

The Rationality of Deductive Inference

Declan Smithies

NYU/Warwick
dgs224@nyu.edu

1. The Question

Deductive inference is one of the ways in which it can be rational to form beliefs. What makes it the case that it is rational to form beliefs in this way? And in particular, what is the role of *consciousness* in explaining why it can be rational to form beliefs in this way?

2. Fodor's Language of Thought Model

I want to approach this question by considering the Fodorian idea that the language of thought hypothesis explains the rationality of deductive inference. More precisely, I want to consider the claim that the language of thought hypothesis provides a *constitutive* explanation of the rationality of deductive inference, rather than merely a causal explanation in terms of "enabling conditions".¹

The explanation takes the following shape: (1) beliefs (and other propositional attitudes) are identified with physical structures that are causally efficacious; (2) deductive inference consists in the occurrence of causal transitions between these physical structures; (3) the causal transitions between them are sensitive only to their formal (i.e. non-semantic) properties; (4) there is an isomorphism between the formal properties of these physical structures and the semantic properties of their propositional contents; (5) the causal transitions between these physical structures are sensitive to their formal properties in such a way as to preserve relations of truth-functional consequence between their propositional contents.² Fodor summarizes this line of explanation in the following passage:

We can therefore build machines, that have...the following property:

The operations of the machine consist entirely of transformations of symbols;

In the course of performing these operations, the machine is sensitive solely to the syntactic properties of the symbols;

¹ In fact, Fodor himself maintains what he calls a "meretricious metaphysical neutrality" in his writings (e.g. 1998, Ch.1); but with respect to adopting the stronger, constitutive commitment, he is prepared to say, "...you could do so if you were so inclined" (1987, p.156 fn.6).

² Fodor (1987) does not make it clear why an explanation of the rationality of deductive inference requires the language of thought (LOT) hypothesis, rather than merely what Fodor calls the representational theory of mind (RTM). (The issue is whether the formal properties of the relevant physical structures are required to have constituent syntax.) However, in Fodor and Pylyshyn (1988), there is an argument that LOT is required to explain the systematicity of deductive inference.

And the operations that the machine performs on the symbols are entirely confined to altering their shapes

Yet the machine is so devised that it will transform one symbol into another if and only if the propositions expressed by the symbols that are so transformed stand in certain semantic relations – e.g. the relation that the premises bear to the conclusion in a valid argument. (Fodor 1987, p.19)

Intuitively, however, we could build one of Fodor's machines without thereby building a rational thinker capable of engaging in genuine deductive inference. If this intuition is right, then the proposed conditions are not sufficient to explain the rationality of deductive inference. What, then, does this account leave out?³

3. Searle's Chinese Room

This question is closely related to the one that Searle considers in his famous article, "Minds, Brains and Programs". Searle's question is whether purely computational conditions (of the kind that Fodor articulates) are sufficient for genuine *intelligence*, which we can gloss in terms of rational thought. The conclusion that he reaches on the basis of the Chinese Room thought-experiment is that purely computational conditions are not sufficient: the man in the Chinese Room satisfies all the relevant computational conditions, but he fails to understand the meanings of the Chinese symbols he is processing, and so the processing is not genuinely intelligent.

Searle's diagnosis is that what is lacking in this case of the Chinese Room, as well as in many other cases of so-called "artificial intelligence", is the presence of genuine intentional content:

Because programs are defined purely formally or syntactically, and because minds have an intrinsic mental content, it follows immediately that the program by itself cannot constitute the mind.

In later work (Searle 1990, 1992), he goes on to fill out this conception by arguing (1) that genuine, as opposed to merely "as if", intentionality requires consciousness, in the sense that every genuinely intentional state must be available to consciousness; and (2) that the satisfaction of purely computational conditions does not suffice for availability to consciousness. It follows from these claims that the satisfaction of purely computational conditions does not suffice for genuine intentionality. And given that genuine intentionality is required for genuine intelligence, it follows that the satisfaction of purely computational conditions does not suffice for genuine intelligence.

A number of comments are relevant here. First, it is not at all clear that all the relevant computational conditions have in fact been met in Searle's case. For instance, he does not consider the point that genuine intelligence must be structured and systematic in the right kind of way.⁴ As he describes

³ It is worth noting that Fodor's proposed explanation is reliabilist to the extent that he takes it to be sufficient for an explanation of the rationality of deductive inference that one should explain its distinctive kind of reliability, i.e. its validity.

⁴ Compare Block (1980).

the case, the competence of the Chinese Room is due to its occupant's use of a mere look-up table. Suppose, then, that we adjust the case so that various states of the system make a systematic causal contribution to the behavioural responses exhibited by the Chinese Room. Now we can challenge Searle's contention that genuinely intentional states must be available to phenomenal consciousness. For instance, we can argue that intentional contents need to be ascribed to the states of the system in order to explain the relational properties of its behavioural responses, regardless of whether or not those states are available to consciousness.⁵ On this line, the ascription of intentional content to states which are not available to phenomenal consciousness need not be a purely instrumentalist matter, but may be answerable to perfectly objective constraints on explanation.

