

# The Roots of Reason

---

*Philosophical Essays on Rationality,  
Evolution, and Probability*

DAVID PAPINEAU

CLARENDON PRESS • OXFORD

# CONTENTS

Introduction	I
1. Normativity and Judgement	6
2. The Evolution of Knowledge	39
3. The Evolution of Means–End Reasoning	83
4. Probability as a Guide to Life <i>Co-authored with Helen Beebe</i>	130
5. Causation as a Guide to Life	167
6. Uncertain Decisions and the Many-Minds Interpretation of Quantum Mechanics	212
<i>Index</i>	240

# I

## *Normativity and Judgement*

### *1 Introduction*

It is widely assumed that the normativity of conceptual judgement poses problems for naturalism. Thus John McDowell urges that ‘The structure of the space of reasons stubbornly resists being appropriated within a naturalism that conceives nature as the realm of law’ (1994: 73). Similar sentiments have been expressed by many other writers, for example Robert Brandom (1994: p. xiii) and Paul Boghossian (1989: 548).

In this essay I want to dispute this widespread view. I am not going to deny that there are norms of judgement. My thesis will only be that such norms raise no special problems for naturalism. The normativity of judgement can as easily be accommodated by naturalists as by anybody else.

Let me begin by fixing an agreed subject-matter. By norms of judgement I mean true statements about the way judgements *ought* to be made.<sup>1</sup> I take it to be uncontentious that there are many truths like this.

<sup>1</sup> And by judgements I mean beliefs and their linguistic expression

However, while I accept that there are many such truths, I do not accept that there is a species of normativity peculiar to judgement. My general strategy in this essay will be to argue that all norms of judgement are derived from moral or personal values, and do not involve any *sui generis* species of judgemental normativity.

A central role here will be played by the end of truth. I shall contend that the most significant<sup>2</sup> norms of judgement can be viewed as prescriptions to the effect that, in order to achieve the truth, you ought to judge in such-and-such ways. In my view, there is nothing constitutively normative about the end of truth itself. So I take the force of these prescriptions to derive from independent moral or personal reasons for attaching value to truth.

Thus truth may in some (or possibly all) contexts have a *moral* value. Alternatively, people may sometimes simply *want* the truth, for whatever reason. In either case, these evaluations will imply truths of the form ‘you ought to judge in such-and-such ways’. In the former case, these will be derived moral oughts, about how to achieve the morally valuable end of truth; in the latter case, they will be derived hypothetical imperatives, about how to achieve the personally desired end of truth.

Of course, both these ‘oughts’ raise further questions for the naturalist. Most obviously, in the moral case the naturalist needs to explain how moral value itself is possible. There is also a less obvious difficulty about the second case, namely, to explain the status of the ‘you *ought* to *Y*’ that seems to follow from ‘you *want* *X*’ and ‘action *Y* is necessary for *X*’.<sup>3</sup>

(though until Section 7 I shall use ‘judgement’ to refer specifically to beliefs). In this essay I shall stay silent on the issue of exactly which mental attitudes qualify as beliefs, and in particular on whether moral attitudes do.

<sup>2</sup> At the end of Section 5 I shall say something about norms which relate to truth-independent aspects of judgement.

<sup>3</sup> When I first started working on this essay I assumed that these hypothetical ‘oughts’ at least could be taken for granted by naturalism. More

However, neither of these difficulties is peculiar to the context of conceptual thought. The issue of moral value arises generally, not just in connection with norms of judgement, as does the more subtle problem of understanding hypothetical 'oughts', as in 'you ought to tune your car well, if you want it to go fast'. So, if there are problems facing the naturalist here, they are not problems occasioned by the special nature of conceptual judgement.

Moreover, both these problems arise for philosophers of all kinds, and not only for naturalists. The nature of moral values presents a puzzle to all philosophers, as does the problem of understanding the 'oughts' in hypothetical imperatives. So naturalists will have no more problems than any other philosophers, if they can reduce the normativity of judgement to derived moral or hypothetical prescriptions.

specifically, I assumed that there is a sense in which (i) 'A ought to X' is an analytic consequence of (ii) 'A wants Y' and (iii) 'X is necessary to Y'. Given that there is nothing naturalistically problematic about facts of kinds (ii) and (iii), I therefore took it that there can be nothing naturalistically problematic about claims of kind (i): maybe such claims don't feature a *moral* 'ought' (maybe they involve a kind of pun on 'ought'), but even so they are analytic consequences of (ii) and (iii). More recently, however, I have been persuaded by a paper by John Broome (forthcoming) that this doesn't work. Broome shows that there is no special non-moral 'ought' that can be analytically derived from (ii) and (iii), because no 'ought' of any kind can be so derived. The only prescription in play is 'You ought, if you want Y (and X is a means), to perform X'. The non-conditional prescription 'You ought to X' cannot be detached. Broome draws an analogy. 'Necessarily, if John is married, he has a wife.' We ordinarily then say, knowing that John is married, that 'John must have a wife'. But this obviously doesn't mean that 'John *necessarily* has a wife'. It is simply an elliptical statement of the necessary conditional. No non-conditional modal claim can be detached. Similarly, Broome shows, 'You ought to X' is often elliptical for 'You ought, if you want Y (and X is a means), to perform X', and not itself detachable. This thus leaves us with the task of explaining the unreduced 'ought' which governs this prescriptive conditional.

For both these reasons, I will be content if I can persuade you that norms of judgement are derived prescriptions orientated to moral or personal values. I accept that there are philosophical problems attached to both these kinds of derived prescriptions. But these problems arise across the board, and for all philosophers.

The plan of this essay is as follows. In the next section I shall explain how my approach to the normativity of judgement presupposes a naturalist view of content. In Sections 3 to 5 I shall argue that my approach avoids three problems facing alternative views. Section 6 elaborates some of these points within the context of specific theories of content. Section 7 then asks whether the linguistic embodiment of judgements raises further problems for the naturalist.

Questions relating to linguistic norms will be postponed to this final section. Until then I shall assume that judgements are mental items, and moreover that their normativity is independent of any norms of language. This is a matter of expository strategy, not of any metaphysical commitment to a realm of pre-linguistic thought. As the final section will make clear, I am happy to allow that some judgements are in part constituted by norms of language. However, it will be helpful to proceed in stages, first showing how naturalism might account for non-linguistic norms of judgement, and then adding in the complexities that arise with language.

## *2 Norms and the Theory of Content*

The approach I have just outlined assumes there is nothing essentially evaluative about truth itself. Truth is in the first instance a descriptive property, like car-speed or celibacy. Like car-speed or celibacy, it may be personally desired or morally valuable, in given contexts, but that is additional to its nature.

