A matter of opinion*

Keith Frankish

Abstract

This paper sets out the case for a two-level theory of human psychology. It takes its start from Daniel Dennett's distinction between belief and opinion, arguing that it has the power to account for a number of tensions within our commonsense concept of belief. It argues, however, that Dennett's account is seriously inadequate, particularly in its treatment of the role of opinion in practical reasoning. The paper goes on to sketch an alternative proposal which retains the virtues of Dennett's suggestion, while providing a richer and more satisfying account of the cognitive role of opinion.

1. Introduction

Most theorists writing about belief share a common assumption. This is that belief is a *unitary* phenomenon – that whenever we ascribe a belief to a person, creature, or system, we ascribe essentially the same *kind* of state. Of course, no one denies that belief has varied aspects and manifestations – it is widely accepted that beliefs can be both occurrent or standing-state, tacit or explicit, conscious or non-conscious, and so on. But it is generally assumed that these are different aspects or variants of a single fundamental kind of state. So occurrent beliefs can be thought of as *activations* of standing-state beliefs, tacit beliefs as *dispositions* to form explicit beliefs, conscious beliefs as beliefs that are the object of *higher-order* beliefs and so on.

There have, however, been dissenting voices, suggesting that this apparent uniformity masks an important psychological distinction. Some writers distinguish passive belief from active *judgement*. Epistemologists, too, commonly mark a distinction between graded belief and flat-out *acceptance*. And Daniel Dennett has argued that we should distinguish nonverbal beliefs from a class of language-involving cognitive states which he calls *opinions*. The distinction between the two states is, he claims, a very important one:

> My hunch is that a proper cognitive psychology is going to have to make a sharp distinction between beliefs and opinions, that the psychology of opinions is really going to be rather different from the psychology of beliefs, and that the sorts of architecture that will do very well by, say, nonlinguistic perceptual beliefs ... is going to have to be supplemented rather substantially in order to handle opinions. (1991b, p. 26; cf. 1991c, p. 143)

A view of this kind has some attractions. For example, it can be argued that full-blown intentionality requires certain conceptual abilities that are unavailable to languageless

* This is the author's version of an article published in *Philosophical Psychology*, 11(4): 423-42 (1998). It may differ in minor ways from the published version.

creatures (Bennett 1964; Davidson 1975, 1982). Anyone impressed by such arguments, yet unwilling to deny that animals have intentional states, will be attracted to the idea that there are distinct kinds or levels of intentionality. Indeed, Dennett suggests that a failure to distinguish these levels lies at the root of many philosophical misconceptions about belief.[1]

There have, however, been few attempts to develop a rounded two-level theory of mind. Few theorists have sought to integrate the literature on acceptance with that on active judgement, or to link either with work on the role of language in thought. (Writers on acceptance tend to focus instead on technical questions about the rationality of flat-out belief and its relation to graded belief.) Dennett is one of the few exceptions here, drawing on a number of sources in a richly suggestive paper on opinion (Dennett 1978, chapter 16). However, he does not work out his ideas in a systematic way, and, like his predecessors, has little to say about the cognitive role of opinions. The central psychological questions remain unanswered. What exactly are opinions? How are they formed and processed? What role do they play in reasoning and decision-making? And why do we need them as well as beliefs?

This paper aims to remedy some of these omissions. I begin by outlining some contrasting aspects of belief which indicate the need for a two-level theory. I then look at Dennett's proposals and highlight some of their limitations. The principal of these, I argue, is a failure to find any real cognitive role for opinion. I go on to outline an alternative proposal which remedies this defect. I describe a mechanism by which opinions could influence reasoning and action and identify a distinctive cognitive role for them. A final section returns the discussion to Dennett, comparing the position sketched here with his and arguing that it is superior.

2. The bifurcation of belief

The idea that the concept of belief is ambiguous has some attractions. The concept is a rich – some would say messy – one, with roots in subjective experience, ordinary language and commonsense psychology. And there is little agreement about the nature of the thing it refers to. There are deep and seemingly intractable disputes about the ontology, semantics, and epistemology of belief, about its possession conditions, causal role, and relation to language. Reading the literature, it is hard to avoid the suspicion that opposing theorists sometimes have different phenomena in mind and are simply talking past each other. I want to begin by focusing on four contrasting aspects of belief, each of which suggests an underlying duality.[2] None of them provides conclusive proof, of course, and there are ways of incorporating them into a unified account, but they provide the raw material for a two-level alternative.[3]

---

1 See the references to 'opinion' in Dennett 1987; see also Dennett 1991c, p.143, and Dennett 1994, p.241.

2 For another approach to the ambiguity of belief, involving a fourfold classification, see Horst 1995.

3 Note that all the distinctions below are drawn in terms of *function* or *constitution*, not *content*. In distinguishing different kinds of belief, I shall assume, for the sake of argument, that the same range of contents is available to both.

Activation

Propositional attitudes, such as belief, can manifest themselves both as episodic events and as persisting states. We are often conscious of actively thinking that something is the case. For example, just now I was reflecting that the weather is unseasonably warm and thinking that it would be wise to turn the thermostat down. Episodes of active thinking like this are sometimes called *occurrent beliefs*. Of course, at any given moment we have lots of beliefs that are not occurrent. For example, a few moments ago I firmly believed, among other things, that cider is made from apples, that Paris is the capital of France and that my surname begins with an 'F' – though I was not in fact thinking of any of those things. Persisting doxastic states of this kind are usually called *standing-state* or *dispositional beliefs*. Now not all theories of belief give equal weight to both its occurrent and standing-state forms. Some theories – we might call them *occurrentist* ones – treat occurrent belief as primary. They hold that beliefs are cognitively inert until occurrently activated, and treat standing-state beliefs as dispositions to form occurrent ones (a disposition which might be realized in the form of a memory trace, say).[4] By contrast, another group of theories – let us call them *dispositionalist* – treat standing-state belief as primary. They identify such beliefs with dispositions to overt behaviour, and hold that they can manifest themselves directly in action without prior occurrent activation.[5]

Which view is the commonsense one? Intuitions pull both ways. Occurrent beliefs often seem to play a crucial role in the production of action. I am driving to work as usual. Suddenly, it occurs to me that roadworks are due to start today, and I decide to take a different route. I do so, moreover – or so it seems to me – precisely *because* the thought about the roadworks consciously occurred to me. If it had not done so, I would not, ceteris paribus, have changed course. Yet many beliefs seem to influence action without occurrent activation. Think, for example, of the numerous beliefs that guide your routine behaviour when driving: that one should drive on the left (in Britain, at any rate), that one should stop at a red light, that it is illegal to exceed the speed limit, and so on for countless others. It is because one has these beliefs that one drives as one does, and it would be perfectly natural to cite them in explanation of one's driving behaviour. Yet few of us entertain them occurrently as we drive – at least not consciously. Again, it is plausible to think that animals have beliefs and desires which influence their behaviour; rather less plausible to think that they have occurrent beliefs. (As Malcolm puts it, it would sound funny to say of an animal that a thought *occurred to him*, or *struck him*, or *went through his mind*; see Malcolm 1973).