It seems to me that Searle has misdiagnosed the intuitive significance of his own case. What is intuitively lacking in the Chinese Room is not intentionality *per se*, but more specifically *understanding* of the Chinese symbols being processed. The claim that understanding requires consciousness (whether we consider linguistic understanding or merely nonlinguistic conceptual understanding) is much more plausible than Searle's own claim that intentionality requires consciousness. After all, there is a compelling intuition that a mere machine without any conscious states could not be credited with any kind of linguistic or conceptual understanding, even if it could be attributed intentional states of some other kind. But some kind of linguistic or conceptual understanding is surely a prerequisite for forming a belief rationally on the basis of deductive inference. Since the Fodorian account of the rationality of deductive inference is silent on the issue of consciousness, this suggests the beginning of an answer to the question of what that account leaves out. However, it is no more than a beginning: the claim that understanding requires consciousness needs both a sharper formulation and a principled explanation.

4. Brewer's Non-Reflective Access Requirement

In the following passage, Brewer makes an intriguing suggestion about *why* consciousness might be a requirement for rational deductive inference:

Any purely mechanical dispositional account of the matter is unacceptable. To start to see why, consider what is involved in following a valid deductive argument with real understanding, in a way that yields knowledge of its conclusion.... In cases where I do not go wrong, I am correctly compelled by the argument to believe its conclusion. In following and fully understanding the argument, this compulsion is not simply a blind and mysterious manipulation of my beliefs by some reliable mechanism, however well established by evolution, benevolent hypnosis, or whatever. I am not just a machine which runs along those rails. For if my following the argument is really to extend my knowledge, then my understanding of it must give me some appreciation of why I am right in believing its conclusion. I have to have some grip on how I

⁵ Compare Peacocke (1994).

thereby know the conclusion, and my belief should be guided by this understanding. A disposition to take beliefs on board in parallel with the steps of the argument is on its own insufficient for the argument to provide me with genuine knowledge. For such beliefs would come as a succession of mere hunches, wholly unsubstantiated *for me* by the de facto validity of the argument propelling my endorsement of them. (Brewer 1995, p.242)

The suggestion is that forming a belief on the basis of deductive inference in a way that is rational, and thus a potential means for extending one's knowledge, requires *understanding* the argument that one is following. Understanding the argument, in turn, involves an appreciation of how the conclusion follows from the premises, and hence an appreciation of why it would be rational to believe the conclusion, given belief in the premises. And since such an appreciation requires consciousness, it follows that genuinely rational deductive inference requires consciousness.

More specifically, Brewer proposes that for a belief to be rationally formed on the basis of deductive inference, the formation of the belief must actually be *guided* by a conscious appreciation of why it would be rational to form it, which is provided by an understanding of the argument in question. If the belief is not formed on the basis of any such appreciation, then the suggestion is that it will seem, from the subject's own perspective, to be no different from blind guesswork or an ungrounded hunch. This, in turn, prompts an intuition that it can be nothing more than a matter of mere luck, from the subject's perspective, if the belief is true. But then, of course, the belief cannot be rational. This point emerges in the following passage:

More generally, the problem for any purely mechanical dispositional account is that it is bound to ignore this sense of why one is right in exercising the capacity in question as one does, which is crucial if this is to make *cognitive* contact with the truth in the relevant area. It is the fact that the capacity has as its point ascertaining the truth on some matter...which sets up the norms for its correct exercise. Some appreciation of how what one is up to is onto this truth, is sensitive to the resultant norms, is therefore essential if exercising it is to be more than a blind mirroring of the norms, extrinsically, and in this sense only incidentally, in contact with the truth. (1995, p.243)

So, in general, if a belief is to be rational, then it must be formed in a way that is guided by some conscious appreciation of why it would be rational to form it.

It is on the basis of such considerations that, in later work, Brewer imposes the following generalized *access* requirement on rational belief and action:

If a person's reasons are to be cited as *her reasons* for believing or doing what she does, then she necessarily recognizes them as such. (1999, p.166)

And moreover:

...she is guided in making the transition by her recognition of her reason *as a reason for doing so*. (1999, p.165)

The difficulty here is to see how to avoid a familiar kind of regress problem. In fact, there are two kinds of regress problem to be avoided. Suppose that the way in which the subject recognizes her reason-giving states as such is by forming second-order beliefs about them to the effect that they give her reasons for belief or action. Now, these second-order beliefs must themselves be formed in a way that is rational. But then the reason-giving states on which these second-order beliefs are based must themselves be recognized as such, which requires the subject to form third-order beliefs about these states to the effect that they give her reasons. And these third-order beliefs, too, must be rational. So, we are embarked on an infinite regress. Perhaps the regress is not strictly vicious, but it does have the vastly implausible consequence that the rationality of a single belief depends on an infinite hierarchy of higher-order rational beliefs.