Here I am taking an implicit stand on the analysis of truth-conditional content. In particular, I am taking a stand against those

approaches to content which place normativity *inside* the analysis of content, in the sense that they presuppose *sui generis* norms governing judgement in explaining truth-conditional content, and hence truth.<sup>4</sup> Let us call theories of this kind ‘non-naturalist’ theories of content. Neo-verificationist and Dummettian approaches to content are of this kind, since they take content to depend on the conditions in which you are *entitled to assert* a claim. So also are Davidsonian theories of content, which take content to depend *inter alia* on facts about when it is *reasonable to form* a belief.

From the point of view of such non-naturalist theories of content, the approach to normativity outlined in the last section will seem deficient, for that approach takes truth-conditional content and truth as given, prior to any issues of normativity, and then analyses norms of judgement as arising only when moral or personal value is attached to truth. Non-naturalist theories will regard this approach as necessarily ignoring the prior norms which play a role in constituting truth. Instead of explaining truth in terms of normativity, non-naturalists will complain, I am trying to explain normativity in terms of truth.

Still, this charge does not worry me, given the availability of a number of accounts of truth-conditional content which do not assume normativity in explaining truth-conditional content. I shall call theories of this kind ‘naturalist’ theories of content, and have

<sup>4</sup> Let us agree, against the deflationists (cf. n. 6 below), that a substantial account of truth-conditional content is presupposed by the notion of truth. We can also agree that, once we have such a substantial account of truth-conditional content, then truth itself falls out without further ado. Thus those who are prepared to reify truth conditions can simply say that a judgement is true just in case its truth condition obtains. And those who would rather not reify truth conditions, like myself, can note instead that a substantial account of truth-conditional content will imply, together with further facts, claims of the form ‘this judgement is true iff *p*’, for any given judgement, and will therewith specify what is required for any given judgement to be true. (Cf. Papineau, 1993: 83–6.)

in mind here such theories as indicator semantics (Stampe, 1977; Dretske, 1981), success semantics (Whyte, 1990, 1991), teleosemantics (Millikan, 1984, 1993; Papineau, 1984, 1993),<sup>5</sup> and Fodor's 'asymmetric dependence' theory of content (Fodor, 1987, 1990). All of these theories offer to explain truth-conditional content without any commitment to prior norms governing judgement.<sup>6</sup>

<sup>5</sup> Some readers may be wondering why I have classified teleosemantics as a theory that analyses content independently of norms. For many teleosemanticists take the view that their theory does indeed trade in norms, in particular 'biological norms', and that for this reason teleosemantics offers an immediate solution to the apparent conflict between naturalism and normativity (cf. Millikan, 1990; McGinn, 1989). However, I find this implausible. Although I myself uphold the teleosemantic theory of content, and so regard it as the right source for the sort of naturalist content I need for my own approach to normativity, I don't think that the teleosemantics itself says anything about norms, nor therefore anything about how to reconcile normativity with naturalism. Whatever norms are, I take it that they must involve some kind of *prescription*, some kind of implication about what *ought* to be done. This simply isn't true of the biological facts on which the teleosemantic theory rests. There is nothing wrong, I suppose, in talk of 'biological norms', if by this you mean simply that we have been designed by natural selection to operate in a certain way. But these biological 'norms' aren't norms in any prescriptive sense. It is a vulgar, and indeed dangerous, error to infer, from the premise that *X* has been biologically designed to *Y*, that in some sense *X ought* to *Y*. My knuckles have arguably been biologically designed to hit people with, but it doesn't in any sense follow that I ought so to use them. Again, a number of human male traits have undoubtedly been designed to foster sexual predatoriness, but it doesn't follow that men ought to be sexually predatory. Similarly with teleosemantics and judgement. As a teleosemanticist I hold that our beliefs have been biologically designed to track their truth conditions. But I don't think that this does anything to show they *ought* to do this.

<sup>6</sup> What about 'deflationary theories', by which I mean theories which combine the thought that truth falls out of content (cf. n. 4 above), with the suggestion that the device of disquotation (plus the idea of meaning as use) removes any need for a substantial account of truth-conditional



Given that views about normativity interact with general questions about content in this way, an obvious dialectical problem faces this essay. It seems that I won't be able to say anything useful about normativity without first establishing the right overall theory of content. Yet clearly there is no question of adjudicating fully on the right theory of content in this relatively short essay. Because of this, my primary thesis will perforce be conditional: *if* you have a naturalist theory of content, *then* you can explain norms of judgement as derived prescriptions orientated to the end of truth. Still, having said this, I do think that the line developed in this essay gives us some purchase at least on the larger issue of which approach to the theory of content is correct.

Most obviously, my line directly blocks the argument that naturalist theories of content are inadequate because they cannot account for the normativity of judgement. From the perspective of this essay, this argument begs the question. I am happy to agree that *non-naturalist* theories of content imply the existence of prior judgemental norms that resist appropriation 'within a naturalism that conceives of nature as the realm of law', in that non-naturalist theories postulate a peculiar species of content-constituting norms which appear mysterious from a naturalist perspective. But naturalist theories of content don't accept the existence of such special content-constituting norms, since naturalist theories explain content without any reference to such prior norms. They place the norms of judgement *outside* the theory of content, and hold that content and truth are constituted independently of facts of the form 'you ought to judge in such-and-such ways'. Given this, they can then account for such norms of judgement straightforwardly, holding that all the 'oughts' involved are derived oughts,

content? (Cf. Horwich, 1990.) I doubt that such theories will work (Papineau, 1993: 84–5), but as far as I can see there is no reason why deflationists should not agree with me on all the issues in this essay. After all, they don't assume anything in explaining content, so *a fortiori* don't assume prior norms.

arising when moral or personal value is attached to the independently constituted aim of truth.<sup>7</sup>

So the line defended in this essay directly blocks an argument *against* naturalist theories of content, namely, that they cannot account for the normativity of judgement. However, I also feel that it provides some positive support *for* such theories. This is because the naturalist approach dissolves a number of awkward problems about normativity facing non-naturalist theories of content. This isn't of course a conclusive argument. There may be further overriding considerations in favour of non-naturalist theories of content, in which case we will simply need to face up to the problems of normativity they generate. But it surely counts against non-naturalist theories of content that they generate difficult problems that do not arise for their naturalistic rivals.

In the next three sections I shall identify three such problems. The first relates to the *universality* of the requirement to seek the truth. The second relates to the *status* of norms of judgement. The third relates to animals, children, and other *unrefined beings*.