Now if you hold that beliefs have to be occurrently activated in order to influence action, then you will have to say that actions like these, apparently initiated without occurrent thought, are in fact the product of *non-conscious* analogues of it. This is a worrisome claim, however. For one thing, it is not clear that there would be *time* for all

---

4 Most versions of the *representational theory of mind* are occurrentist, in this sense (see e.g. Fodor 1987). For a very clear statement of an occurrentist position, see Goldman 1970, chapter 4.

5 This means, of course, that the dispositionalist must deny that beliefs *cause* actions in the way that events cause their effects. For a well-known dispositionalist position, see Dennett 1987.

the necessary non-conscious occurrent beliefs to occur. (Think, again, of all the beliefs that are relevant to your moment-by-moment actions when driving in traffic.) Dispositionalist theorists are in a rather better position here: a huge number of dispositions can manifest themselves simultaneously in a single action. (Indeed, a disposition will *only* manifest itself against a background of many others.) However, the dispositionalist has a corresponding problem in accounting for occurrent belief, understood as a *precursor* of action. If beliefs just *are* behavioural dispositions, then it is hard to see how they could be activated in any way short of overt behavioural manifestation.

We have a stand-off, then: some cases seem to be best described in occurrentist terms, others in dispositionalist ones. Given this, it is tempting to wonder if the two theories might each characterize a distinct *level* of cognition – to wonder, that is, if there might be two distinct *kinds* of belief, one requiring occurrent activation, the other not. At any rate, it is a hypothesis worth investigating.

Control

It is widely accepted that we have no direct control over what we believe. Indeed, it is sometimes argued that the opposite claim – that beliefs can be acquired at will – is not only false, but *incoherent* in some way (see e.g. O'Shaughnessy 1980, Vol. 1, pp.21-28; Williams 1973). Now it is undeniable that *some* beliefs are passively acquired. One has only to think of beliefs that derive from perception, memory and simple inferential processes. Yet there is a long tradition of thinking that we have the power to *decide* what attitude to take towards a proposition, through an act of deliberate *judgement*. (Indeed, many philosophers would have identified what I have been calling *occurrent beliefs* with acts of deliberate judgement.) Aspects of commonsense mentalistic discourse also lend support to such a view. We speak of *trying* to believe, of *refusing* to believe, of *having a duty* to believe, of *needing* to believe, and so on. These idioms are often dismissed as pointing, at most, to the possibility of *indirectly* inducing belief, by exposing oneself to belief-forming influences. But there is another commonsense idiom which suggests that we are capable of more than this. This is our talk of *making up* and *changing* our minds. We often refer to such episodes, and speak of them as free intentional actions (we urge the indecisive to make up their minds and blame the inconstant for changing them). Of course, sometimes, when we speak of a person having made up their mind, we mean that they have made a decision to *do* something – that they have formed an intention, not a belief. But this does not exhaust the idiom. We also speak of making up or changing our minds about matters of fact – about the truth of a theory, say, or the safety of a course of action, or the honesty of a politician (see Baier 1979). That is to say, we allow that some *doxastic* attitudes can be objects of decision. And since not all such attitudes are formed in this way, this again points to a bifurcation in our notion of belief.

Degree

When we think about the role of belief in guiding action, it is very tempting to regard it as a continuously graded state, reflecting degrees of confidence. The thought goes like this. We frequently have to decide what to do without being completely sure of the facts. For example, suppose you are trying to decide what to buy for dinner. You are tempted to buy beef, but believe there is a small risk it may be infected with a deadly disease. [6] How do you decide? Well, buying beef will have very different outcomes depending on whether or not it is infected. So the rational way to proceed would be to consider how *desirable* or *undesirable* you find each of these outcomes and how *confident* you are that the meat is or is not infected. That is, you should consider the *utility* of buying beef relative to each background condition (infection present vs. infection not present), weighting this value in each case by your *confidence* in the existence of that condition. Summing these values will give you a measure of the *expected utility* of buying beef. If this is lower than that for not buying it, then you should not buy. This strategy can be generalized to yield a global procedure for rational decision-making – *Bayesian decision theory*, as it is called. Of course, few people actually make decisions by doing Bayesian calculations – at least not at a conscious level. Nevertheless, people can be *interpreted* as adhering to Bayesian principles: their preferences, as manifest in their behaviour, can interpreted as maximizing expected utility, in Bayesian fashion, relative to certain assignments of utility to outcomes and of confidence to propositions (see, for example, Maher 1993, pp.9ff).

The important thing to notice about this is that Bayesian theory speaks only of degrees of confidence and utility; it finds no role for a notion of *flat-out* or *unqualified* belief or desire.[7] It makes no sense to ask whether a person *believes* flat-out that beef is safe, only how *confident* they are of its safety – and, consequently, how much they are prepared to stake on it. Some theorists conclude that belief talk should be interpreted as, or replaced by, talk of degrees of confidence (Jeffrey 1970, pp.160-1). This view, however, seems to conflict with the commonsense view of the matter. For it means that we *fully* believe very few things and completely *disbelieve* almost nothing.[8] Yet we often want to make claims to knowledge; and it would seem odd to say that one *knows* that *p*, but does not fully *believe* it. Again, we would normally think it irrational to believe both a proposition and its contradictory. Yet if beliefs are degrees of

---

6 The presence of BSE, or 'mad cow disease', in British cattle, and the recent discovery of evidence that the disease can spread to humans, has lead some British consumers to form this belief.

7 Distinguish confidence assignments (or *subjective probabilities* as they are sometimes called) from flat-out beliefs in objective probabilities. To say that Jack has confidence of degree 0.8 in proposition p (or, equivalently, assigns a subjective probability of 0.8 to p) is quite a different thing from saying that he believes flat-out that the objective probability of p is 0.8. The two attitudes differ both in form and content.

8 According to the Bayesian, you should not assign a probability of 1 to a proposition unless you are willing to bet everything you have on its truth. For, you should then attach a probability of 0 to its negation – and should therefore discount *entirely* the disutility of any outcome contingent on its falsehood.

confidence, we do this all the time. It is a requirement of rationality that, if you have confidence of degree *n* in proposition *p*, then you should have confidence of degree 1-*n* in the proposition *Not-p*. So, excepting rare cases of complete certainty, no rational person will believe a proposition without simultaneously believing its contradictory to some degree. The Bayesian account also seems to have no place for the attitude manifested in deliberate judgement. The point of making up your mind about something is to settle categorically what you think of it – at least until further consideration. And this attitude is binary, not graded: for any proposition *p* you either have, or have not, made up your mind that *p*. Of course, you might make up your mind that *p* has a certain *objective probability* – but then it is the *proposition* that is graded, not your attitude to it. A final problem for the Bayesian is what to say about conscious reasoning. As I said, this very rarely takes a Bayesian form – people find statistical procedures hard to apply and generally prefer classical non-Bayesian methods. For example, we often use the *practical syllogism*. We reason that we want *p*, believe that *p* requires *q*, and so decide to try to bring it about that *q*. Here we treat beliefs as ungraded, binary states. Even when we do make explicit assessments of probability, we tend to use these as premises in classical reasoning. If the Bayesian is right, however, classical reasoning of this kind is *irrelevant* to the rationality of our actions (see de Sousa 1971, p.57).