However, we get a regress which is more clearly vicious if the only plausible account of the rationality of higher-order beliefs makes it dependent on the rationality of beliefs lower down the hierarchy. But certainly in the logical case, this can seem to be the only plausible account. For we might suppose that recognizing the rationality of a belief based on modus ponens requires recognizing that modus ponens is a valid rule of inference. And we are supposing that this recognition must be a matter of rationally formed belief. But the rationality of a belief that modus ponens is valid must surely depend on making a modus ponens step, as in the following piece of reasoning:

An argument is valid just in case, necessarily, if its premises are true, then its conclusion is true. Now, suppose that p is true and that *if p , then q* is true. Well then q must be true, as a matter of necessity, whatever p and q stand for. So, an argument of the form – p ; *if p , then q* ; *therefore, q* – is valid.

However, the rationality of making such a modus ponens step is precisely what we are trying to explain; so it can seem that any purported explanation along these lines is guaranteed to presuppose what it is trying to explain.

Brewer's strategy for avoiding such regress problems is to propose a distinctive account of the way in which the access requirement is satisfied. According to his proposal, the access requirement need not be satisfied at the *reflective* level by means of higher-order beliefs about one's reason-giving states; rather, it can be satisfied at the unreflective level by means of one's reason-giving states themselves. On this view, it is an essential feature of reason-giving states that being in them suffices for satisfaction of the access requirement; no further, reflective thoughts about those reason-giving states are required. So, his version of the access requirement is a *non-reflective* access requirement.

Unfortunately, it is not clear how to make sense of the claim that reason-giving states embody a built-in conscious recognition *of their own reason-giving status*. This seems to require a kind of self-referentiality in the contents of reason-giving states which is obscure, at best, and possibly even incoherent.

In fairness, Brewer's (1999, Ch.6) account of the rationality of perceptual-demonstrative judgements suggests that he has something less demanding in mind. According to this account, perceptual experiences

provide reasons for perceptual judgements about the external world because they embody an appreciation on the part of the subject that the experience is causally dependent on the state of the external world in relevant respects. It is in virtue of embodying such an appreciation that perceptual experiences provide reasons for perceptual judgements. It could be objected that this does not suffice, strictly speaking, for the claim that perceptual experiences embody an appreciation of the fact that *they themselves provide reasons* for perceptual judgements. But it could equally be replied that, loosely speaking, it suffices for the claim that perceptual experiences embody a sense of “why it would be right” to make perceptual judgements on their basis.

The analogous claim in the case of deductive inference would be that a thinker’s understanding of the premises and the conclusion of an argument embody a conscious appreciation of the fact that the conclusion follows from the premises. Loosely speaking, this can be described as a sense of “why it would be right” to believe the conclusion of the argument on the basis of believing its premises, even in the absence of an explicit recognition of the fact that my belief in the premises provides me with a reason to believe the conclusion. The proposal, then, would be that the rationality of deductive inference depends on the fact that it is made on the basis of a conscious appreciation of the fact that the conclusion of the argument follows from its premises. This proposal has recently been endorsed by Christopher Peacocke, who expresses it as follows:

How can we elucidate the rationality of the thinker’s judgements? One intuitive account is that in making a rational transition to a judgement that p [on the basis of deductive inference] a thinker must know what it is for it to be true that p , must appreciate that his grounds or reasons for the transition to the conclusion that p suffice for the truth of p , and must be making the judgement because of his appreciation that these grounds or reasons so suffice. (2003, p.176)

However, neither Peacocke nor Brewer says much about how the relevant kind of conscious appreciation might be constituted so as to play this role in grounding the rationality of deductive inference. But this issue turns out to be crucial in the evaluation of the proposal.

5. Boghossian’s Challenge

Suppose that appreciating the validity of a form of argument is a matter of explicitly articulating the form of the argument and rationally judging it to be valid. If this is what it is involved, then there seem to be compelling objections against the claim that a deductive inference is rational only if the subject makes the inference on the basis of some such appreciation of its logical validity. The first objection is that the requisite conceptual capacities are lacked by most normal adults, not to mention animals and children; but it would be intuitively implausible to deny on this basis that they have any capacity for rational deductive inference. The second objection we have already considered, in effect. If we suppose that the rationality of judgements about the validity of forms of argument itself depends on rational deductive inference, then it follows that the rationality of deductive inference cannot

depend on an appreciation of the validity of a form of argument; rather, the direction of dependence must be the reverse.⁶

One line of response to these objections would be to claim that we have some kind of direct and non-inferential, quasi-perceptual access to the validity of forms of argument. The claim would be that such a quasi-perceptual faculty of “rational intuition” is possible in the absence of a sophisticated capacity for articulating general forms of argument and using the resources of deductive inference to evaluate their validity. Moreover, it could be claimed (by analogy with the perceptual case) that judgements about validity can be non-inferentially justified on the basis of the deliverances of such a faculty, without requiring any additional inferential backing.