In this section I have explained how my naturalist approach to normativity stands opposed to non-naturalist theories of content, theories which explain content and truth in terms of norms of judgement, and so cannot in turn explain those norms in terms of truth. However, my approach to normativity will also be opposed by some philosophers who do not embrace such non-naturalist theories of content. I have in mind here those epistemologists who defend 'internalism', in the sense of upholding normative

<sup>7</sup> To avoid terminological confusion, let me make it clear that this is meant as an explanation of how naturalists *ought* to deal with normativity, not of what they actually say about the issue. I don't want to count a writer like Millikan as a non-naturalist, just because she adopts what I take to be the mistaken view that judgemental norms are present *inside* their theory of content (cf. n. 5 above). In my terminology, then, a theory of content is naturalist if it doesn't in fact commit us to prescriptive norms, whatever its proponents may say about this.

requirements on knowledge or justification that cannot be explained in terms of the pursuit of the external aim of truth. Such philosophers need not suppose that these 'internal' norms play any part in constituting content and truth (indeed they could embrace one of the naturalist theories of content listed above). But they would still deny my claim that the normativity of judgement can be fully explicated in terms of the pursuit of truth.

My main target in this essay is the resistance to naturalism arising from non-natural theories of content. In the space available I cannot also seriously address the motivations for epistemological internalism, and I shall not refer explicitly to such internalist views in what follows. Even so, some of the problems I shall pose seem to me to raise difficulties for epistemological internalists as well as for non-naturalists about content. In particular, the next two sections seem to me to raise such difficulties, even though the third problem, of unrefined beings, does not.

### ***3 Is There a Universal Requirement to Seek the Truth?***

I have contrasted naturalist theories, which place normativity *outside* the theory of content, with non-naturalist theories, which place it *inside*. In this section I want to point to a prima-facie reason for thinking that the naturalists are putting normativity in the right place.

Consider people who aim deliberately to mislead themselves. Suppose an elderly man realizes that he is likely to be upset if he learns about the real probability of his developing cancer, and so arranges to avoid any evidence that might undermine his sanguine belief that this probability is low. Or suppose an adolescent youth learns that people with an inflated view of their own worth are generally happier and more successful, and so deliberately seeks out evidence which will make him think overly well of himself. Of course, there are familiar psychological difficulties about deliber-

ately arranging to have false beliefs, but examples like this suggest they are not insuperable.<sup>8</sup>

Are these people acting wrongly? Of course, they aren't doing what they need to, *if* they want their beliefs to be true. But by hypothesis they don't want their beliefs to be true. So is there any other sense in which they are proceeding improperly?

It is not obvious to me that there is. I would say that it can sometimes be quite proper, as in the above examples, not to be moved by the aim of truth. However, if this is right, then there is a problem for those non-naturalist theories of content that explain content in terms of normativity, and so place the norms governing judgement inside the theory of content. For such theories make it constitutive of your possessing a belief with a certain truth condition that you be subject to norms which are apt to guide you to the truth.<sup>9</sup> In particular, if you believe that you have a low chance of cancer, or that you are blessed with above-average attributes, then the norms which constitute these beliefs will apply to you, and so you won't be thinking as you ought to, if you have formed those beliefs with due regard to normatively authorized evidence.

Let me be clear about the difficulty I am raising here. The complaint is not that non-naturalist theories of content suppose an *indefeasible* requirement to seek the truth at all costs. Prescriptions can compete, and there is no reason why the prescriptions required by non-naturalist theories of content should not be overridden by

<sup>8</sup> Note that these examples involve the manipulation of belief by the deliberate avoidance or pursuit of evidence, not any deliberate refusal to tailor your beliefs to the evidence you already have. I agree that cases of this latter kind are of doubtful psychological possibility. It is intriguing to consider exactly why this should be so.

<sup>9</sup> Here I assume that the non-naturalists' putative content-constituting norms will include truth-conducive norms. Some possible non-naturalist theories of content might lack this feature, as I shall observe in Section 6. But it is common to all actual non-naturalist theories, and in particular to theories in the Dummettian and Davidsonian traditions.

other considerations, such as the desirability of not upsetting yourself unnecessarily. My objection is to the weaker claim, to which non-naturalist theories of content do seem to be committed, that there is always *some* reason to seek the truth, even if it can be overridden. For even this weak claim seems to me implausible. Why should we suppose that the elderly man, or the adolescent youth, are in any sense obliged to comport their beliefs to the evidence? It doesn't seem to me that they are violating any prescriptions at all by adopting their entirely sensible strategies.

Note how there is no difficulty here if you adopt the naturalist view that all 'oughts' relating to judgements are derived oughts, arising because personal or moral value is attached to truth. For there is no obvious reason to suppose that derived 'oughts' of either of these kinds apply to cases like the elderly man or the youth. They don't want the truth, and it is not obvious that they are transgressing any moral boundaries in avoiding it. The naturalist view thus makes space for the possibility that people who aim to avoid the truth may nevertheless be acting entirely properly.

Perhaps some readers will wonder whether it is as easy as I am suggesting to lack personal or moral reasons for seeking the truth. Isn't there always at least some personal or moral reason for wanting the truth? I am doubtful about this, but note that, even if I am wrong here, this won't really help non-naturalism. For non-naturalists are committed to the view that the elderly man and adolescent youth are proceeding improperly, *whether or not* there are moral or personal reasons for them to seek the truth. After all, non-naturalists believe in *sui generis* norms of judgement, which prescribe the pursuit of truth independently of moral or personal considerations. It is this underlying commitment to *sui generis* norms that I am objecting to, not the idea that there may always be reasons of some kind for pursuing the truth.

Still, even though it is not directly relevant to my line of argument, it will be instructive to digress briefly, and consider whether it is indeed possible to lack any moral and personal reasons for

seeking the truth. On the question of personal reasons, it might seem that there is at least one sense in which everybody has a personal reason to pursue the truth (even if they themselves suppose, like the elderly man or adolescent youth, that they will be better off without it). For doesn't it remain the case that everybody will be better able to satisfy their desires, whatever they may be, if their beliefs are true? Well, I agree that *if* you act appropriately on true beliefs, then your actions are guaranteed to satisfy your desires, and indeed I take this pragmatic connection to be a crucial component in the analysis of truth-conditional content (Papineau, 1993: ch. 3.6–3.7). And this pragmatic connection does mean that there is always a species of derived personal value to having true beliefs that are relevant to action, for such truth will always help you to find a way of satisfying whatever desires you have.

However, this prescription only applies to those beliefs that *are* relevant to your actions, and so provides no reason at all for your wanting truth in beliefs which are not so relevant.<sup>10</sup> For example, if the elderly man cannot do anything about cancer, then his belief about his chance of developing it is unlikely to influence any of his actions, and it won't therefore matter to his success in satisfying his desires whether this belief is true or not. More generally, while we can agree that everybody has an interest in seeking the truth on action-relevant matters, such pragmatic considerations do nothing to show that people cannot quite sensibly avoid the truth on other issues.