So here is another tension in our commonsense notion of belief. On the one hand, it seems to be a continuously graded state, like length or mass; on the other a binary one, like being more than 3 feet long. Many theorists conclude that we need two distinct doxastic concepts: a concept of *partial belief*, reflecting degrees of confidence, and a concept of *flat-out belief* – usually called *acceptance* – which is binary. Now this is not yet to say that these concepts pick out distinct psychological phenomena: perhaps we could think of acceptance as a certain *level* of confidence – say, exceeding 0.5. It is, however, widely agreed that this will not do. It is easy to show that, given certain plausible assumptions about rationality, there is *no* level of confidence which could be either necessary or sufficient for acceptance.[9] Acceptance seems to be genuinely distinct from partial belief.

Language

It is sometimes claimed that natural language can act, not only as a vehicle for the *expression* of thoughts, but as itself a *medium* of thought. We can coin thoughts, it is suggested, in the very act of articulating them – often in the form of subvocalized, self-directed speech (recent advocates include Carruthers 1996; Dennett 1991a; Gaulker 1994). There is some introspective evidence for this view, and a powerful argument can be run for the view that conscious propositional thinking occurs in natural language (Carruthers 1996, 1998). Of course, it is implausible to suppose that *all* thinking occurs in natural language – animals and prelinguistic infants can think

[9] The arguments invoke the paradoxes of the Lottery and the Preface; see Kaplan 1996, chapter 3; Maher 1993, pp.133ff.

after all. So here, again, there is a motive for distinguishing two different kinds of thought – linguistic and nonlinguistic.

3. Dennett's opinions

In four different respects, then, the concept of belief bifurcates. Can we see a pattern here? Do the divisions indicate the presence of two distinct kinds of state? Dennett suggests so; ordinary language, he claims, conflates beliefs with *opinions*. Beliefs, according to Dennett, are dispositional states: to believe that *p* is simply to be disposed to behave in the way that a rational *p*-believer would, given one's other beliefs and desires. Both humans and animals have beliefs. Opinions, on the other hand, are a more sophisticated kind of state. In characterizing them, Dennett draws on Ronald de Sousa's work on acceptance (de Sousa 1971). It is, as I said, implausible to think of acceptance as a state of confidence. Instead, de Sousa suggests, we should think of it as a *behavioural* state, initiated by an act of *assent*. This, he explains, is a linguistic activity: to assent to a sentence is, as it were, to make a *bet on its truth* – to add it to one's stock of sentences labelled 'True'. Episodes of *occurrent belief* are in fact, de Sousa suggests, acts of assent, motivated by our low-level beliefs and desires – specifically our *epistemic* ones. We have, as it were, a *lust for truth* – a hankering for objects of unqualified epistemic virtue – and it is this which prompts us to assent to plausible sentences. A theory of this kind, de Sousa notes, can reconcile a Bayesian model of reasoning with the classical one. Each operates at a different *level*; Bayesian theory applies to low-level deliberative processes, the classical model to verbalized thinking, which proceeds by leaps of outright assent.

Dennett endorses de Sousa's proposal, and suggests that it also helps to illuminate the processes of change and making up of mind. In general, he suggests, these involve a sort of *behavioural commitment*. This is most obvious in the cases of practical decision – where one makes up one's mind to *do* something. But theoretical makings up of mind, he suggests, are similar: they too involve commitment – the leap of epistemic faith which de Sousa calls assent. In making such a leap, one passes from a generic desire for truths to a commitment to a particular presented sentence. The state generated by such an act – what Dennett calls an *opinion* – is not one of belief, but more like commitment or ownership, and it may outlast or diverge from the agent's low-level beliefs, as manifest in their actions. (It is when such divergences occur, Dennett suggests, that we fall into self-deception and akrasia; Dennett 1978, p.307). Although Dennett follows de Sousa in emphasizing the active dimension of opinion formation, he is careful to add that not *all* opinions are the products of active decision. We can also collect sentences unthinkingly – often, he suggests, because we see them as *sure bets*. Most of us, for example, routinely accept what we are told, unless we have reason to suspect deceit. What is distinctive of opinion is, not its aetiology, but the attitude of personal commitment it manifests.

Dennett's proposal is attractive. It neatly accounts for the competing intuitions we noted above. We have low-level beliefs, which are passively formed, graded, and non-

verbal, and which manifest themselves as behavioural dispositions. And we also have opinions, which display the opposite characteristics. Despite this, I think the account needs substantial revision. I have two small worries and one big one. The first concerns the role of language. I suspect that Dennett is right to say that change and making up of mind involve language. But it is implausible to think that they involve betting on *particular* sentences. I may remember having made up my mind about some matter (that a certain politician is untrustworthy, say), without being able to remember precisely the words I used to frame the thought. If I were bilingual, I might even forget which *language* I had used. Likewise, it seems possible to form indexicalized opinions (to decide, for example, that *that* politician is untrustworthy). But then it would be necessary to *modify* the sentence on which one was prepared to bet in order to accommodate changes in the way its subject was presented. This would involve something more like the adoption of a *policy* than a bet on a sentence.

Secondly, Dennett follows de Sousa in identifying episodes of occurrent belief with acts of opinion formation – that is, of judgement or making up of mind. Now I agree that some occurrent beliefs are such; but not all of them. Recall my occurrent belief about the roadworks, which led me to take a different route to work. In entertaining that thought I was not *making up my mind* about anything, but *recalling* something I already knew. It is not clear how episodes of this kind fit into Dennett's picture.

Now to the big worry. It concerns the *cognitive role* of opinion. What exactly does one commit oneself to, in betting on a sentence? Here there is a deep tension in Dennett's position. Although he insists that one is committed to one's opinions, he consistently writes as if the commitment involved is slight, and predominately verbal.[10] Thus, he often identifies opinion formation with the sort of intellectual assent that lacks real conviction (1978, pp.307-8; 1987, p.19n), and denies that opinions play any direct role in guiding behaviour:

> It is my beliefs and desires that predict my behaviour *directly*. My opinions can be relied on to predict my behaviour only to the degree, normally large, that my opinions and beliefs are in rational correspondence, i.e., roughly as Bayes would have them. (1978, pp. 306-7)

Opinions, it seems, are just sentences which one has picked up and is disposed to avow as true.[11]

Now in this Dennett is following the majority of acceptance-theorists. Acceptance is usually thought of as an intellectual attitude, motivated by narrowly cognitive ends

---

[10] His favourite example of an opinion is a bit of arcane information he once picked up from a play. It is, he implies not exactly something he *believes*, but, given suitable inducement – in a quiz, say – he might bet on its truth (Dennett 1978, p.306).

[11] That Dennett thinks of opinion as a rather superficial phenomenon is confirmed by his later references to it. He typically invokes the concept to account for features of human cognition which conflict with his basic commitment to an ascriptivist view of the mind. See the references to 'opinion' in his 1987.