However, there are compelling objections against assimilating our knowledge of logical validity too closely to the perceptual paradigm. The first is that perception of an object or an instantiated property requires the holding of a causal relation between the perceiving subject and the object or instantiated property. Valid forms of argument, however, are abstract objects, and as such, are incapable of entering into causal relations. So, there is no coherent notion of perceiving the validity of a form of argument. The second objection is that perception alone can only yield knowledge of contingent truths.⁷ But knowledge of logical validity is knowledge of necessity. Therefore perception alone is incapable of explaining this modal dimension of our logical knowledge. The upshot is that the notion of *genuinely* perceptual access to the validity of forms of argument is incoherent; but without further elaboration, the notion of *quasi*-perceptual access is merely obscure.

In addition, Paul Boghossian (2001, 2003) has argued that the appeal to rational intuition cannot do the work required of it. Suppose we grant that deductive inference is not required for forming a rational judgement to the effect that a certain form of argument is valid; still, he argues, it is required in order to bring any such judgement to bear on the rationality of a particular inference. So, if we judge on the basis of rational intuition that any inference of the modus ponens (MPP) form is valid, then we still need to make the following inference:

- (i) Any inference of the form MPP is valid.
- (ii) This particular inference, from (1) and (2) to (3) is of MPP form.

Therefore,

- (iii) This particular inference from (1) and (2) to (3) is valid. (2003, p.233)

A similar objection applies if we grant that what rational intuition provides is not direct access to the validity of a general form of argument, but rather to the validity of a particular instance of that form.⁸ The thought is that inference

⁶ See Boghossian (2001), (2003) for a statement of these objections.

⁷ See Peacocke (2003) for an account of a posteriori knowledge of necessity on which it decomposes into a priori knowledge of necessity combined with a posteriori knowledge of contingency.

⁸ The suggestion was made, in response to Boghossian (2001), by Wright (2001).

is still required in order to bring the judgement to bear on the rationality of making a deductive inference in the particular case, along the following lines:

- (i) This particular inference from (1) and (2) to (3) is valid.
- (ii) If an inference is valid, then anyone who is justified in believing its premises and knows of its validity is justified in inferring its conclusion.

Therefore,

- (iii) Anyone who is justified in believing the premises of the argument is justified in inferring its conclusion.
- (iv) I am justified in believing the premises (1) and (2).

Therefore,

- (v) I am justified in inferring (3). (2003, p.234)

6. Dummett on the Perception of Semantic Patterns

Despite the force of Boghossian's objections, it seems to me that rational intuition is both a genuine phenomenon and one to which we can usefully appeal in giving substance to the Brewer/Peacocke account of the rationality of deductive inference. In arguing the point, I will draw on some illuminating remarks made by Michael Dummett on the matter.

In the following passage, Dummett draws a distinction between *understanding* a proof and merely checking that it is correct:

...an understanding of a proof demands more than an ability to recognize that it is correct. To verify that every line of a formal proof follows from earlier lines by one of a list of transformation rules is to be convinced, within the limits of human error, that it is correct; but it takes one very little way towards understanding the proof. The proof has an architecture that must be comprehended as a whole; but the first necessity for gaining such comprehension is to be intuitively convinced, for each step, that it genuinely follows from the earlier lines from which it was derived. On Frege's account, this will in general require a creative act. It is not enough merely to grasp the thought expressed by each line of the proof; in addition, one must perceive patterns common to those thoughts and others, patterns which are not given with the thoughts as a condition for grasping them but which require a further insight to apprehend. (1991, p.197-8)

To understand the proof, it is not enough to understand each line of the proof and to have checked that each line follows from the earlier lines by application of the inference rules. In addition, it is required that one should *perceive certain patterns* which hold between the propositions or thoughts expressed by the various lines of the proof. This "perception of pattern" is not something which is built into the requirements for understanding the various lines of the proof, but requires an additional, creative insight. Thus, on Dummett's account, *following* a proof is as much a creative matter as *discovering* a proof.

Dummett contrasts the case of understanding a proof with the imaginary case of an obedient community under the governance of an International Academy of Logic, a body which lays down decrees concerning which logical laws are, until further notice, to be treated as valid or invalid. He observes:

...we could obey these decrees; but we should lose the sense that we any longer understood what we were saying. The *rules* of the language-game would be clear enough; but its *point* would now escape us. (1991, p.207)

Presumably, what would be lacking in such a community is not merely the capacity for perception of pattern – its members could be perfectly capable of perceiving the *formal* patterns holding between the lines of a proof – but rather the capacity for perception of *semantic* patterns: for example, that the truth of the conclusion of an argument follows from the truth of its premises. The idea seems to be that in the absence of a capacity for representing such semantic relations, the subject cannot have any sense of the *point* of his inferential practices: he will be merely following the rules without understanding *why* he is following them.