So our personal desires need not always give us reason to pursue the truth. What about morality? Might it not be that the truth is always to be morally valued? Well, while it is uncon-

<sup>10</sup> Note also that the pragmatic connection with successful action is by no means the only possible reason for personally wanting truth. You could want the truth for any number of reasons, or indeed for none. (In Section 7 below I shall explain how the institution of language in particular provides one specific reason for wanting true beliefs.)

tentious that the truth is morally valuable in some contexts (it is clearly morally valuable that we should have true beliefs about the dangers of smoking, for example), it is by no means obvious that it is *always* morally valuable. Does it matter *morally* if musicologists are mistaken about the drummer on Elvis's first Sun session? Or again, does it matter morally if the elderly man remains hopeful that he won't get cancer, or that the youth thinks he is more estimable than he is?

Of course this issue deserves further discussion. It may well be that reflection on the overall role of truth in human life, its importance for social institutions, and so on, might show us that we always have a moral duty to respect the truth. Still, as I said earlier, even if this is so, it won't really help my non-naturalist opponents. For they are committed to truth having a value independently of any such considerations. They put the value of truth inside the theory of content, not outside with moral values. And this seems wrong. If further reflection were to show us that truth is *not* always of moral value, then in what sense would the elderly man or adolescent youth necessarily be acting wrongly in ignoring it?

#### 4 *The Status of Judgemental Norms*

I am aware that the points made in the last section will carry little dialectical force against non-naturalists. For clear-headed non-naturalists can simply repeat that their putative 'oughts' of judgement are special, quite different in kind from moral or personal 'oughts'. So, whether or not the elderly man and adolescent youth are not doing anything wrong by ordinary moral or personal standards, they will still be acting wrongly in a further sense, by violating the *sui generis* norms of judgement.

Even so, it has been worth rehearsing these issues, for they serve to highlight the peculiar status of the non-naturalists' putative norms of judgement. If these norms are quite distinct from

moral or personal 'oughts', then where do they come from? What kind of fact is it that we categorically 'ought' to reason in certain ways? And whence does the guiding force of these 'oughts' derive—why *should* we reason in these ways? I know that I should do what is morally required. And there is an obvious sense in which I should do what will get me what I want. But I find myself in difficulty understanding why I should be moved by the non-naturalists' putative *sui generis* norms of judgement.

Thinkers in the non-naturalist tradition have in effect addressed this issue, but it cannot be said that they have arrived at any canonical answers. I have in mind here the discussion surrounding Wittgenstein's 'rule-following considerations'. This is not the place to unravel all the different strands of that discussion, many of which are to do with issues of language. But I take it that one central issue which remains, after we put issues of language to one side (cf. Boghossian, 1989: 514), is the question of why we ought to judge so-and-so on specific occasions. Why ought we to judge that *these things are green*? Why ought we to judge that 1,002 is the result of adding 2 to 1,000?

When we focus on this issue, non-naturalist theories seem to have little room to manoeuvre. It doesn't seem right to explain these 'oughts' in terms of individual dispositions. The fact that I am typically disposed to judge that certain things are green, say, doesn't make these judgements right, for I may be systematically prone to error. Nor does it help to widen the issue to some larger community, for whole communities too can be prone to error. In the end there seems to be little alternative for the non-naturalist except to hold that the norms of judgement are primitive and not to be further explained (McDowell, 1984). This seems unsatisfactory. Even if we can't reduce judgemental norms to other kinds of facts, it is surely desirable that we should have some kind of understanding of the peculiar force that judgemental norms are supposed to exert on us.

Once more, none of these difficulties about the status of judge-



mental norms arises on the naturalist line I am advocating. On my view, what makes it the case that you ought to judge in certain ways on specific occasions is that this will be a means to your judging *truly*. There is nothing circular about this analysis, provided truth itself can be analysed without appealing to norms of judgement, and the adoption of truth as an aim in turn explained by reference to moral or personal value attached to truth.

From my perspective, then, difficulties about rule-following are peculiar to non-natural theories of content. If you want to explain truth-conditional content in terms of norms governing judgement, as non-natural theories of content do, then you need to explain the source of these norms before you explain content, and it is not obvious there is any good way of doing this. But this problem simply does not arise for naturalist theories of content. They can explain content without mentioning norms. And then, given this, they can happily place norms of judgement outside the theory of content, as prescriptions which follow once moral or personal value is placed on truth.

Some readers may be wondering whether the naturalists' norms can suitably *guide* thinkers. We need norms of judgement not just to set a standard, but also to show thinkers how to conform to this standard. But bare prescriptions such as 'you should judge truly', or even such more specific imperatives as 'you should judge *green* of green objects', 'you should judge  $n$  plus 2 when asked to add 2 to  $n$ ', would seem not to fit this bill. They may articulate the aim of judging truly, either generally or on more specific matters, but they are not going to help thinkers who need guidance on how to achieve this aim.

This query calls for much detailed analysis, some of which will be sketched in Section 6 below. But at this stage let me merely observe that the adoption of a given end will standardly generate derivative norms about appropriate means to that end. Just as the adoption of car-speed as an end generates the prescription 'you ought to tune your car well', so will adoption of truth as an end

generate prescriptions about means of arriving at truth. The reliabilist tradition in epistemology becomes relevant at this point. This tradition is concerned to identify belief-forming methods that generally produce true beliefs as output. These are just the kind of methods that you *ought* to use, if value is attached to truth. I would argue that inductive logic, deductive logic, standard observational procedures, and many other modes of reasoning are reliable methods in this sense. So, by derivatively prescribing these methods as means, the naturalist approach to normativity does not just assume truth as an aim, but also generates advice about how to achieve it.

At this point it is worth observing that reliability for truth is not the only desideratum that can be placed on belief-forming methods. At least two others offer themselves. It will normally also be desirable that our belief-forming methods are *important*, in the sense of delivering informative beliefs on matters of concern, and *economical*, in the sense of not using large amounts of time or other resources. Both these desiderata will generally pull against the requirement that belief-forming methods should deliver only true conclusions. Moreover, any moral or personal considerations that call for reliability are likely also to call for importance and economy, so that some joint optimization of these mutually conflicting desiderata will usually be required.

For naturalists, then, reliability for truth is simply one potentially valuable feature of belief-forming processes, alongside importance, economy, and possibly others. This marks a further contrast between the naturalists and their opponents, a further sense in which those norms of judgement which relate to the aim of truth are not *sui generis* for naturalists. For naturalists, truth-conducive norms of judgement constitute just one dimension among a number of possible ways of evaluating belief-forming methods. These norms have no special status which distinguishes them from evaluations orientated to further desiderata like importance and economy. This might seem odd to philosophers brought

up to focus on truth-related issues, but it seems to me just the right attitude. There seems no good reason to privilege the desideratum of reliability for truth over others when we evaluate methods of belief-formation. True, an economical method that generates lots of important falsehoods is not generally worth much. But neither is a reliable method that consumes costly resources in generating trivialities.

