

Open Research Online

The Open University's repository of research publications and other research outputs

Natural language and virtual belief

Book Section

How to cite:

Frankish, Keith (1998). Natural language and virtual belief. In: Carruthers, Peter and Boucher, Jill eds. Language and thought : Interdisciplinary themes. Cambridge, UK: Cambridge University Press, pp. 248–269.

For guidance on citations see [FAQs](#).

© 1998 Cambridge University Press

Version: Accepted Manuscript

Link(s) to article on publisher's website:

<http://www.cambridge.org/catalogue/catalogue.asp?isbn=0521637589>

Copyright and Moral Rights for the articles on this site are retained by the individual authors and/or other copyright owners. For more information on Open Research Online's data [policy](#) on reuse of materials please consult the policies page.

oro.open.ac.uk

Natural Language and Virtual Belief*

KEITH FRANKISH

Abstract

This chapter outlines a new argument for the view that language has a cognitive role. I suggest that humans exhibit two distinct kinds of belief state, one passively formed, the other actively formed. I argue that actively formed beliefs (*virtual beliefs*, as I call them) can be identified with *premissing policies*, and that forming them typically involves certain linguistic operations. I conclude that natural language has at least a limited cognitive role in the formation and manipulation of virtual beliefs.

1 Introduction

It is sometimes claimed that we can think *in* natural language – that natural language sentences can act as vehicles of thought and that we can think by tokening them – usually in the form of silent auditory or articulatory images (see, for example, Bickerton, 1995; Carruthers, 1996; Harman, 1973.) There is some introspective evidence for this view (which is sometimes called the *cognitive conception* of language), and, given a certain view of what it is for thought to be conscious, a powerful argument can be run to the conclusion that conscious propositional thinking occurs in natural language (see Carruthers, 1996).

Now, defence of the cognitive conception usually proceeds in two stages. First, it is argued that thought-processes are computational in character, the computations in question being understood as occurring at a subpersonal level. It is then argued that, in some thought-episodes at least, the *medium* of computation is natural language. In this chapter I shall outline an alternative way of arguing for the cognitive conception. The argument will not require a defence of computationalism, but will start with the claim that some intentional states are *actively formed*.

The chapter falls into two halves. The first aims to show that some beliefs, or belief-like states, are actively formed. I begin by describing two psychological phenomena which seem to have an active dimension. Then, following a suggestion by Daniel Dennett, I go on to develop a unified account of these phenomena, based around the idea that there exists a class of *virtual*

* This is an e-print of a chapter published in P. Carruthers and J. Boucher (eds), *Language and Thought: Interdisciplinary Themes*, Cambridge University Press, 1998 (pp. 248-69). It includes some minor revisions.

beliefs, formed by actively adopting policies of *premissing*. The second half of the chapter aims to show that virtual belief formation will typically involve natural language. I argue that, although there is no *conceptual* link between premissing and natural language, there are grounds for thinking that premissing – and hence virtual belief formation – will in fact usually take a linguistic form. Some uses of language, I conclude, have a cognitive role in the formation and manipulation of virtual beliefs. A final section briefly compares this defence of the cognitive conception with the more standard one, and suggests that it is to be preferred.

2 Virtual Belief

2.1 Change of mind

In this section I want to tease out a relatively neglected strand of our commonsense psychological discourse – our talking of *making up* and *changing* our minds. We frequently refer to such events, and tend to speak of them as free, intentional actions (we urge the indecisive to *make up* their minds and blame the irresolute for *changing* them). Changes and makings up of mind have, however, received little philosophical attention – perhaps because it has been assumed that they can be identified with the formation and revision of intentions to action.¹ It is, of course, true that many changes and makings up of mind occur in practical reasoning. So, for example, I might *make up* my mind to go to Italy this year, or *change my mind* about which political party to vote for. And such cases do seem to involve simply the formation or revision of action-plans. It would be hasty, however, to assume that all changes and makings up of mind are such. We can make up our minds about what to *think*, as well as about what to *do*. Thus I can *change my mind* about the truth of a theory, or *make up my mind* that a certain politician is untrustworthy.²

If this is right, then some changes of mind are changes in belief, or in something very like belief. But not all changes in belief are changes of mind. As Annette Baier emphasizes, in one of the few philosophical analyses of the subject, one does not make up, or change, one's mind whenever one acquires new beliefs or revises old ones (Baier, 1979). Baier lists several distinguishing

¹ There has, in fact, been a good deal of discussion of the conditions under which deliberate belief revision would be *rational* – particularly in the literature on *acceptance* (see below). But there has been little attempt to explain exactly *how* such revisions can be carried out, or how they are related to other cognitive states and processes.

