoughts and attitudes are normally seem to involve some malfunction or rational deficiency. There are exceptions—the cases of unconscious, modular attitudes that are not accessible to reflection. One could easily make brute errors about these. There are judgments about one's emotions, character, or deep motives, that seem hard to get right. I leave open whether these might sometimes involve brute errors. But it seems that we make mistakes about many attitudes that are accessible to reflection primarily when we are subject to some failure of rationality or defect in our cognitive powers.⁸

I stated that I would argue that the specific role of knowledge of our thoughts in critical reasoning requires that it be associated with a distinctive sort of epistemic entitlement that necessitates a non-contingent, rational relation between the relevant first-person present-tense judgments and their subject matter or truth. Why need self-knowledge be in any way special? Why is it not enough that it be pretty reliable observation? Some knowledge of our own mental states and events is empirical in the sense that it is based either on imaging, remembering, or reasoning about *sensed* inner-goings-on, or on observing our own behaviour and hearing about it from others. Simplicity tempts some to hold that all self-knowledge is like that.

Let me elaborate this temptation. It is commonly held that beliefs about others' attitudes must be based on inferences from or criteria for observation.⁹ On the model at issue, beliefs about one's own attitudes differ only in that one need not always infer those beliefs,

8 Our epistemic entitlement to judgments about our present attitudes is a general right and is compatible with our making various mistakes about our attitudes even in the course of critical reasoning. (Of course, then we are, in a sense, not critically reasoning *with* the attitudes we are mistaken about.) We make mistakes of haste, bias, and self-deception. Some attitudes are hard to get at, except with discipline, and even maturation or therapy. In some cases, other people are better at knowing our attitudes than we are. So one might demand further specification of our entitlement. One might ask under what conditions it is overturned or insufficient to give us knowledge. And one might inquire in more depth into the conditions under which errors arise. These issues are complex. I think that when our judgments about a certain class of our thoughts and attitudes are in a certain sense immediate (which entails that they are neither inferred nor otherwise biased by other attitudes), and when our minds are not subject to malfunction, we do not make errors. But there is no recipe for insuring that our judgments are immediate or that they are about the relevant class. There is no internal recipe for avoiding error.

9 I do not accept this view, but I need not question it here. Cf. 'Content Preservation' *op. cit.* Certainly one's beliefs about others' thoughts are often based this way.

because one is the closest witness. There is no authority in self-knowledge, other than the authority of inner observational presence, practice, and familiarity.

This simple observational model does not account plausibly for *cogito*-like thoughts. Such thoughts are logically special in their self-verification and epistemically special in their clear dependence for entitlement on intellection and understanding, not on any sort of observation. But *cogito* thoughts do not constitute the full range of thoughts that enter essentially into critical reasoning.

The simple observational model is encumbered with the obscurity of the notion of inner observation as applied to thoughts and attitudes. Unlike sensations or images, thoughts and attitudes lack distinctive presentations or phenomenologies. The model is phenomenologically implausible for many immediate judgments about one's own beliefs or current thoughts. But I want to show that there is a deeper problem if the model is taken to cover all cases.¹⁰

Before presenting the argument, I will say what I take to be fundamental to the simple observational model. The model need not claim any phenomenological presentation in self-knowledge, though waiving such a claim weakens the analogy to observation. The fundamental claim is that one's epistemic warrant for self-knowledge always rests partly on the existence of a pattern of veridical, but brute, contingent, non-rational relations—which are plausibly always causal relations—between the subject matter (the attitudes under review) and the judgments about the attitudes. This claim is compatible with holding that from the point of view of epistemology, observational judgments are often immediate and non-inferential, requiring no background causal hypothesis on the part of the individual about their source.

¹⁰ Hume is, I think, a proponent of the simple observational model. A more recent proponent is D. M. Armstrong, *A Materialist Theory of the Mind* (London, Routledge & Kegan Paul, 1968), pp. 323–338. The rationalist tradition, in its emphasis on the role of self-knowledge in rationality, and the role of understanding (not sensory observation) in self-knowledge, is the source of my view. Kant develops this tradition in a particularly deep way, although his epistemology left him with what was, in my opinion, an implausibly restrictive account of cognition of one's own thoughts, one indeed that overrates the role of inner sense. A more recent non-observational account that emphasizes the role of self-knowledge in reasoning may be found in Sidney Shoemaker, 'On Knowing One's Own Mind' *Philosophical Perspectives* 2, (1988), pp. 183–209.