(more precisely, by a desire to *maximize cognitive utility* – that is, to assert maximally true and comprehensive theories). So scientists may, after due investigation, decide to accept a theory as true for the purposes of future research, and may categorically assert its truth in their writings. But, theorists claim, this should have no influence on how they act outside the context of enquiry: as Maher insists, the decision to accept a theory should not affect one's willingness to act as if it is true in practical contexts (Maher 1993, p.150).[12] The reason for this is simple. The fact that you have accepted a proposition does not make it more probable, and should not, therefore, make you more confident of its truth (excepting cases of self-fulfilling prophecies, such as 'I won't get to sleep tonight'). Nor, in most cases, will it alter your practical utilities (deliberation is about *achieving* one's ends, not about *changing* them). So if the rational action to choose is the one determined, in Bayesian fashion, by one's assignments of confidence and utility, then acceptance should not alter a rational person's choices. To allow it to do so is, for the Bayesian, to lapse into irrationality. As the remark quoted above indicates, Dennett takes a similar view: one's opinions may *predict* one's behaviour, but they do not actively *guide* it.

Now if opinion were an attitude which manifested itself only in the context of academic inquiry, then this would be unobjectionable. But that is not Dennett's view. For he identifies opinion formation with making up of mind and occurrent belief. And these seem, on the face of it, to have immense practical influence. We can make up our minds about mundane matters, such as whether beef is safe to eat, the salesman trustworthy, or the weather threatening enough to justify taking an umbrella. And such decisions have significant behavioural consequences. If I make up my mind that eating beef is unsafe, then it is natural to suppose that this will affect, not only what I *say*, but what I *eat*. The same goes for occurrent belief. We entertain all manner of occurrent thoughts, and they seem to have a profound effect on what we do (recall my occurrent belief about the roadworks, which caused me to take a different route to work). So there is a tension here. On the one hand, Dennett wants to deny that opinions have a direct role in the guidance of action; on the other, he wants to identify them with conscious occurrent beliefs – episodes which do have such a role. He could diffuse the tension, of course, by denying that conscious occurrent beliefs directly influence action. Indeed, there are passages in his writing which suggest such a view (see, e.g., Dennett 1969 p.123, p.154; 1978 chapter 9, and 1987, p.91). But it is, I think, an implausible one.

4. A cognitive mechanism for opinion

Is there a way of squaring this circle? – that is, of identifying a role for opinions in rational decision-making, without giving up Bayesianism? We have already seen that

[12] This is not to deny that there is a *connection* between what a person accepts and how they behave in practical contexts; both, after all, are determined by the same set of underlying assignments of confidence and utility.

it would be irrational to allow opinions to feed into regular Bayesian reasoning.[13] Perhaps, then, we should think of opinion-based reasoning as constituting a separate system, which interacts with the belief-desire one only at certain points. (Such a distinction of levels would be more in the spirit of de Sousa's original proposal.) I think this is right, but it immediately raises a host of questions. What is the relation between the opinion system and the belief-desire one? How and when do they interact? And how does opinion-based reasoning work? What procedures does it use? We have identified opinions with conscious occurrent thoughts, but conscious reasoning is radically enthymematic. We think 'Beef is unsafe' and thereupon avoid it. We do not add 'Consuming unsafe food damages one's health', 'Good health is desirable', and so forth. Are the suppressed premises supplied somehow? And if so, by what means? (If the system's database contains only those propositions which have been explicitly assented to or consciously articulated, then it seems unlikely to contain them.)

I have a suggestion. It is that opinions are actively processed at a personal level. Forming an opinion, I suggest, involves committing oneself to an extended *policy* or *strategy* of reasoning – to taking the accepted proposition as a *premise* in one's deliberate conscious reasoning, both practical and theoretical. Executing such a policy would also involve various personal activities: keeping track of the premise, working out what conclusions it entails or which intentions it justifies, then deliberately accepting these conclusions and executing those intentions. Similarly, I suggest, one can form opinion-level desires by adopting some state of affairs as a *goal* in one's conscious reasoning, and regulating one's future deliberations accordingly. In effect, what I am suggesting is that the mechanisms of opinion-based reasoning are *intentional*. Opinions influence our behaviour because we *want* them to do so – because, having given them the status of premises and goals, we *want* to perform the actions they dictate and to draw the conclusions they license. [14]

This suggestion is really just an extension of de Sousa's original proposal. According to him, the processes of assent are driven by our low-level beliefs and desires. As enquirers after truth, we want to collect true sentences, and this desire prompts us to assent to likely candidates when they present themselves. Of course, these epistemic desires are usually not conscious. Nevertheless, they motivate acts of assent, and must be cited in explanation of them. My proposal is similar, but more wide-ranging. I suggest that we desire, not only to assent to true propositions in the context of inquiry, but also to take them as premises in our practical reasoning. (I shall say more shortly about why we have this desire and what ends it serves.) And if

---

[13] Indeed, if opinions were flagged as unqualifiedly true, it would be potentially disastrous. If you assign a proposition a probability of 1 then, according to the Bayesian, you should be prepared to stake *anything at all* on its truth.

[14] The account outlined here is inspired in part by Jonathan Cohen's work on acceptance (Cohen 1992). Cohen, too, thinks of acceptance as involving commitment to a policy of premising. However, like other acceptance-theorists, he remains ambivalent about its cognitive role. Our acts of acceptance, he suggests, play no *causal* role in the production of our actions, and serve merely to *rationalize* them (1992, p.64).

we believe that a previously adopted premise mandates a certain action – either another act of acceptance or some overt activity – then we will want to perform it *precisely because* it is so mandated. Again, like de Sousa, I am not suggesting that the low-level beliefs and desires which drive the opinion system are entertained consciously. (Indeed, I contend that those mental states which issue in conscious occurrent thoughts are not beliefs at all, but opinions.) Note, too, that it is not essential to this account that all opinions are the product of active decision. I agree with Dennett that opinions can be formed unthinkingly, as a result, say, of uncritically accepting something one is told (as I accepted what I was told about the roadworks). The defining mark of an opinion is not its origin, but one's attitude to it. An opinion is a proposition to whose truth one is committed, and which one is disposed to employ as a premise in one's conscious reasoning.

Let me illustrate all this with an example. Take the case where I make up my mind that beef is unsafe. This leads to my acquiring some new low-level beliefs – *not* beliefs about beef, but beliefs about myself and my cognitive commitments.[15] For one thing, I come to believe that *I have formed the opinion that beef is unsafe*.[16] (I do not *actively* form that belief, of course. All I *do* is decide to endorse the proposition; I come to believe that I have done so as a result of normal subpersonal mechanisms of self-awareness and belief-fixation.) Now suppose I am debating what food to buy. I recall my opinion of the danger of beef. I see that, given my background beliefs and desires, it dictates that I should not eat beef, and, therefore, since I do not want to buy food I will not consume, that I should not buy it either. So I believe that I have an opinion which commits me to choosing something other than beef. Since I have a general desire to act upon my opinions, I decide to do just that.