Having said this, naturalists will do well to concede a significant terminological point to their opponents. When we consider whether a belief is '*justified*', or constitutes '*knowledge*', we evaluate the methods behind it purely from the perspective of reliability for truth, and in abstraction from such issues as importance and economy. Imagine a man who spends a month counting the individual blades of grass in his garden. We will no doubt feel this is a complete waste of time, and that the conclusion is of no possible interest to anyone, yet we will not say on this account that he does not *know* how many blades of grass there are, nor that his belief in their number is not *justified*.

It is striking how these central epistemological notions are exclusively truth-focused in this way. I am inclined to think that this tells us something about the central importance of truth-seeking for human society (cf. Essay 2, Section 10). Still, we can recognize this importance, and the epistemological terminology it carries in train, without supposing that there is some special type of normativity attaching to the aim of truth.

## 5 *Unrefined Thinkers*

Let me turn now to another respect in which naturalist theories of content fare better than non-naturalist ones. On the face of it, many beings who do not observe truth-conducive norms of judgement nevertheless seem to have contentful beliefs. The most obvious examples are higher mammals and very young children, who

often act in ways that seem to call for the attribution of beliefs, but seem in no obvious sense to be observing norms of judgement. This poses a prima-facie problem for those non-naturalist theories which make sensitivity to such norms a precondition of the possession of contentful beliefs.

In addition, some grown-up human beings pose a similar problem for non-naturalist theories of content. I am thinking here of the large body of psychological literature which suggests that normal humans characteristically violate principles of logic, elementary probability theory, and so on (cf. Cherniak, 1986; Stein, 1996). In this case it is not just that the believers lack the sophisticated abilities required to observe norms. They have the intellectual wherewithal, yet simply fail to apply it to the pursuit of truth. If this is right, then again we have a prima-facie problem for those theories which make sensitivity to truth-conducive norms constitutive of the possession of contentful beliefs.

Once more, non-natural theories of content seem to have little room to manoeuvre. Some philosophers, it is true, have argued that the psychological research just alluded to should not be taken at face value, and that, when the matter is looked into fully, the supposed violations of rational norms can be explained away (Cohen, 1981). I myself doubt that anybody without philosophical preconceptions would find this line plausible, but, even if it is right, the non-naturalist still has the infants and animals to deal with.

Here the non-naturalist has two options. Either infants and animals are guided by norms, or they lack beliefs. Neither option seems attractive. It is hard to see how animals and pre-linguistic children could in any sense be said to be observing norms. So non-naturalists have tended to go for the other option, suggesting that animals and infants do not really have beliefs. ('Dumb animals are natural beings and no more.' McDowell, 1994: 70.) There seems little to recommend this move. It is one thing to recognize that there are significant differences between the beliefs of non-linguis-

tic creatures and those of normal adult humans. It is another to inflate this into a fundamental division, with genuine thought on one side of the line, and nothing but mechanical response on the other.

Again, note how none of this is a problem for the naturalist. Naturalist theories of content regard judgement as prior to norms, and so have plenty of conceptual space for beings who have beliefs yet do not observe norms. Of course, specific naturalist theories of content will have specific views about the nature of judgement, and it will depend on the details exactly whom they recognize as believers. But by definition no naturalist theories will make sensitivity to norms itself a requirement for believing. (Moreover, for what it is worth, the specific naturalist theory I favour, teleosemantics, has no difficulty allowing that animals, infants, and systematically irrational adults all have beliefs.)

From the naturalist point of view, sensitivity to norms of judgement is an addendum to the possession of beliefs itself. Such sensitivity can arise in beings with beliefs, but does not have to. More specifically, given what I have said so far, it will arise in beings who adopt the end of true beliefs, for whatever reason, and consider what steps they need to take to achieve this end.

As it happens, I think adult human beings do this a great deal, and indeed that this plays a crucial role in human culture (cf. Essay 2, Section 5). But note that the deliberate pursuit of true beliefs involves a sophisticated meta-representational concern *about* beliefs, and is therefore not something that we should expect to be present in all beings who *have* beliefs. Since young infants, and probably all animals, lack the notion of true belief, they will be incapable of sensitivity to such norms. And even adult human beings, who do have the notion of true belief,<sup>11</sup> don't always take

<sup>11</sup> Footnote 4 implies that I am here taking ordinary adults to have a 'substantial account of truth-conditional content'. However, this needn't mean a sophisticated philosophical theory of content, like a teleosemantic

pains to ensure that they are getting their beliefs from reliable sources, and often simply take them as they come.

These points now allow a better understanding of the psychological literature on widespread human irrationality. In particular, they suggest an answer to one obvious query about this literature, namely, that it seems to be belied by the success of human enterprises which call for high levels of doxastic accuracy ('If we're so dumb, how come we sent a man to the moon?'). The answer is that humans normally reason badly, but can set themselves to do better. We can suppose that humans normally get their beliefs from 'quick and dirty' processes, which no doubt worked well enough in our evolutionary history, but all too often lead us astray in the modern world, and in particular in the experimental circumstances studied by psychologists. On the other hand, this is only part of the story. For humans can also deliberately set themselves to use reliable methods instead of the 'quick and dirty' tricks. Even if humans don't always follow norms that offer reliable routes to the truth, they are perfectly capable of this, and characteristically set themselves to do so when it matters.

According to the naturalist position, then, infant and animal believers lack the intellectual wherewithal to follow truth-conducive norms of judgement, and even adult human believers, who are intellectually capable of observing such norms, don't always do so. This position may be philosophically unfamiliar, but it has the virtue of plausibility. The usual condition, even for adult humans, is surely to think and act without worrying about what you are doing. Sometimes, I agree, we adults do make an effort to observe truth-conducive norms of judgement, and it is important that we be able to do this. But it overintellectualizes ordinary mental life to suppose that this kind of effort accompanies all judgement.

theory, say, or a Davidsonian analysis, but only some way of telling which people have which beliefs with which contents.

## 6 *Some Specific Theories of Content*

My remarks so far have contrasted ‘naturalist’ and ‘non-naturalist’ theories of content. In this section I want to attend to some different species within these broader genera, and note briefly how the details affect some of the points made so far. I think these refinements are of some intrinsic interest, but they will not affect my overall line of argument, so those who have been comfortable with the level of abstraction adopted so far may wish to skip to the next section.