My view about perceptual entitlement is more specific than the fundamental claim of the simple observational model. I believe that our entitlement rests partly on our being perceivers, which entails that we—or our species-perceptual systems—are or have been in brute, contingent, non-rational but veridical relations to objects of perception, and the kinds that our perceptual judgments specify.¹¹ It is necessarily constitutive of the content of our observational or perceptual beliefs about physical objects, and of the very nature of our perceptual systems, that we be veridically attuned to the environment through causal relations to it—either in our learning histories or indirectly in the evolution of our perceptual systems.¹² Entitlement to observational physical object beliefs rests partly on

11 A tempting oversimplification is to claim that these constitutive veridical causal relations are always *reliably* veridical. Such a claim is tempting because in so many cases our perceptions are reliable. Perhaps many types of perception must be. But the claim is oversimplified because some perceptual intentional types in some perceptual systems are established through the systems' reliable avoidance of false negatives rather than through their reliable achievement of true positives. It is more critical to a hare's perceptual system that it not fail to register a predator when one is there than that it be reliable in its registration of predators. So the system could commonly indicate the presence of predators falsely—and be broadly unreliable in its perceptions—as long as it was reliable in correctly indicating present predators. It remains, however, constitutive of the systems' perceiving predators as predators that some veridical perceptions played a role in the evolution-fashioned function or in the actual use of the system. Of course, reliability is more important for perceptions of safety than for perceptions of danger. These qualifications on reliability, of course, complicate any account of the relation between perceptual-content constitution and perceptual entitlement. For presumably epistemic entitlements are *prima facie* comprised by constitutively unreliable perceptual deliverances. I believe that a perceptual system in any agent, however, is constitutively associated with reliable perceptions in a range of cases. But these are issues for another occasion.

12 I am inclined to think that it is a *conceptual* necessity that there be *causal* relations in perception. But for purposes of my argument, the fundamental feature is that the entitlement to observational beliefs necessarily rests on some pattern of brute, contingent, non-rational relations between observed and observer, regardless of whether the contingent relations are causal. It is common to my view and the opposed observational view of self-knowledge that in many of the cases under dispute, there is a causal mechanism that relates attitudes to judgments about them. What is in dispute is the nature of the epistemic entitlement that one has to such judgments, not the existence of a psychological mechanism. On the simple observational model, our entitlement to self-knowledge always rests partly on the brute, contingent, non-rational causal relations. On my view, in some important cases, it does not: Christopher Peacocke has pointed out to me that in some, though I think not all, cases of special self-knowledge, the entitlement may specify some causal relation between subject matter and judgment. But not all causal relations are brute, contingent, non-rational ones. (For example those involved in a person's deductive inference are not.) Where a causal relation is not merely a background enabling condition, but an element in the relevant entitlement to self-knowledge, it will on my view never be a brute, contingent, non-rational one. It will be associated in the entitlement with norms for transfer of reasons.

this necessity. But in particular instances of perception, the relations on which one's entitlement to perceptual judgments rests are brute, contingent, non-rational.¹³ It is this claim that forms the paradigm for the simple observational model.

The existence of veridical beliefs involving *de re* causally based relations to the environment is necessary to and constitutive of something's being a perceptual system. But the individual relations are brute, contingent, non-rational. The brute contingency of these relations in individual instances of veridical perceptual judgment is fundamental to observation. Different conditions could have caused a perceptual judgment that was internally-indistinguishable (indeed I think the same judgment-type) but non-veridical, without loss of entitlement, even as the system functioned optimally well.¹⁴ Thus

13 The simple observational model is inspired by a comparison of self-knowledge to observational judgments about physical objects. I operate with a commonsensical conception of such judgments. There are non-common-sensical conceptions that take the basic relation that underlies our epistemic right to be one between the observer and some sort of mental item, a sense datum or an appearance. I ignore such theories not because I regard them as mistaken (although I *do* regard them as mistaken). I ignore them because they model observation of physical objects on knowledge of one's mental events. The model I am attacking proposes to illumine self-knowledge through an independent model. I do think that knowledge of our pains and other sensations—as contrasted with knowledge of our propositional states and events—is empirical in the sense that it depends for its entitlement on sensory experience or sensory beliefs. Judgments that constitute such knowledge just are sensory beliefs. Although I believe that brute error is possible in certain judgments of this sort, such cases are marginal. Understanding even these empirical judgments will, I think, owe more to the kinds of considerations I am elaborating than to reflection on ordinary perceptions of physical objects. But I regard knowledge of one's sensations as requiring separate treatment from knowledge of one's thoughts and attitudes.