It may seem odd to suggest that my behaviour in this case is motivated, not by beliefs about *beef*, but by beliefs about my *opinions* and their commitments. Surely the reason I avoid beef is not that I believe my opinions commit me to avoiding it, but simply that I think it unsafe? Well, yes and no. We are supposing, remember, that the commonsense concept of belief is *ambiguous* between low-level belief and opinion. And this means that psychological explanations will be ambiguous, too – sometimes citing low-level beliefs, sometimes opinions. First-person explanation, I suspect, is very often of the latter kind. We tend to cite the content of our opinions rather than our beliefs, since it is these that are present to our conscious minds. And this is just what I would do in the present case: I would say simply that I thought beef was unsafe. And this explanation would be correct. That opinion played a crucial role in the generation of my action. If I had not formed it, or had subsequently forgotten it, I would not have acted as I did. Moreover, the opinion-invoking explanation is a genuinely *intentional* one. My opinion influenced my action precisely in virtue of its *semantic content*. If it had had a different content, it would not, ceteris paribus, have

---

[15] There is a ongoing debate as to whether accepting a proposition, p, requires or produces *belief* in *p* (See Cohen 1992; Clarke 1994). I propose to remain neutral on this; all I want to insist on is that accepting *p* normally produces belief in the proposition *that one has accepted* p.

[16] Of course, the belief might not be conceptualized in quite this way.

generated that action. For I would then not – again, ceteris paribus – have *believed* that it warranted the action, and would therefore not have been moved to perform it. So there is, as we would expect, a perfectly good intentional explanation of my action which cites only my attitude to the content *beef is unsafe*. There is also, however, another, quite different, explanation of my action – and one rather less easy for me to formulate – couched in terms of my low-level beliefs *about* my opinions and their commitments.

But doesn't the existence of this second explanation undermine the first? If my action can be adequately explained in terms of my low-level beliefs and desires, doesn't this mean that the opinion-based one is redundant – that my opinion was, in effect, causally idle? No: no more than the existence of a physical explanation for my action means that the psychological one is idle. The two are pitched at quite different levels. The opinion system is not an adjunct to the belief-desire one: it *supervenes* on it. Our opinions come to have the role they do in the aetiology of action precisely *in virtue of* our low-level beliefs and desires about them. My opinion about beef leads me to refrain from buying beef precisely because I *believe* it warrants such a course of action. We should not think of the opinion system as *interacting* with the belief-desire one, but as *realized in* it. It is, as it were a *softwired* feature of the human brain, implemented in more basic, low-level intentional processes.[17]

There are, then, two distinct strategies of psychological explanation available to us. Suppose you know my opinions and goals. Then you could predict my behaviour directly from them. You would just have to think of them as simple functional states with roles like those of beliefs and desires. Alternatively, you could make the same prediction on the basis of my beliefs and desires *about* my opinions. Each predictive strategy would have its advantages. The opinion-level route would be quicker and easier, the belief-desire one both more comprehensive (it would predict actions that were not opinion-generated, as well as those that were) and more reliable (it would work in cases where more powerful beliefs and desires interfered, preventing me from acting on my opinions). In this way, opinion talk fits nicely into the layered framework of explanatory stances that Dennett elsewhere describes (1987, chapter 2). Just as some events, in addition to having physical explanations, also have teleological and intentional explanations, so some events have yet a fourth kind of explanation. The trade-off between simplicity and accuracy is typical of the relation between different explanatory stances (see Dennett 1991d).

We can now deal with the worries mentioned above. The most serious concerned the rationality of opinion-based action. Won't it be irrational for people to allow their opinions to influence their behaviour – won't it involve a departure from Bayesian norms? No: not if the influence is exerted in the way described above, via their low-level beliefs and desires *about* their opinions. For the same action can be justified *both* on classical grounds, as dictated by one's opinions, *and* on Bayesian grounds, as warranted by one's opinion-related probabilities and desirabilities. All we have to

[17] For this reason I have elsewhere referred to it as a *virtual processor* and to its states as *virtual beliefs*. See my 1998.

suppose is that opinion formers habitually attach a high desirability precisely to *acting upon their opinions*. Then, in so doing, they will not be departing *from* Bayesian norms; but displaying their adherence *to* them. The fallacy in the original worry was to suppose that an act of assent could influence action only by raising the agent's confidence in the proposition assented to. And this is not so. Assenting to *p* should not make one more confident that *p*, though it should make one – massively – more confident *that one has assented to* p. Given some appreciation of *p*'s semantic powers, and a general desire to act on one's opinions, this may well be sufficient to motivate courses of action that hitherto seemed less than optimal.

The reader may suspect some sleight-of-hand here. If it would be irrational to assign *p* maximum probability in one's practical reasoning, how could it be rational to assign a high desirability to *acting* as if it were unqualifiedly true? Well, it certainly would not be rational to act as if *p* were true *in all contexts*. It would not be rational for me to avow that beef is unsafe in a context where false assertions were penalized with instant death and true ones rewarded with a halfpenny. But this is not to say that in such a case I would *abandon* my opinion. I might continue to regard it as the soundest view of the matter, and still use it as a basis for action in less momentous contexts. It is just that in this particular case the desire to act on my opinions would be overridden by prudential considerations.

What about the other problem I mentioned earlier – the enthymematic nature of conscious reasoning? Again, I think my account offers a solution. For it does not require opinions to be, as it were, *self-motivating*. Their influence upon action is *mediated* by our low-level beliefs about them and their warrant. And the mediating beliefs need not relate solely to our explicitly entertained opinions. What leads me to refrain from buying beef is not that I believe that my opinion that beef is unsafe warrants such action *on its own*, but that I believe it warrants it *given certain obvious background assumptions* – that beef is a food, that unsafe food damages one's health and so on. We could think of these assumptions as objects of *implicit acceptance – tacit opinions*, if you like. They are suppressed premises in our explicit reasoning – premises to which we would immediately assent if questioned.[18]

Finally, let us return to the two small worries about Dennett's position mentioned in the last section. The first concerned language. I said it was implausible to think of opinions as attitudes to particular sentences. But this feature is not essential to the account, at least as developed here. In forming an opinion one does not so much collect a *sentence* as a *propositional attitude*, committing oneself to regulating one's conscious reasoning in ways characteristic of the attitude adopted. Now in fact I believe that language will usually play an important role in this process: conscious reasoning, I believe, typically requires a linguistic vehicle.[19] But this does not mean that each opinion must be associated with a unique sentence: the same opinion might

---

[18] Distinguish tacit opinions in this sense from opinions which have been taken on trust – and so not actively assented to – but which have nonetheless been explicitly entertained (my belief about the roadworks, for example).

[19] For the argument, see my 1998.

assume different linguistic vehicles on different occasions. The second objection concerned the role of occurrent thoughts which (like my belief about the roadworks) involve the *recollection* of an opinion rather than its formation. Dennett's account had no role for such thoughts. But mine has. The conscious recollection of an opinion is just a prelude to its employment in conscious reasoning.