The distinction between natural and non-natural theories of content interacts with that between ‘output-based’ and ‘input-based’ theories. By ‘output-based’ I mean those theories, like success semantics and teleosemantics, which analyse the content of beliefs in terms of the *actions they prompt*; and by ‘input-based’ I mean those theories, like verificationism and Dummettian theories, which analyse the content of beliefs in terms of the *conditions that produce them*. In more detail, output-based theories aim to equate the truth conditions of beliefs with those conditions that guarantee that resulting actions will satisfy desires (with success semantics then differing from teleosemantics in its account of desire satisfaction),<sup>12</sup> while input-based theories aim to equate truth (or assertibility) conditions with the circumstances that typically give rise to beliefs.

When I first started thinking about these matters, I saw the conflict as primarily between ‘output-based’ naturalistic theories and ‘input-based’ non-naturalistic ones. This was because the naturalist theories I favour are in fact output-based, while the most familiar non-naturalist theories are in fact input-based. However,

<sup>12</sup> See Whyte (1991), Papineau (1993: ch. 3.8). In this essay I am ignoring philosophical questions about satisfaction conditions for desires. There are non-naturalist as well as naturalist answers to these questions, as there are about the truth conditions of beliefs.

further reflection showed me that I was running two issues together here. The input versus output issue is orthogonal to the natural versus non-natural issue. After all, there are also naturalistic input-based theories, like indicator semantics and Fodor's asymmetric dependence theory. And, even if there is no clear-cut example of a non-natural output-based theory, there are elements of this position in Davidsonian approaches to content.

These distinctions complicate what I have said so far. Take first the division of *naturalist* theories of content into output-based (teleosemantics, success semantics) and input-based (indicator semantics, asymmetric dependence). I have suggested that naturalist theories of content can regard truth-conducive rules guiding belief-formation as specifications of means appropriate to the valued end of truth. But this comes out somewhat differently for output-based and input-based naturalist theories. For output-based theories, all such means will be *contingently* connected to the end of truth: since what makes a belief true are circumstances picked out by the belief's relation to action, it will be an independent further fact that any given way of arriving at the belief should track the truth. With input-based theories it works differently. Since the circumstances that make a belief true will now be equated with some subset of the possible causes of that belief, any recommendation that this belief should issue only from those causes, while still construable as a derived imperative orientated to the valued end of truth, will now be recommending a course that is *constitutively* connected with the end of truth.

Given this, we might expect a related contrast on the output side. That is, we might expect output-based naturalist theories to view some aspect of output as constitutively related to truth, which input-based natural theories regard as only contingently connected. There is indeed such a contrast. To see this, note first how input- and output-based naturalist theories are a kind of mirror image of each other. Just as an input-based theory needs to distinguish, among all the circumstances which can actually cause



beliefs, those canonical conditions which constitute their truth conditions (this is why Fodor's asymmetric dependence, or something similar, is needed), so also does an output-based theory need to distinguish, among all the behavioural choices which can actually be occasioned by beliefs, those canonical behavioural consequences that play a role in constituting truth conditions: for it is the specific circumstances in which *those* canonical choices lead to successful actions that constitute the truth conditions. (The point is that we need to winnow out choices occasioned by rushes of blood, acting in haste, and so on, or we will get the truth conditions wrong. It is a challenge for a naturalistic output-based theory to make this distinction naturalistically, a challenge which I think is best met by teleological considerations. Cf. Papineau, 1993: ch. 3.7.7.)

Now consider the recommendation that people with true beliefs *ought* to choose in those canonical ways, *if* they are to satisfy their desires, first from the perspective of an output-based naturalist theory, and then from the point of view of an input-based one. From the output-based perspective, this hypothetical imperative falls constitutively out of the nature of truth: if your belief is true, and you so choose, then, by the nature of truth, you will satisfy your desires. For, on the output-based theory, true beliefs are in effect identified as those which lead to desire satisfaction specifically when they combine with such canonical choices. On an input-based theory, by contrast, this hypothetical imperative (that you *ought* to choose in the canonical ways, *if* you are to satisfy desires) will depend on contingent facts. For there is nothing in an input-based analysis of content (truth conditions = canonical *sources*) to ensure that if you have true beliefs, and choose in the canonical ways, your desires will be satisfied. This may be true, even for an input-based theorist, but if so it will be because of contingent facts about the world, and not because of the nature of truth.

Let me now quickly consider these matters from the point of view of non-naturalist theories of content. On input-based non-naturalist theories, like verificationism or Dummettian theories,

recommendations about appropriate sources for beliefs will be taken to be primitive norms which fix truth (or assertibility) conditions for those beliefs. These theories are thus akin to input-based naturalist theories like Fodor's, in fixing truth conditions by reference to appropriate sources, except that they hold that there is an essential normativity involved in the reference to 'appropriate'. Moreover, they seem to come out exactly like Fodor-style theories on the output side. When it comes to recommendations that people with true beliefs *ought* to choose in such-and-such canonical ways, if they are to satisfy their desires, an input-based normative account offers no constitutive rationale. Just as with Fodor, they can still say that these choices are sensible because they will ensure satisfaction of desires, if your beliefs are true. But there is no constitutive tie here, and this claim will be true, if it is, because of contingent facts about the world.

What about output-based normative theories? Nobody seems actually to hold such a theory, but it is easy enough to imagine one. It would be like success semantics or teleosemantics in fixing truth conditions as those circumstances in which appropriate choices will succeed, except in holding, again, that 'appropriate' here is primitively normative. So on the input side this theory wouldn't disagree with its output-based naturalist counterpart: it would view belief-forming recommendations as derived imperatives specifying that certain sources are an effective means, as a matter of contingent fact, when value is attached to the end of truth.

Even if there are no exclusively output-based non-natural theories of content, the Davidsonian approach to content can be viewed, as I said, as embodying elements of this approach. For we can view Davidson as in effect upholding a non-naturalist theory of content which is *both* output-based and input-based, in that he identifies beliefs as states which ought primitively to arise from such-and-such epistemological inputs *and* which ought primitively to lead to such-and-such choices.

Won't a Janus-faced theory like Davidson's at some point face an embarrassing doubling of primitive and derived norms? In its initial identification of beliefs, the theory invokes both primitive input and output norms. But there seems no barrier to a Davidsonian also recognizing derived norms: derived input norms arising whenever moral or personal value is attached to the end of truth, and derived output norms about which choices you ought to make, if you are to satisfy your desires. This seems like an overabundance of norms. Surely we do not want to allow *both* primitive *and* derived norms prescribing effectively the same strategies for forming beliefs and making choices.