It is worth noting that a view that we must 'inferentially' base judgments about physical objects on observations of sense data would also normally be committed to holding that one's entitlement to those judgments rests on brute, contingent, non-rational relations to the physical objects that always allow for brute error. The same point applies to inference-to-the-best-explanation views of our warrants for perceptual beliefs about physical objects.

14 This gloss on the brute *contingency* of the relations, apart from the parenthetical remark, is less committal than my own view of the contingency involved in observational relations. I think the same perceptual object could, with different external auxiliary conditions, have caused a different non-veridical judgment. And I think that a different perceptual object, or perhaps none at all, could have combined with different external auxiliary conditions to cause a perception or perceptual judgment of the same type, though perhaps one with a different token demonstrative element, making it non-veridical. (I do not depend on these views in my argument here.) These different conditions, in individual cases, need not affect the individual's entitlement to the perceptual judgment; nor need they affect the well-functioning of the individual's perceptual-cognitive apparatus. The sense in which the relevant relations are *non-rational* is complex. Perhaps it can suffice here to note that since in the case of ordinary perception the perceptual objects are physical kinds or physical individuals, there can in

perception is always subject to brute error. The object or conditions of perception could lead us into misperception without there being any failure of entitlement and without there being any malfunction of our cognitive or perceptual systems.

The objectivity of perception depends on the possibility of epistemically entitled misperception. Perceptual justification and criticism necessarily presuppose a distinction between a person's cognitive perspective and the objective, physical subject matter. They further presuppose this unremarkable possibility of contingent mismatches in individual cases that in no way impugn the individual's epistemic entitlements or perceptual-cognitive functioning. Rational and epistemic evaluation fixes on the individual's perceptual judgments and perspective, not on their physical subject matter. For this is only brute contingently related, in individual cases, to epistemic entitlement.

A consequence of interpreting all self-knowledge on the simple observational model is that in any given case brute errors—errors that do not reflect on the rationality or sound functioning of the reviewing judgment—are possible. I intimated earlier that brute errors do not seem to threaten some instances of judgments about attitudes. I propose to show why this must be so.

Not all one's knowledge of one's propositional attitudes can fit the simple observational model. For general application of the model is incompatible with the function of knowledge of one's own attitudes in critical reasoning. The main idea is that such application would entail a dissociation between cognitive review and the thoughts reviewed that is incompatible with norms of epistemic reasonability that are basic to all critical inquiry, including empirical, mathematical, philosophical, and practical inquiry.

Rational evaluation of attitudes commonly applies to and within a perspective or point of view. The argument will make reference to this fact. Different people have different points of view. My judgment that your beliefs are irrational may be reasonable from my point of view. But it does not follow that there is reason from your perspective to change your beliefs. I may have made some brute

that case be no question of a rational relation between them—which have no intentional content at all—and perceptions or perceptual judgments.

error about what your beliefs are, or your perspective may have different associated reasons or background information from mine.

There can be different perspectives or points of view within a given person. What is reasonable for a person at a given time may be different from what is reasonable from the perspective of the person's memory back on that time. What is reasonable on reflection may differ from what is reasonable in modular cognitive processes, or in an instant practical reasoning, or in subconscious reasoning. My argument hinges on how reasons transfer across points of view.

Suppose that all one's knowledge of one's propositional mental events and states fit the simple observational model. Then one's entitlement to instances of such knowledge would always rest on purely contingent relations between any given judgment about one's mental states and the subject matter of the judgment. What is more, brute error would be possible in any given case. Normative evaluations of reasonability and epistemic entitlement in critical reasoning—in checking and evaluating one's reasoning—would apply within the perspective of the judgments, but not immediately within the perspective of the subject matter of the judgments, except insofar as it contingently conformed to those judgments, and except insofar as it happened to be embedded in a perspective relevantly similar to the perspective from which the judgments were made. For the subject matter might, in any given case, fail to conform to the judgments through no failure of justification or entitlement in the judge, and through no malfunction of the relevant faculties.