Note, by the way, that the account developed here justifies a commonsense intuition about occurrent thought. It is natural to think that conscious occurrent beliefs influence action precisely in virtue of the fact that they are conscious. (If I had not consciously recalled that the roadworks were due to start, I would not have changed course.) The present account explains this. For if conscious occurrent beliefs are opinions, and if opinions influence action in the way suggested, then conscious recollection will be a precondition for such influence. An opinion is a proposition that one is committed to employing as a premise in one's conscious reasoning. And it is trivial that *becoming conscious* is a precondition for *being employed in conscious reasoning*. Opinions can influence reasoning and action *specifically as opinions* only if they are consciously recalled.

The reader may spy a worry here. I am claiming that opinions have to be consciously recalled in order to influence action. But I have also claimed that opinions influence action *in virtue of* mediating low-level beliefs and desires about them. And I do not want to claim that low-level states need to be occurrently entertained – certainly not *consciously* entertained – in order to influence action. How can I reconcile these two claims? If all the appropriate mediating beliefs and desires were in place, why couldn't an opinion do its work without conscious tokening? The question is understandable, but misrepresents my proposal. The desire which drives the opinion-system is not simply a desire to treat suitable propositions as unqualifiedly true – a desire which might indeed manifest itself without conscious thought – but a desire to treat them as such *in one's conscious reasoning*. The desire is specific to the context of reflective conscious thought. Of course, it might have cognitive effects outside of this context – as will other opinion-related low-level states. But it will do so as an independent state, not as part of the mechanism of opinion.

To sum up: we have low-level beliefs and desires, which are non-linguistic, graded, and passively formed, and which manifest themselves without occurrent tokening. And we have opinions, newly construed as premising policies, which have the opposite properties. They need linguistic vehicles; they are binary (for any proposition p, one either is, or is not, committed to using p as a premise); they can be actively formed through acts of deliberate judgement, and they manifest themselves in conscious occurrent beliefs, which are in fact episodes in the formation or execution of premising policies.

5. The function of opinion

I have claimed that we attach a high desirability to forming and acting upon opinions – to taking propositions as premises in our practical reasoning and to performing the actions they warrant. But why should we do this? What is the point of the enterprise?

The answer, I suggest, is that it gives us *personal control* of our reasoning processes. Much of our behaviour is spontaneous – the product, not of conscious deliberation, but of non-conscious inferential processes to which we have no direct access and over which we have no direct control. Think, for instance, of the processes that guide one's moment-to-moment behaviour when driving or navigating one's way down a busy street. Often, this is all for the best (drivers know how disconcerting it can be to give conscious thought to what they are doing). And if our non-conscious inferential processes were optimal, then we could leave all our decisions to them. In every case we would spontaneously do the right thing, unthinkingly. But of course they are not optimal. We often find ourselves perplexed. And it is here that active judgement is called for. For perplexity is disturbing; we like to know what to do next. Indeed, we would often prefer to make a decision, even if it is not optimal, than to continue in uncertainty. (Dennett gives a homely example. You are in a restaurant choosing from the menu. It is very hard to say what you *really* want most; but given the pressures of time and etiquette you have to avow some categorical desire; Dennett 1987, p.20) So we make up our minds. We select a few salient propositions, of which we feel fairly confident, and a few goals, which we are content to pursue, accept them as categorically true or desirable, and feed them into simple decision procedures, such as the practical syllogism. With a bit of luck this yields a clear prescription which gets us moving again.

So here we have a powerful motive for forming opinions. Adopting explicit premises and goals can help to simplify complex problems and overcome deliberative log-jams; it furthers the ends of practical reflection. [20] Systematically followed through, the practice will create a distinct level of cognition which can be activated whenever low-level processes fail to yield a ready solution. Note, too, that once we have decided to adopt a particular premise or goal, we will generally want to stick with it and to use it as a basis for future deliberation. This is not because the decision affects our assessment of the proposition's probability or desirability, but simply because psychological consistency is itself valuable. Continually changing one's mind is a recipe for confusion and renewed perplexity. (Incorrigible vacillators can become *incapable* of acting.) So once you have plumped for the chocolate cheesecake, you immediately gain an extra interest in trying to satisfy that desire.

---

[20] Of course, people can decide to take propositions as premises for other reasons. For example, a lawyer might decide for professional purposes to take it as a premise that her client is innocent. Such an decision, however, would not mark the formation of an *opinion* in my sense (the lawyer has not *made up her mind* that her client is innocent, and we would not describe her, when she entertained that proposition, as having an occurrent *belief*). What disqualifies the lawyer's act is, I suggest, that it is motivated by local nondeliberative concerns – namely, a desire to do her job – rather than by a desire to further the ends of practical reflection.

There is another way in which opinion formation facilitates deliberative control. For opinion formers can select, not only the *inputs* to their conscious reasoning, but also the *procedures* they employ in such reasoning. I argued that opinions are processed at a personal level – that they motivate action in virtue of our beliefs about them. So, as opinion formers, we need to have beliefs about what warrants what. But we do not have to rely on native wit here. We can engage in activities designed to *generate* such beliefs (executing algorithms, applying heuristics, running thought-experiments, etc.). Moreover, we can evaluate these procedures themselves, modifying or abandoning bad ones and acquiring new and better ones. That is to say, we can control not only *what* we think about, but *how* we think about it. One consequence of this is that we can to some extent overcome the limitations of our hard-wired cognitive systems. For example, we seem to be naturally disposed to make certain systematic errors in reasoning with conditionals (these show up most clearly on the Wason selection task; see Wason 1966). It is likely that this is a consequence of hard-wired features of our brains. (It is notable that people's performance improves markedly when the task is presented in such a way as to draw on their native skills at detecting social free-loaders; see Cosmides and Tooby 1992.) However, with a little training in elementary logic, we can overcome this frailty. Such training, I suggest, does not involve a rewiring of the original subpersonal system (even logicians find the old Adam reasserting itself occasionally). Rather, it results in the subject learning new *personal* techniques for assessing the logical powers of explicitly entertained conditionals. One might, for example, learn to visualize the truth-table for the material conditional, thereby equipping oneself with a reliable source of beliefs about the cognitive commitments of opinions with a conditional form. The acquisition of such a skill would thus contribute to the installation of a soft-wired opinion-based cognitive processor.

6. Dennett compared

I have argued that Dennett's claims about opinion are vitiated by a failure to find any real cognitive role for opinion, and have sketched a rival account which remedies this defect. In this final section, I return the discussion to Dennett and compare the resulting account with his.

First, I want to head off a possible objection. It may be thought that the cost of the new account is just too high. Won't it involve rejecting out-of-hand the sort of *interpretationalist* view of the mind which Dennett advocates? The thought is this. I have argued that opinions influence reasoning and action in virtue of our low-level beliefs about them. So if you act upon the opinion that *p*, there will be a low-level intentional explanation of your action which cites the *belief* that you have that opinion. And this seems incompatible with interpretationalism about low-level belief. Take my food-buying activities. Surely, interpretational constraints would require you to ascribe to me the *simplest* low-level beliefs and desires that could adequately explain my behaviour – in this case, the belief that beef was unsafe and the desire to stay

healthy – and not sophisticated beliefs about my opinions and their commitments? So, it seems, we can have either interpretationalism about low-level belief or our new theory of opinion, but not both.