It is here that an account of the rationality of deductive inference links up with an account of the thinker's understanding of logical concepts, which consists in knowledge of their contribution to truth-conditions. After all, it is presumably in virtue of the thinker's knowledge of the contribution made to truth-conditions by logical concepts that he is capable of appreciating the semantic properties of logical forms, such as the validity of a form of argument. Dummett's idea seems to be that what is lacking in his imaginary community is precisely the knowledge of reference in which the understanding of logical concepts consists:

...the missing component of understanding is not to be stigmatized as a 'mere' feeling.... What he lacks is not the *feeling* of understanding, but the *knowledge* that is an essential component of understanding. It is that knowledge that we should lack if we were compelled to reason in accordance with principles that appeared to us invalid or gratuitously restricted: we could rightly confess that we no longer knew what we were saying. (1991, p.208)

7. Meeting Boghossian's Challenge

In summary, Dummett's idea is that understanding a proof requires perceiving the semantic patterns holding between the different lines of the proof or, in other words, perceiving the validity of the proof. But how are we to understand talk of "perceiving" in this context? We have already seen compelling objections against literal interpretations of such talk, but perhaps it can be understood metaphorically.⁹ It seems to me that the perceptual metaphors suggest themselves because there are clear similarities between discerning the patterns in a proof and a certain kind of perceptual phenomenon. Indeed,

⁹ After all, we use perceptual vocabulary in a wide variety of settings that do not involve sensory perception, for example: seeing that a certain course of action is the one to take, seeing that one has made a mistake in a proof, and so on.

it is plausibly one and the same kind of phenomenon which occurs both in the context of perception and in the context of pure thought. This is the phenomenon that Wittgenstein labeled “*seeing as*” in the second part of the *Philosophical Investigations*. Just as one can be struck by certain spatial relations between objects in the perceived array, so one can be struck by certain semantic relations between thoughts in occurrent consciousness. In some cases, this may involve a kind of conceptual structuring of the contents of conscious experience, but the basic phenomenon is pre-conceptual – it is a matter of directing one’s attention or having one’s attention directed in such a way that the relations in question become salient in consciousness.

This basic attentional phenomenon has a number of features which make the perceptual metaphors particularly apposite. First, it involves a characteristic kind of *passivity*, which is given expression in the idea of being *struck* by the relations in question. Second, and relatedly, being struck by such relations is not something which stands in need of reasons or justification, any more than does the passive reception of perceptual experience itself. Third, the phenomenon is belief-independent. Just as we are subject to cognitively impenetrable perceptual illusions, such as the Muller-Lyer illusion, which may persist in the phenomenology of experience even when we know about the illusion, so we are sometimes subject to illusions of validity which persist in the same way even when we know about them.¹⁰

These points bring out a sense in which the analogy between rational intuition and perception need not be obscure, but can actually be quite illuminating. They also defuse Boghossian’s threat of the reappearance of Carrollian circularity. To the extent that the representations of logical validity that we are considering are *belief-independent*, it cannot be assumed that *inference* must be involved in bringing them to bear on the rationality or justifiability of a particular inference. On the contrary, it can be argued that the subject’s conscious attention to the fact that the conclusion of an argument follows from its premises may be directly causally and rationally implicated in the fact that he makes the inferential transition from a belief in the premises to a belief in the conclusion.

So, despite Boghossian’s challenge, the notion of rational intuition may have an important role to play in an account of the rationality of making a particular deductive inference. In my view, it certainly plays an indispensable role in giving a plausible account of propositional knowledge of logical truths, including knowledge of which forms of argument are logically valid. Consider what is involved in the rationality of judging that a particular argument of a certain form is valid. On one view, which we can call the *deductive* model, this requires giving an explicit articulation of the general form of the argument in question and inferring its validity from more general principles concerning what it is for a form of argument to be valid. But there is an alternative view, which we can call the *rational intuition* model, on which judgements of validity in the particular case have epistemological priority and can be rationally based directly on the deliverances of rational intuition. On this view, one’s conscious awareness of the fact that the conclusion of an argument follows

¹⁰ Ayers (1991) gives the example of Zeno’s paradoxes in discussing illusions of validity.

from its premises, which is explained by one's understanding of the logical concepts that figure in the argument, provides one with a non-inferential rational basis for judging that the argument is logically valid.

The comparison with the epistemology of moral and epistemic judgements is instructive. If we consider what is involved in the rationality of judging that a particular case exemplifies knowledge, or goodness, then our options are structurally similar. On the deductive model, it requires giving an explicit articulation of the conditions that are necessary and sufficient for the exemplification of knowledge, or goodness, and showing that these conditions are satisfied in the particular case in question.¹¹ On the rational intuition model, by contrast, judgements about the exemplification of knowledge, or goodness, in particular cases have epistemological priority and can be rationally based directly on the deliverances of rational intuition. On this view, one's conscious awareness of the fact that a particular case exemplifies knowledge, which is explained by one's understanding of the concept of knowledge, provides one with a non-inferential rational basis for judging that the case exemplifies knowledge.