But this picture is nonsense if it is applied to all judgments about one's own propositional attitudes. For it is constitutive of critical reasoning that if the reasons or assumptions being reviewed are justifiably found wanting by the reviewer, it *rationaly follows immediately* that there is prima facie reason for changing or supplementing them, where this reason applies within the point of view of the reviewed material (not just within the reviewing perspective). If the relation between the reviewing point of view and the reasons or assumptions being reviewed always fit the simple observational model, there would never be an immediate rationally necessary connection between justified rational evaluation within the review, on one hand, and its being prima facie reasonable within the reviewed perspective to shape attitudes in accord with that evaluation,

on the other. For the relation between the perspective of the review and that of the reviewed attitudes would always be purely contingent, even under canonical descriptions of them, for purposes of rational evaluation. (The attitudes reviewed would be to the reviews as physical objects are to our observational judgments. They would be purely 'objects' of one's inquiry, not part of the perspective of the inquiry.) It would be reasonable for the person from the point of view of the review that a change in the reviewed material be made. But this reason would not necessarily transfer to within the point of view of the attitudes under review, even though that is a point of view of the same person. Its transferring would depend on brute, contingent, non-rational relations between the two points of view.

In critical reasoning, however, the connection is rationally immediate and necessary. Justifiably finding one's reasons invalid or one's thoughts unjustified, is normally *in itself* a paradigmatic reason, from the point of view of the thoughts being reviewed (as well as from the perspective of the review), to alter them.

If in the course of critical reasoning I reasonably conclude that my belief that a given person is guilty rests entirely on unreasonable premises or bad reasoning, then it normally follows immediately both for the perspective of the review and for the perspective of the reviewed belief that it is reasonable to give up my belief about guilt or look for new grounds for it. In such second-order reasoning, I am not normally reasonable in altering my first-order views about guilt or innocence only with the proviso that they are embedded in reasons that contingently match those associated with my reviewing perspective. I do not normally have the sort of excusing condition that allows for rational error that hinges on the contingent relation that the subject matter bears to my judgments about it. Rather my checking my belief and finding it wanting normally *itself* provides immediate prima facie reason to change it from within the perspective of the review. This is because the first- and second-order perspectives are the *same* point of view.

The reviewing of reasons that is integral to critical reasoning includes the review and the reviewed attitudes in a single point of view. The simple observational model treats the review and the system being reviewed as dissociated in a way incompatible with the norms of critical reasoning. It makes the reviewed system an object of investigation, but not part of the investigation's point of

view. So the model fails to account for the norms of critical reasoning.

A closely related point centres on epistemic responsibility. We are epistemically responsible only because we are capable of reviewing our reasons and reasoning. And we are paradigmatically responsible for our reasons when we check and review them in the course of critical reasoning. But the simple observational model implies that in carrying out reviews of one's reasoning, one is epistemically responsible not primarily for the thoughts being reviewed but primarily for the review. The model implies that we are in reviewing our reasons only derivatively responsible for objects of review, as one might be responsible for the actions of one's child or dog—but fully and primarily responsible only where one's knowledge and control contingently matched what one is justified in believing about them.

But one is not epistemically responsible for the thoughts one reflects upon in critical reasoning in the way one is responsible for something one owns or parents. One's responsibility in reflecting on one's thoughts is immediately for the whole point of view. The simple observational model fails to account for the fact that critical reasoning is carried out within a single multi-level point of view.

Of course, we are sometimes disunified. Sometimes to our own good and efficiency, sometimes to our misfortune, we fail to know our motives or reasons, or know them only through observation and empirical reasoning. Sometimes from the point of view of our self-conscious reviewing selves, we are indeed epistemically responsible only derivatively for attitudes that we know only empirically. But in these cases, we are not reasoning critically with those aspects of ourselves that we know only in these ways.

Theoretical knowledge of one's modular attitudes is one sort of purely observationally-based 'self-knowledge'. When attitudes *cannot* be known in a direct non-observational way, one commonly enters some qualification on the sense in which the attitudes are *one's own*. In these cases failure to know the attitudes non-observationally is no sign of dissociation. But when one knows only observationally unconscious attitudes which are in principle accessible to non-observational self-knowledge, there is some dissociation of self, constituted by a divide between the point of view of one's critical reasoning and the attitude known only observationally.