This is too hasty. We must remember that the interpreter does not aim to make sense of isolated actions, but of *patterns* of activity extending over time. And on the proposed account, the various *mental* acts involved in the formation and execution of the premising policies that constitute opinions will *themselves* be candidates for intentional characterization. So in the beef-buying case what we have to explain is, not just that I buy something other than beef, but that I do so as a result of an episode of conscious reflection upon the content *beef is unsafe* (itself part of a larger complex of such episodes). When this wider context is taken into account, the simplest interpretation of my action will be one that adverts to second-order beliefs of the kind mentioned. At any rate, that is the claim.

This objection does, however, point up another possible difference with Dennett. For our account requires us to take a *realist* view of certain kinds of unobservable mental behaviour. Forming opinions involves performing certain conscious mental *actions* – specifically, undertaking and executing strategies of reasoning. And these actions need have no overt behavioural manifestations – indeed, to a casual observer it might be hard to tell whether I have formed the opinion that beef is unsafe, or simply believe it with a high degree of confidence. And this seems to rule out the sort of anti-realism about the mind which Dennett favours.

Again, this is too hasty. It is true that Dennett is sceptical – rightly so in my opinion – of the existence of certain putative mental acts. In particular, he denies the existence of the sort of internal executive acts that are sometimes supposed to figure as the routine antecedents of overt action (1991a, chapter 8). But he does not deny that other kinds of mental acts occur – including silent verbalization, calculation, and visualization (see, e.g., Dennett 1991a, p.197). Now I suspect that the reason he finds acts of these kinds unobjectionable is that they can be regarded as *internalizations* of bits of overt behaviour (so inner speech is just overt speech with the articulatory component suppressed). And the acts involved in opinion formation are of just this innocuous kind. The ability to perform explicit inferential operations in one's head is just a refinement of the ability to perform them publicly, by manipulating external symbols. And private, self-directed commitments can be thought of as internalizations of certain sorts of social commitment. So there need be no principled resistance to speaking of such actions. Dennett would insist, of course, that we may be quite wrong about the *nature* of our mental acts. (Just because we speak of *rotating mental images*, for example, it does not follow that there exist rotatable mental items; see Dennett 1991a, chapters 4, 10.) But the present account does not suppose otherwise. The claim that occurrent beliefs are acts of premise-adoption is an explanatory hypothesis, not a piece of conceptual analysis.

Let us turn now to another contrast. I have recommended my story as an account of how conscious occurrent beliefs could have a cognitive role. But, of course, it is not the only such account around. Indeed, Dennett himself has recently developed a

rather different one. Nowadays, he talks less about *opinion* and more about the *Joycean machine* – the stream of self-directed inner speech which some novelists have attempted to reproduce. Conscious occurrent thoughts, he suggests, are just items of inner speech. And as such, he claims, they have an important cognitive role. Inner speech is channelled through a feedback loop linking the speech production and comprehension systems. Self-generated sentences are processed like externally produced ones, and tend to produce similar behavioural effects. So asking yourself a question may prompt an instinctive verbal reply containing information which you would otherwise have been unable to access. In this way, occurrent thought can help to focus attention and facilitate certain executive tasks. For example, we can engage in 'positive thinking' – deliberately repeating encouraging or admonitory phrases to ourselves in order to help to modify our attitudes. We can comment on our own behaviour, thereby heightening our awareness of what we are doing. We can engage in strategies of self-reminding, rehearsing the benefits of irksome tasks and the perils of superficially attractive ones. And we can facilitate information-retrieval by practising mnemonic tricks – repeating words and phrases in order to build up associations that will aid recall (Dennett 1991a, pp.224-5, 277-8, 301-2). Dennett's conjectures receive support from the literature on the role of private speech in children – especially that inspired by the work of Lev Vygotsky (Vygotsky 1962). A number of studies have confirmed that the frequency of children's private speech – that is, of their self-directed overt vocalization – does tend to predict performance on certain sorts of puzzle tasks (Berk 1992). And it seems to do so by exerting a self-regulatory function, in much the way that Dennett describes. (It takes the form of self-commentary, self-guiding remarks, and suchlike.) So here we have a rival account of the cognitive role of conscious thought – and a rather more economical one at that (one that does not mention second-order attitudes, for example).

Now I do not want to deny either the existence or the importance of the processes Dennett describes. His account of the Joycean machine is, I think, an important contribution to our understanding of the conscious mind. But I do want to maintain that it is not the whole story. Indeed, it is doubtful whether Joycean processes should count as genuinely *cognitive* at all. The Joycean machine, as Dennett characterizes it, works by association and self-stimulation.[21] Inner verbalizations do not influence behaviour in the way that beliefs do – by directly initiating actions they rationalize. Rather, they influence action only *indirectly*, by evoking conditioned behavioural responses. [22] It is not *saying* the words to yourself that has a cognitive role; it is *hearing* and *processing* them at a subpersonal level. Nor does the account accommodate one-off judgement and making up of mind. Telling yourself 'Beef is unsafe' does not

---

21 For an interesting anticipation of Dennett's account, see Skinner 1957, chapter 19.

22 It may be objected that opinions in my sense do not influence behaviour directly either: their influence is mediated by low-level beliefs about them. This is true, but the mediation here is of a special kind. Opinion-based processes are *realized in* low-level psychological ones, just as these in turn are realized in neurological ones. *At their own level* opinions directly influence behaviour, just as beliefs and desires do at *their* level.

constitute the formation of a new doxastic state – though it may help to cajole your neural subsystems into seeing that you decline the *steak tartar*.

So there is still a role for opinion as I have characterized it. Indeed, I think the account sketched here meshes quite nicely with the literature on private speech. There is, first of all, evidence that private speech is *multi-functional* (see Berk 1992 pp.41ff). So some private utterances could be expressions of opinion, even if others are self-stimulations. Secondly, it is widely accepted among Vygotskian theorists that private speech is a development of social speech – some even suggest that it exhibits *dialogic* properties (see, e.g., Ramirez 1992). This view, I think, harmonizes well with my story about opinion, and may help to dispel worries about its over-intellectualization. I said that if we are to act upon our opinions, then we need to know, or be able to work out, which actions and inferences they license. That is, we need to be able to distinguish sound inferential moves from unsound ones. How do we learn to do this? Well, in the course of linguistic interaction with our peers; the skills we need for conscious private reasoning are just those we need for engaging in reasoned argument with others. Indeed, we can think of the abilities involved in opinion-based reasoning as internalizations of skills first manifested in inter-personal argumentation (making assertions, defending them, calculating their commitments, accepting or rejecting these commitments, striving for consistency, and so on). The processes of opinion have an inherently dialogic character.