The cases of moral and epistemic judgements bring out particularly clearly the intuitive appeal of the rational intuition model, as compared with the deductive model. But its intuitive appeal is supported by considerations of more theoretical nature. The problem with the explicit inference view is not *just* that it makes the acquisition of moral, epistemic and logical knowledge an extremely difficult and arcane matter; it also makes it viciously circular. In order to know that a particular case exemplifies knowledge, I already need to know a great deal about knowledge – indeed, I need to have explicit knowledge of which conditions are necessary and sufficient for knowledge. But how am I to acquire this knowledge? At this point, it is hard to deny that we acquire knowledge of necessary and sufficient conditions for knowledge on the basis of an attempt to systematize and explain the judgements that we are pretheoretically inclined to make in particular cases, and which we are (defeasibly) entitled to presume are in good epistemic standing. Certainly, this has been the methodology which characterizes the mainstream in the literature on the conceptual analysis of knowledge. But the explicit inference view makes it viciously circular.

For similar reasons, it is quite implausible to require that I need to have articulated, or even to be capable of articulating, the general form of an argument in order to be rational in judging that a particular argument of that form is valid. Consider the following “Days of the Week” argument:

- (1) Today is Monday.
- (2) If today is Monday, then tomorrow is Tuesday.

Therefore,

- (3) Tomorrow is Tuesday.

¹¹ See McDowell (1979) for an early critique of this view. I hope in the future to discuss Peacocke's (2003) view that we have tacit knowledge of moral principles which are brought to bear in making judgements in particular cases.

I can just see that this argument is valid, and be rational in judging that it is, without being able to articulate the general form of the argument. It is not just that the capacity to articulate the general form of the argument is somewhat arcane and technical; it also depends on a more fundamental capacity to evaluate particular arguments as valid or invalid. Suppose I form the belief that the Days of the Week argument is valid and go on to try to justify my belief on the basis of the claim that the argument instantiates the following form which I also claim to be logically valid:

- (1) P
 - (2) If P, then Q
- Therefore,
- (3) Q

Then I am open to apparent counterexamples, in the style of Van McGee, involving instances of this form which have conditionals embedded in the antecedent of the conditional premise. Now, I might try to explain away the intuitive plausibility of these apparent counterexamples and, who knows, I might even be successful, but their initial plausibility cannot be denied. The crucial point is that my judgement that the Days of the Week argument is valid (as opposed to my attempt to give an explicit justification of that judgement) should not depend on the outcome of this debate. Rather, it is plain that the debate itself is a debate about how best to systematize and explain the judgements we are pretheoretically inclined to make in particular cases, which we are defeasibly entitled to presume are in good epistemic standing.

A purely deductive model of our acquisition of logical knowledge is unacceptable; it needs to be recognized that the acquisition of logical knowledge depends both on rational intuition and on a broadly inductive process of inference to the best explanation of judgements based on rational intuition.

8. On Brewer's Non-Reflective Access Requirement

For these reasons, then, it seems to me that rational intuition is both a genuine phenomenon and one that has an important role to play in the epistemology of logic. However, it is a further question whether it can be invoked in a plausible defence of Brewer's claim that the rationality of deductive inference depends on the fact that the inference is made on the basis of a conscious appreciation of its validity. Brewer's proposal seems to be that a subject's belief in the premises of a deductively valid argument provides him with a reason for believing its conclusion in virtue of embodying an appreciation of the logical validity of the form of argument in question. According to this line, the appreciation of logical validity is actually embodied in the very beliefs which serve to provide the premises of the inference. But this is not Dummett's conception of the matter. According to him:

It is not enough merely to grasp the thought expressed by each line of the proof; in addition, one must perceive patterns common to those thoughts and others, patterns which are not given with the thoughts as a condition for grasping them but which require a further insight to apprehend. (1991, p.198)

Indeed, he takes this point to be crucial in the explanation of how it is that deductive inference can be informative, in the sense of providing a source of new knowledge:

Deductive reasoning is thus in no way a mechanical process, though it may be set out so as to be checkable mechanically: it has a creative component, involving the apprehension of patterns within the thoughts expressed, and relating them to one another, that are not required for or given with a grasp of those thoughts themselves. Since it has this creative component, a knowledge of the premises of an inferential step does not entail a knowledge of the conclusion, even when we attend to them simultaneously; and so deductive knowledge can yield new knowledge. Since the relevant patterns need to be discerned, such reasoning is fruitful; but since they are there to be discerned, its validity is not called into question. (1981, p.42)

If believing the premises of a valid argument were sufficient to provide you with an appreciation of how the truth of the conclusion follows from the premises, then deductive inference would not be informative. On Dummett's account, by contrast, it requires a creative insight to see how the conclusion of an argument follows from its premises, and so the informativeness of inference is preserved.