Knowledge through therapy of one's unconscious, before full integration of that knowledge, provides one sort of example. One may know the attitudes on the basis of observationally based therapy, but the unconscious attitudes may provide a point of view of their own into which the meta-evaluations of them may not transfer. Those evaluations may not speak to unconscious considerations that are integral to the unconscious pathology; or the unconscious point of view may not have 'taken in' matters that are integral to the rationality of the meta-, therapeutic point of view.

Psychoanalytic cases are not the only sort that illustrate the relevant dissociation. One may know from experience or theory that one will act a certain way, and yet rational meta-considerations may not penetrate to the system of attitudes that motivate the action. Some self-admitted compulsions provide examples. One knows one has or will have the relevant intention, but knows the intention only as object; it is then not the product of critical deliberation. Sometimes rational considerations from the meta-point of view may not have the same rational force and relevance within the point of view that includes the observationally known attitudes.

There are cases of knowledge of one's beliefs like this as well. One may know from self-observation that 'underneath' one believes something because one needs to believe it, while feeling sincere rational urges to assert the contrary. The system of underlying practical beliefs that motivate the needed belief may form a point of view that does not recognize as sufficient the rational meta-reasons that one can offer oneself for giving up the need-based belief. The person's meta-perspective may correctly condemn the need-based belief as epistemically irrational. But the belief may be dissociated from the point of view of his observational knowledge of that belief. It may be rational 'overall' for the person to give up the belief, but the practical rationality of the limited, need-driven perspective may exclude or outweigh the considerations that count against the belief.

The relevant psychological dissociation is, I think, sometimes partly to be explained in terms of the fact that a second 'point of view', or system of attitudes with its own internal coherence, has gotten set up within the person, in such a way that reasons from the point of view of the person's critical rationality do not automatically transfer to within the second point of view, rather as reasons from

my perspective do not always apply as reasons from yours. Being known from the perspective of the critical reasoner *only* as an object, on the basis of observation and theory, is sufficient for an attitude to be dissociated in this way. Attitudes that are part of such a dissociated point of view may provide us with reasons, even operative ones, for doing things. But insofar as we know them only observationally, they are not part of our critical reasoning.

Where we know our thoughts or attitudes only by observation, the question of *means* of control—of effective application of reasons to them—arises, at least from the perspective of our observational knowledge. Where we know our attitudes only as empirical objects, not only are our rational evaluations of those attitudes relativized to contingencies associated with the knowledge. But our ability to apply our reasons (those associated with the point of view of the observational knowledge) must acknowledge the contingency of our rational control over those attitudes. We must face a question of how, by what means, to make those reasons effective in view of the contingent relation between the point of view of the self-knowledge and rational evaluation, on one hand, and the observationally known attitudes, on the other. Again, this is not critical reasoning. In critical reasoning, such questions of means and control do not arise, since one's relation to the known attitudes is rationally immediate: they are part of the perspective of the review itself.

Sometimes observation-based self-knowledge enables one to assimilate an attitude into one's critical point of view, and to take direct critical control and responsibility for the attitude. I may learn through observing my behaviour or through reasoning in therapy that I believe that a friend is untrustworthy. I may 'internalize' this belief so that it is no longer merely an object of observational knowledge. This process is sometimes immediate, sometimes hard, requiring deep personal change.

Although much reasoning and rational attitude formation occurs outside the purview of critical self-knowledge, or indeed any self-knowledge, critical reasoning remains central to our identities as persons. So no reasonable account of self-knowledge can ignore the role and entitlements critical reasoning gives to self-knowledge.

The argument I have given against the simple observational model indicates that the relations between knowledge and subject matter on which one's entitlement rests cannot always be causally

brute, contingent, non-rational ones. In some particular instances of self-knowledge, the connection must be a rational one. For conclusions about the reasonability of one's thoughts based on self-review *directly* yield reasons within the point of view of those reviewed thoughts to alter or confirm them. The relation between self-knowledge and subject matter is that they must normally and rationally *be part of the same theoretical and practical point of view*—elements of a single theory or plan.