I think this is indeed a genuine problem, but it is not peculiar to Davidson's theory. Rather it is a problem that arises for all non-naturalist theories, and indeed is simply a special case of the problem I raised in Section 3. We can put the point like this. Naturalists (whether input- or output-based) want to construe all norms relating to beliefs as derived imperatives: input norms are orientated to moral or personal values placed on truth, while output norms are orientated to the satisfaction of desires. Non-naturalists, by contrast, hold that at least some norms (input, or output, or both) are prior, *sui generis* and non-derived, and moreover that these prior norms underpin the constitution of content for beliefs. But this means that non-naturalists of any kind are going to be threatened with some doubling of norms. For there will be always be room for corresponding derived norms, orientated to the value of truth, or to the satisfaction of desires, alongside whichever norms non-naturalists take to be primitive and content-constituting. (Of course non-naturalists, by definition, will suppose that these primitive norms come first, so to speak, in that they are needed to constitute truth-conditional content; but the trouble is that, once content has been so constituted, non-naturalists won't be able to stop extra derived norms arising, because certain ways of reasoning are additionally valued as a means to the end of truth, or because certain ways of choosing are additionally valued as a means to the satisfaction of desires.)

In Section 3 I pressed the implausibility of there being primitive norms for truth-seeking, as required by input-based non-naturalist theories, in addition to derived norms which arise when value is attached to truth. We can now see that the point generalizes. An output-based non-naturalist theory would have primitive output norms (you ought to choose in certain canonical ways) *and* the hypothetical recommendation that you ought to so choose, if you are to satisfy your desires. And a Davidsonian Janus-faced theory will have primitive norms for both outputs and inputs, *and* in addition both derived input and output norms.

I take this generalized point to argue against non-naturalist theories of any kind, just as the special case in Section 3 argued against non-naturalist theories with primitive input norms. Primitive norms are an embarrassment, given the extent to which derived prescriptions are available anyway.

## 7 Norms of Language

In this final section I want briefly to consider the linguistic dimension of judgement. Obviously, in the space available, much will have to be taken for granted. Let me cut through a number of issues by adopting the following simple two-part model for an assertoric language  $L$ . (1) *Knowledge of truth conditions*. For all sentences  $s$  in  $L$ , speakers agree on a claim of the form:  $s$  is true if and only if  $p$ .<sup>13</sup> (2) *A rule of truth*. Utterances are governed by the prescription, 'You should utter  $s$  only if (you know that)  $p$ '.<sup>14</sup>

Given anything like this model, linguistic judgements will

<sup>13</sup> Perhaps they need also to know that such claims are derived canonically from facts about the semantic roles of  $s$ 's parts, if this is to constitute knowledge of truth conditions. Alternatively, they may derive such claims from claims of the form that  $s$  means that  $p$ , or  $s$  expresses the belief that  $p$ , or  $s$  would be true iff  $p$ . I shall skip these refinements.

<sup>14</sup> See Williamson (1996) for arguments in favour of inserting the knowledge requirement here.

uncontroversially be constituted in part by a normative rule of truth. I have argued that content for non-linguistic judgements is not constituted by norms. By contrast, assertoric utterances would clearly not have their contents without a rule of truth. If it were not the rule that *s* is only to be asserted when *p*, then utterances of *s* would not have the significance they do. (Maybe *s* would still in some sense be *true* if and only if *p*, given its agreed truth condition, but it is only in virtue of the rule of truth that we can regard those who utter *s* as *committed* to *p*, as offering the judgement that *p*—as opposed to conjecturing that *p*, say, or denying that *p*, or whatever.)

So I accept that linguistic judgements are in this sense normatively constituted, if anything like the simple model above is right. However, I do not accept that there is anything in this to trouble naturalism. To start with, note that the issue of normativity arises rather differently in the present context from the way it did in earlier sections of this essay. In earlier sections I was concerned to explain how naturalists can *uphold* claims of the form ‘you ought to judge in such-and-such ways’. That is, I wanted to show how naturalism can accommodate the *truth* of such normative claims. But there is no need for me, as a naturalist theorist, to *uphold* the norm which I take to constitute linguistic assertions. What matters is that the speakers of the relevant language accept the rule of truth for assertions, not that this rule is true. It doesn’t matter whether they *really* ought only to say what is true. Their assertoric practice will still contain a rule of truth, which will constitute certain utterances as assertions, as long as they *think* they ought only to say what is true.

In sociological parlance, what matters is that the rule of truth is a *social fact* in the linguistic community in question, not that it is in fact a good rule. As such, it is akin to social facts proscribing theft, extra-marital sex, broken promises, driving on the right, or splitting infinitives. Such social facts are constituted by the evaluative attitudes of the members of the relevant community, quite inde-

pendently of the normative truth. Even if extra-marital sex is in fact quite proper in many circumstances, it would still be true that in many societies there is a social norm prohibiting it, in the sense that the members of those societies disapprove of it and will exert pressure to prevent it.

Now, there are many interesting philosophical questions about the constitution of social facts. For example, there seems reason to distinguish social facts which depend on the community's moral attitudes (to sex, theft, promises), from those which rest on feelings about etiquette or style (splitting infinitives), and again from those which derive from a common interest in coordination (driving on the left). And it is by no means clear in which of these three categories the rule of truth ought to be placed. But, wherever we put it, there seems nothing to trouble the naturalist. Social facts of all kinds are constituted by the psychological attitudes of the members of the relevant society. In the earlier sections of this essay I have sought to rebut the most common argument for thinking that naturalism cannot accommodate psychological attitudes. If I am right, and naturalism can accommodate psychological attitudes, then it should be able to accommodate social facts in general, and the rule of truth in particular.

I have argued that the constitution of languages presupposes only that the relevant rule of truth is a social fact, and not that it is true. But it is worth observing that, if a rule of truth is indeed a social fact, then members of linguistic communities will generally have some kind of pro-attitude towards saying only what is true, and this in itself will imply hypothetical imperatives about what they ought to do to get the truths they want. Moreover, if it is morally valuable that languages exist, as it arguably is, then this in turn will imply that people ought morally to avoid false utterances.

It is also worth observing that any such prescriptive truths about *utterances* will also generate prescriptions about the avoidance of false *belief*. If there is moral or personal value in speaking

truly, then there will be derived moral or personal value in believing truly, since aiming to believe truly is a presupposition of aiming to speak truly. In preceding sections of this essay I have considered what follows if true beliefs have moral or personal value. We now see how the institution of language can add to the reasons for which true beliefs might have such value.

Let me finish by briefly addressing the issue of the relation between thought and language. My overall expository strategy may have seemed to commit me to the priority of thought over language. After all, I have explained representation, in the first instance, by applying naturalist theories of content to beliefs, without attending to any possible dependence beliefs may have on language. And, more specifically, when I have turned to language, I may seem to have credited speakers with language-independent abilities to *think* that  $p$ , in explaining their understanding of sentences in terms of beliefs of the form *sentence  $s$  is true iff  $p$* .