So, on Dummett's account, it is possible to understand the premises of an argument without thereby having any conscious appreciation of the fact that the conclusion of the argument follows. We therefore have to consider the case in which a subject believes the premises of an argument and makes the transition in thought to believing its conclusion, but *without* being guided by any conscious appreciation of the fact that the conclusion follows from the premises. According to Brewer, such a transition cannot be rational. But this seems to me to be setting the hurdle too high. Brewer and Dummett make a convincing case that following a proof with a genuine sense of understanding requires an appreciation of its validity. But following a proof is a relatively sophisticated activity. There are much more mundane cases of deductive reasoning in which it is not so plausible to demand an appreciation of validity, or even the capacity for appreciation of validity, as a requirement for rationality.¹²

9. A More Permissive Proposal

I want to make a more permissive proposal according to which a deductive inference is rational, in the most basic cases, just in case the form of the inference is a valid form, and the thinker makes the inference *because* it is of that form.

What is it to make an inference *because* it is of a certain form? I have in mind the following kind of account. At the personal level, it is just a matter of exercising a *disposition* to make inferences of that form, but at the subpersonal level, the disposition must be grounded by some state that is a common cause of the various exercises of the inferential disposition, and

¹² I have in mind cases of deductive inference in animals and children, but also in unreflective adults and reflective adults in their more unreflective moments.

hence which embodies tacit knowledge of the form of inference in question.¹³ On this account, there is no requirement on representation of the form of inference at the personal level: in particular, there is no requirement that the subject should be guided in making the inference by any conscious appreciation of the fact that the form of inference in question is a valid form.

But if this all there is to be said, then what, if anything, is left out on Fodor's account of the rationality of deductive inference? The suggestion we have been considering is that deductive inference requires understanding, which requires consciousness; but Fodor leaves out any mention of consciousness in his account. The basic idea behind this suggestion seems to me to be correct, but it can be deployed in different ways. The Dummett/Brewer/Peacocke line is that making a deductive inference in a way which involves genuine understanding requires a conscious appreciation of the validity of the form of inference. My objection was that, while this requirement may be appropriate in relatively sophisticated cases, such as following a proof, it is inappropriate in more mundane cases of deductive inference. But I think there is a different way to make the connection between inference, understanding and consciousness.

A deductive inference is certain kind of (non-deviant) causal transition between *beliefs*. And in order to have a belief with a certain content, it is necessary to *understand* that content, which requires understanding its component concepts. Now, the crucial question is this: what is required for understanding a logical concept? And the crucial claim is that it is not sufficient to have tacit knowledge of what are in fact valid logical forms. Tacit knowledge of logical forms is just a matter of having a full-blooded disposition to make *transitions between intentional states* whose contents instantiate those forms. There is no requirement for understanding the contents of the states involved in the relevant transitions. (Compare: the intentional contents involved in low-level computations in the visual or grammar modules.) Understanding a logical concept, on the other hand, is a matter of having a full-blooded disposition to make *deductive inferences between beliefs* whose contents instantiate certain logical forms. Since we are now in the domain of belief and inference, possession of the relevant disposition *does* require understanding the contents of the beliefs involved in the inference.

But now we seem to be moving in a circle which is too small to be illuminating. How can we get a grip on what is distinctive about inferences between beliefs as compared with merely inference-like transitions between merely belief-like intentional states, without simply appealing to a notion of understanding which is taken as primitive? How can we get some independent grip on the claim that understanding is involved in the one kind of case, but not the other?

The answer I propose is that the notions of belief, inference and understanding are *normatively* individuated: that is to say, they are individuated by their *rational* role. We can make sense of the distinction between causally isomorphic systems of inference-like transitions between belief-like intentional states on the one hand, and inferences between beliefs on the other, in virtue of the fact that the latter exhibit rational relations,

¹³ See Peacocke (1992) Ch.7.1 for a more detailed proposal along these lines.

whereas the former do not. But we have still not yet reached a satisfactory stopping point. There must be some property of beliefs in virtue of which they are capable of standing in rational relations to one another, a property which might be lacked by a causally isomorphic system of belief-like states. That property, I claim, is availability to consciousness.¹⁴

In summary, understanding a logical concept is a matter of being disposed to make distinctively *rational* inferences between beliefs, where the rationality of these inferences is grounded, in part, by the fact that the beliefs involved are available to consciousness. So, as promised, there is an alternative to the Dummett/Brewer/Peacocke account of the connection between consciousness, understanding and the rationality of deductive inference.

This alternative is more permissive than the Dummett/Brewer/Peacocke account insofar as it denies that a rational deductive inference must be made on the basis of a conscious appreciation of the validity of the form of inference in question. I argued that this requirement might be plausible for relatively sophisticated cases, such as following a proof, but not in more mundane cases of deductive reasoning. But I also think that the permissive alternative provides the materials for explaining how conscious appreciation of validity is possible in more sophisticated cases. According to my proposal, the following are requirements for the rationality of a deductive inference:

- (1) The thinker must understand the contents of the beliefs which serve as the premises and conclusion of the inference; in other words, the thinker must know what it would be for them to be true;
- (2) The beliefs involved in the inference must be available to consciousness.