Connections between reviews and thoughts under review that are fully open to reason and that allow immediate transmission of reasons are necessary to the rational coherence of a point of view. A merely observation-based relation between attitudes insures that they are parts of different points of view. Indeed, it is constitutive of a point of view that failure to follow or understand its connections by the holder of the point of view is a failure of rationality. Thus a point of view is not closed under deductive consequence, and contains rational connections other than deductive ones.

So entitlement to knowledge of one's own thoughts and attitudes is not purely a matter of what one does. It has to do with who one is. One's status as a person and critical reasoner entails epistemic entitlement to some judgments about one's propositional attitudes. It entails some non-observational knowledge of them.

Cogito-like thoughts illustrate one non-contingent rational relation between knowledge and subject matter. In those cases, the reviewed thought is simply a logical part of the review. But *cogito*-like thoughts are in many ways special cases. If we are to understand critical reasoning, the entitlement that I have discussed must apply more broadly. It must include judgments about beliefs, intentions, wants, as well as occurrent thoughts.

So far, I have put little weight on the first-person present tense form of the relevant pieces of self-knowledge. Clearly, for the review and the reviewed thoughts to be part of a practice of critical reasoning, the reviewed thoughts must be capable of becoming part of the reasoner's present array of attitudes. And the special features of *cogito* cases do depend on present tense. But much of what I have said about the dependence of an entitlement on its role in critical reasoning, and about the non-observational character of this entitlement, applies to preservative memory—that type of memory that preserves propositions and our commitments to them in

reasoning.¹⁵ I think that this sort of memory provides us with some non-observational knowledge of our past mental states and events, and is epistemically underwritten by its role in critical reasoning.

The first-person point of view is clearly basic to self-knowledge in critical reasoning. The self-knowledge that I featured differs from observational knowledge of physical objects in that the first-person point of view is deeply relevant to the epistemic status of the knowledge. In observations of physical objects, anyone could have made substantially the same observation with equal right, if the same angle of perception had been available at the same time. But self-ascriptions constitute an epistemic angle in themselves.

What does this metaphor come to? If the reviewing knowledge is to be integral to critical reasoning, if it is to provide immediate rational ground for change in the reviewed material, the review must take up the same perspective or point of view as the act under review—the reasoner's own object-level point of view. The first-person point of view bears a distinctive relation to the relevance of rational norms to rational activity. For a review of a propositional mental event or state to yield an immediate rational ground to defend or alter the attitude, the point of view of the review and that of the attitude reviewed must be the same and must be first-personal.

In evaluating reasoning critically, one must make commitments to attitudes partly on the basis of critical evaluations of them. If one is to fully articulate the rational basis for the application of rational norms within critical reasoning, the commitments to both reviews and reviewed attitudes must be conceptually acknowledged as one's own. For acknowledging them as one's own is taking them as attitudes that one could rationally and directly change or confirm. Acknowledging first-order attitudes as one's own is necessary to articulating the direct rational relevance of one's critical reasons to first-order reasoning (or more generally, reasonable activity). I intend to say more about this matter on another occasion.

I have sketched the environmental neutrality of our entitlement to self-knowledge. The entitlement remains constant under possible unnoticeable variations in environmental circumstances or cognitive content. For it does not depend on the empirical content

15 Cf. 'Content Preservation', *op. cit.*

of the judgments. It does not depend on checking whether our judgments meet certain conditions. It depends on the judgments' being instances of a kind essential to critical reasoning. Critical reasoning presupposes that people are entitled to such judgments. Since we are critical reasoners, we are so entitled.

Epistemic entitlement derives from jurisdiction—from the place of the judgments in reasoning. In *cogito*-like thoughts, this place is coded in the content of the judgments themselves. In other relevant sorts of self-knowledge, which are fallible, the entitlement, indeed one's knowledge, depends only on one's not misusing the judgments and on one's remaining a sane critical reasoner.¹⁶

16 Versions of this paper were given as the third of six Locke Lectures at Oxford in 1993, and the first of two Whitehead Lectures at Harvard in 1994, as well as on several other occasions. The key idea and first draft of the paper dates from 1985. I am grateful for helpful comments on drafts or talks based on this paper to Robert Adams, Kent Bach, Tony Brueckner, Phil Clark, David Kaplan, Christopher Peacocke, Marleen Rozemond, Hilary Putnam, Nathan Salmon, Houston Smit, Barry Stroud, Patrick Suppes, and Corliss Swain.