However, I don't think there is in fact anything worrisome here. Despite any initial appearances to the contrary, I am happy to concede that thought depends on language in ways I have yet to address. But I do not accept that this dependence invalidates my overall analysis. More specifically, it does not preclude the application of naturalist theories of content to beliefs, nor does it preclude the explanation of assertoric language in terms of speakers' beliefs about truth conditions.

It is worth distinguishing two ways in which belief depends on language: psychodevelopmentally and constitutively. To start with, linguistic training obviously exerts a developmental influence on psychological abilities. Many thoughts are clearly too complex for humans whose perceptual sensitivities and inferential associations have not been shaped by a relevant linguistic upbringing. In addition, there is also a more fundamental, constitutive dependence of thought on language. I am thinking here of those now familiar cases where my possession of a concept does not depend on individual factors like my perceptual sensitivities and inferential asso-

ciations, but rather on the structure of my linguistic community. Names of people (*Aristotle*) or kinds (*arthritis*) provide obvious examples. Here it is not that my linguistic upbringing has caused me to develop new internal states. Rather, my possessing such concepts is *constituted* by my being embedded in a linguistic community. If you change the community, while leaving me the same, my concept changes.

I deny that there is anything in either of these kinds of dependence that is inconsistent with the overall story I have told in this paper. The first point to address is my use of naturalist theories to analyse the representational contents of beliefs. Now, such theories need clearly have no difficulty with the fact that linguistic training affects psychological sensitivities and abilities. For they can simply embrace these linguistically developed abilities as part of what fixes the representational contents of beliefs. Moreover, they can make just the same move with respect to the social constitution of content. Again, they can simply embrace this, and agree that certain of your states represent such-and-such, not just because of intrinsic facts about you, but also because of your relationships to other people, and in particular because of your embedding within a linguistic community. Of course, exactly how these points might work will depend on the specifics of different naturalist theories of content. But it is not hard to see in outline how your development of special psychological abilities, or your membership of a linguistic community, might affect the informational significance of your states, or their teleological purposes, or their success conditions. All in all, there seems no obvious reason why naturalist theories of belief content should not simply accept that beliefs depend on language in the two ways I have indicated.

The other point where I might seem to be committed to the priority of thought over language is where I assume that assertoric language depends on speakers' beliefs that sentences *s* are true iff *p*. The worry here is that I apparently need to assume that speakers can think *p* *before* they learn sentences which mean that *p*.



However, I do not in fact need any such assumption of temporal priority. I need only claim that once the linguistic institution is up and running, as it were, and speakers are capable of thinking that  $p$ , then the different parts of the institution, thinking and speaking, are coordinated by beliefs of the form:  $s$  is true iff  $p$ .

True, this will leave us with an interesting kind of circularity in some cases. Speakers have concepts like *arthritis* because they belong to a linguistic community. This community exists in part because speakers follow rules for the linguistic expression of these concepts (they aim to assert 'arthritis' sentences only when certain arthritis facts obtain). But speakers will only be able to follow these rules if they already have the concepts (they won't be able to grasp the rules for 'arthritis' sentences unless they are capable of *arthritis* thoughts). So the concepts are constituted by the rules which are constituted by the concepts.<sup>15</sup>

However, while I accept there is a circle here, I do not think it vicious. It is just a special case of certain phenomena (here, mastery of concepts) being constituted by rules which in turn presuppose those phenomena. Many familiar social phenomena display this feature. There wouldn't be soccer matches unless there were rules about who wins soccer matches. There wouldn't be bank balances were there not regulations about who has access to bank balances. Certain concepts wouldn't exist were there not norms governing the linguistic expression of those concepts. This kind of circularity is certainly interesting, but it does not present any more of a special philosophical problem than soccer matches or bank balances do.<sup>16</sup>

<sup>15</sup> It is specifically the constitutional, not psychodevelopmental, dependence of thought on language that generates this interesting circularity. It is scarcely noteworthy that a new generation's concepts should causally depend on the linguistic practices of older people with those concepts.

<sup>16</sup> Earlier versions (some very early) of this essay were read at the Scottish Postgraduate Philosophy Society, King's College London, University of Nottingham, University of St Andrew's, and the Cambridge

## REFERENCES

- Boghossian, P. (1989), 'The Rule-Following Considerations', *Mind*, 98.
- Brandom, R. (1994), *Making it Explicit* (Cambridge, Mass.: Harvard University Press).
- Broome, J. (forthcoming), 'Relative Rationality'.
- Cherniak, C. (1986), *Minimal Rationality* (Cambridge, Mass.: MIT Press).
- Cohen, L. J. (1981), 'Can Human Irrationality be Experimentally Demonstrated?' *Behavioural and Brain Sciences*, 4.
- Dretske, F. (1981), *Knowledge and the Flow of Information* (Oxford: Basil Blackwell).
- Fodor, J. (1987), *Psychosemantics* (Cambridge, Mass.: MIT Press).
- (1990), *A Theory of Content* (Cambridge, Mass.: MIT Press).
- Horwich, P. 1990. *Truth* (Oxford: Basil Blackwell).
- McDowell, J. (1984), 'Wittgenstein on Following a Rule', *Synthese*, 63.
- (1994), *Mind and World* (Cambridge, Mass.: Harvard University Press).
- McGinn, C. (1989), *Mental Content* (Oxford: Basil Blackwell).
- Millikan, R. (1984), *Language, Thought, and other Biological Categories* (Cambridge, Mass: MIT Press).
- (1990), 'Truth Rules, Hoverflies and the Kripke–Wittgenstein Paradox', *Philosophical Review*, 99.
- (1993), *White Queen Psychology and other Essays for Alice* (Cambridge, Mass.: MIT Press).
- Papineau, D. (1984), 'Representation and Explanation', *Philosophy of Science*, 51.
- (1993), *Philosophical Naturalism* (Oxford: Basil Blackwell).
- (2000), 'The Evolution of Knowledge', in P. Carruthers (ed.), *The Evolution of Mind* (Cambridge: Cambridge University Press).
- Stampe, D. (1977), 'Towards a Causal Theory of Linguistic

Moral Science Club. Thanks are due to all who contributed on these and other occasions, in particular to Robert Black, John Broome, Edward Craig, Peter Goldie, Christopher Hill, Paul Horwich, Alan Millar, Gabriel Segal, Scott Sturgeon, Julia Tanney, Timothy Williamson, and especially Neil Manson.

Representation', in P. French, T. Ueling, and H. Wettstein (eds.),  
*Midwest Studies in Philosophy*, ii.

Stein, E. (1996), *Without Good Reason* (Oxford: Clarendon Press).

Whyte, J. (1990), 'Success Semantics', *Analysis*, 50.

— 1991, 'The Normal Rewards of Success', *Analysis*, 51.

Williamson, T. (1996), 'Knowing and Asserting', *Philosophical Review*, 105.