I have a final, brief, comparison to make. It is often argued that the commonsense concept of belief is a strongly *realist* one – that we conceive of beliefs as functionally discrete states which can be individually formed and lost, and selectively activated in reasoning and decision-making (see Ramsey et al. 1991; Davies 1991). Now Dennett is well-known for his hostility to this view, which is, he suggests, an illusion generated by our opinion-forming habits (Dennett 1994, p.241). But if we take seriously the idea that opinions constitute a genuinely distinct level of cognition, then an intriguing possibility opens. Perhaps our realist intuitions are true at least of *opinions*, if not of beliefs. For opinions, understood as premising policies, fit the realist profile very well. Premising policies can be selectively undertaken (by datable acts of policy adoption) and selectively abandoned (either by deliberate repudiation or just by forgetting them). And since they can be selectively recalled, they can be selectively executed, too. So they will be functionally discrete. Moreover, their discreteness depends, not on facts about the internal architecture of our brains, but upon the shape of our personal-level reasoning strategies. So with our new theory of opinion, we get an extra dose of realism about folk psychology for free. This is, I think, an attractive consequence.

7. Conclusion

I have tried to do two things in this paper: to show that a two-level theory of belief is worth taking seriously, and to sketch a constitutive account of one of those levels. The two aims are connected. For I suspect that one reason for the relative neglect of two-level theories is that it has been hard to see how the levels would interact. It is only

when we think of the higher level as *implemented in* the lower one that this problem becomes really tractable. The present account is, of course, merely a sketch. Properly developed, a two-level theory will have far-reaching implications, some of which must be explored before it can be properly assessed. This is just a start.[23]

References

BAIER, A. (1979). Mind and change of mind. *Midwest Studies in Philosophy*, 4, 157-176.

BENNETT, J. (1964). *Rationality*. London: Routledge and Kegan Paul.

BERK, L.E. (1992). Children's private speech: an overview of theory and the status of research. In R.M. DIAZ & L.E. BERK (Eds) *Private speech: from social interaction to self-regulation*, (pp. 17-53). Hillsdale, NJ: Erlbuam.

CARRUTHERS, P. (1996). *Language, thought and consciousness.* Cambridge: Cambridge University Press.

CARRUTHERS, P. (1998). Conscious thinking: language or elimination? *Mind and Language*, 13.

CLARKE, D.S. (1994). Does acceptance entail belief? *American Philosophical Quarterly*, 31, 145-155.

COHEN L.J. (1992). *An essay on belief and acceptance.* Oxford: Oxford University Press.

COSMIDES L., & TOOBY, J. (1992). Cognitive adaptations for social exchange. In J.H. BARKOW, L. COSMIDES & J. TOOBY (Eds) *The adapted mind* (pp. 163-228). New York: Oxford University Press.

DAVIDSON, D. (1975). Thought and talk. In S. GUTTENPLAN (Ed.) *Mind and language* (pp.7-23). Oxford: Oxford University Press.

DAVIDSON, D. (1982). Rational animals. *Dialectica*, 36, 318-27.

DAVIES, M. (1991). Concepts, connectionism and the language of thought. In W. RAMSEY, S. STICH & D.E. RUMELHART (Eds) *Philosophy and connectionist theory*, (pp. 229-257). Hillsdale, NJ: Erlbuam.

DE SOUSA, R. (1971). How to give a piece of your mind: or, The logic of belief and assent. *Review of Metaphysics*, 25, 52-79.

DENNETT, D.C. (1969). *Content and consciousness.* London: Routledge & Kegan Paul.

DENNETT, D.C. (1978). *Brainstorms: philosophical essays on mind and psychology*. Montgomery, VT: Bradford Books (English edition, Sussex: Harvester).

DENNETT, D.C. (1987). *The intentional stance*. Cambridge, Mass.: MIT Press.

DENNETT, D.C. (1991a). *Consciousness explained*. Boston: Little, Brown. (English edition, London: Viking.)

DENNETT, D.C. (1991b). Mother nature versus the walking encyclopedia: a Western drama. In W. RAMSEY, S. STICH & D.E. RUMELHART (Eds) *Philosophy and connectionist theory* (pp. 21-30). Hillsdale, NJ: Erlbaum.

[23] I am grateful to William Bechtel, George Botterill, Peter Carruthers and two anonymous referees for *Philosophical Psychology* for their comments on earlier drafts of this paper.

DENNETT, D.C. (1991c). Two contrasts: folk craft versus folk science, and belief versus opinion. In J.D. GREENWOOD (Ed.) *The future of folk psychology: intentionality and cognitive science* (pp. 135-148). Cambridge: Cambridge University Press.

DENNETT, D.C. (1991d). Real patterns. *The Journal of Philosophy*, 88, 27-51.

DENNETT, D.C. (1994). Self-portrait. In S. Guttenplan (Ed.) *Companion to the philosophy of mind* (pp. 236-244). Oxford: Blackwell.

FODOR, J. (1987). *Psychosemantics*. Cambridge Mass.: MIT Press.

FRANKISH, K. (1998). Natural language and virtual belief. In P. CARRUTHERS & J. BOUCHER (Eds) *Language and thought: interdisciplinary themes* (pp. 248-269). Cambridge: Cambridge University Press.

GAULKER C. (1994). *Thinking out loud*. Princeton, NJ: Princeton University Press.

GOLDMAN, A.I. (1970). *A theory of human action.* Princeton, NJ: Princeton University Press.

HORST, S. (1995) Eliminativism and the ambiguity of 'belief'. *Synthèse*, 104, 123-145.

JEFFREY, R.C. (1970). Dracula meets wolfman: acceptance vs. partial belief. In M. SWAIN (Ed.) *Induction, acceptance, and rational belief* (pp. 157-185). Dordrecht: Reidel.

KAPLAN, M. (1996). *Decision theory as philosophy*. Cambridge: Cambridge University Press.

MAHER, P. (1993). *Betting on theories*. Cambridge: Cambridge University Press.

MALCOLM, N. (1973). "Thoughtless brutes". *Proceedings and Addresses of the American Philosophical Association*, 1972-73, 46 (November), pp. 5-20.

O'SHAUGHNESSY, B. (1980). *The will: a dual aspect theory*. Cambridge: Cambridge University Press.

RAMIREZ, J.D. (1992). The functional differentiation of social and private speech: a dialogic approach. In R.M. DIAZ & L.E. BERK (Eds) *Private speech: from social interaction to self-regulation* (pp. 199-214). Hillsdale, NJ: Erlbuam.

RAMSEY, W., STICH, S., and GARON, J. (1991). Connectionism, eliminativism and the future of folk psychology. In J.D. GREENWOOD (Ed.) *The future of folk psychology* (pp.93-119). Cambridge: Cambridge University Press.

SKINNER, B.F. (1957). *Verbal behaviour*. New York: Appleton-Century-Crofts.

WASON, P. (1966). Reasoning. In B.M. FOSS (Ed.) *New horizons in psychology* (pp. 135-151). Harmondsworth: Penguin.

WILLIAMS, B. (1973; orig. 1970). Deciding to believe. *Problems of the self* (pp. 136-151). Cambridge: Cambridge University Press.

VYGOTSKY, L.S. (1962; orig. 1934). *Thought and language* (Ed. and trans. E. Hanfmann & G. Vakar). Cambridge, MA: MIT Press.