The fact that there are conditions under which the beliefs involved in the inference are occurrent in consciousness grounds the possibility that the thinker may become consciously aware of certain relations between their contents. And given his knowledge of what it would be for them to be true, it grounds the possibility that he may become aware of certain *semantic* relations between their contents: in particular, that the truth of the conclusion follows from the truth of the premises. In other words, it grounds the possibility of conscious awareness of the validity of the logical form of the inference.

10. Access Internalism

Brewer's non-reflective access requirement is motivated by the need to accommodate the internalist intuitions that are prompted by standard counterexamples to pure reliabilism while avoiding the problems faced by traditional forms of access internalism. According to Brewer's diagnosis, these problems stem from the following pair of commitments:

- (1) to the thesis that a person's reasons are essentially recognizable by her as such;
- (2) to the idea that this recognition can only be a matter of

¹⁴ These claims are elaborated and defended in more detail in my papers, "Rationality and the Subject's Point of View" and "The Autonomy of Personal Level Explanation".

her second-order knowledge that the mental state providing the reason in question is appropriately related to that for which it is a reason, where this is independent of the first-order state itself in that she might have been in just that state yet not had the second-order knowledge required for its status as a reason for her. (1999, p.164)

The pure reliabilist response is to reject (1) altogether, whereas Brewer's suggestion is that we can endorse (1) while rejecting (2). Effectively, then, his proposal is that we should replace the traditional *reflective* access requirement by means of the *unreflective* access requirement.

However, there is another way to characterize the elusive middle ground that we should be seeking. According to this suggestion, what is wrong with traditional forms of internalism is not the reflective access requirement *per se*, but rather the traditional interpretations of that requirement. On the traditional interpretations, the possibility of access (i.e. recognition of reasons) is grounded in such a way as to impose substantive constraints on the reflective capacities of the rational subject in question. This interpretation of the requirement is basically taken over by Brewer, except that the constraints are imposed on unreflective, as opposed to reflective, capacities of the rational subject. On the alternative interpretation, by contrast, the possibility of access is grounded in the nature of the reason-giving states themselves. Thus, even if the subject is not capable of recognizing the rationality of her belief, it may be rational all the same, so long as it is formed in such a way as to ground the possibility (in principle) of recognition.¹⁵

In giving an account of the requirements for rationality, we always have to be sensitive to the fact that there are many different degrees of sophistication in rational thought and action. We should not be so impressed by the more sophisticated kinds of rationality that we end up denying that more mundane kinds qualify as rational at all. But neither should we be so impressed by the more mundane kinds that we rob ourselves of the resources for explaining the more sophisticated kinds. I have tried to give an account that is as permissive as possible, but which provides the resources for explaining how it is that the most mundane kinds of rational thinking make possible the more sophisticated kinds. Thus, it is intended to exhibit what seems to me to be the correct methodology in this area.

Bibliography

Ayers, Michael (1991) *Locke*, volume 1. London: Routledge.

Block, Ned (1980) "Psychologism and Behaviourism"

Boghossian, Paul (2001) "How Are Objective Epistemic Reasons Possible?" *Philosophical Studies*.

Boghossian, Paul (2003) "Blind Reasoning," *Proceedings of the Aristotelian Society, Supplementary Volume*.

¹⁵ See "Rationality and the Subject's Point of View" for further defence of these claims.

- Brewer, Bill (1995) "Mental Causation: Compulsion by Reason", in *Proceedings of the Aristotelian Society, Supplementary Volume 64*.
- Brewer, Bill (1999) *Perception and Reason*, Oxford University Press.**
- Dummett, Michael (1981) *The Interpretation of Frege's Philosophy*. London: Duckworth.
- Dummett, Michael (1991) *The Logical Basis of Metaphysics*. London: Duckworth.
- Fodor, Jerry (1987) *Psychosemantics*. MIT Press.
- Fodor, Jerry and Pylyshyn, Zenon (1988) "Connectionism and Cognitive Architecture: A Critical Analysis," *Cognition* 28.
- Fodor, Jerry (1998) *Concepts: Where Cognitive Science Went Wrong*. Oxford University Press.
- McDowell, John (1979) "Virtue and Reason," *The Monist* 62.
- McGee, Vann "A Counterexample to Modus Ponens"
- Peacocke, Christopher (1994) "Content, Computation and Externalism", *Mind and Language* 4.
- Christopher Peacocke (2003) *The Realm of Reason*, Oxford University Press.
- Searle, John (1980) "Minds, Brains and Programs," *Behavioural and Brain Sciences* 3.
- Searle, John (1990) "Consciousness, Explanatory Inversion and Cognitive Science", in *Behavioural and Brain Sciences* 13.
- Searle, John (1992) *Rediscovering the Mind*.
- Wright, Crispin (2001) "On Basic Logical Knowledge," *Philosophical Studies*.