² Interestingly, this extends to desires, as well. We can make up and change our minds about what we *want*, as well as about what we *think*; and we often tell people – children especially – to do just that.

features of changes of mind. First, they typically occur in *anticipation* of experience, rather than in reaction to it. A person who mistakenly thinks the ice will support them and ends up sinking can be said to have *learned better*, but not to have *changed his or her mind* (Baier's example). To be able to change your mind is, in part, to be able to anticipate experience and to alter erroneous beliefs before the world changes them for you. Secondly, changes of mind are *reflective* rather than routine. Suppose I misread the map, and you correct me (again, the example is Baier's). We would not say that, in accepting the correction, I had *changed my mind* about the route. Part of the reason for this, Baier suggests, is that I never needed to *make up* my mind about it. Where we defer to an authority or trust to a routine calculation, there is no call for us to make up our minds. Similarly, in practical reasoning, changes of plan that are dictated by circumstances or by longer-term plans do not count as changes of mind (taking a different route to avoid a traffic jam, say, or giving up one's job if one wins the lottery). Thirdly, changes of belief that result from large-scale cognitive or emotional upheaval are not changes of mind (Saul did not *change his mind* on the road to Damascus). Fourthly, Baier suggests, change of mind involves a kind of *commitment* – as to a rule, convention, or plan. In cases involving practical reasoning this is easy to see: the commitment in question will be to a plan of action. Baier does not explain, however, what sort of commitment is involved in *theoretical* changes of mind – cases in which one decides, not what to do, but what to think.

A distinctive feature of the various cognitive revisions excluded from the class of changes of mind, Baier suggests, is that they are *passively* undergone or assented to: the changes involved are too unreflective, too routine or too overwhelming for active deliberation to be needed. A change of mind, on the other hand, follows upon a reevaluation of one's options, and involves a kind of free, creative judgement.

If Baier's analysis is on the right lines, then, it seems, we exhibit two rather different kinds of cognitive state, not clearly distinguished in everyday psychological discourse: passive belief, formed in an automatic and unreflective way, and another kind of cognitive state, formed by freely making up or changing one's mind, and involving a kind of commitment.³

2.2 *Acceptance*

The second strand of thought I want to introduce is the philosophical – and, more specifically, epistemological – literature on *acceptance*. 'Acceptance' here is

³ Note that the distinction between these two kinds of state is drawn in terms of *form*, not *content*. I shall assume, for simplicity's sake, that the same range of contents can be entertained both passively and actively.

a technical term for a kind of propositional attitude allied to, but distinct from, belief. There is no universal agreement about the properties of acceptance, but certain claims recur.

First, acceptance is qualitative, not quantitative. When we think about the role of belief in guiding action and decision-making, it is very tempting to suppose that belief is a matter of *degree*, with individual beliefs reflecting subjective probabilities of the sort assigned by Bayesian decision theory. (A probabilistic view of this kind also seems to be required if we are to resolve certain paradoxes of belief – in particular those of the lottery and the preface.) On the other hand, commonsense epistemology seems to treat belief as a qualitative, ungraded state. Unqualified belief is, for example, often assumed to be a necessary condition for knowledge. (It would seem odd to say that one *knows* that *p*, but does not *fully* believe it). And we tend to think of our sincere assertions as claims to *truth*, rather than as expressions of subjective probability. Many epistemologists think that the way to reconcile these intuitions is to suppose that everyday belief-talk fails to distinguish two distinct kinds of cognitive state – *belief* and *acceptance* – one of which is quantitative, the other qualitative (see, in particular, Levi, 1967; de Sousa, 1971; Kaplan, 1981.)

Secondly, states of acceptance, unlike beliefs, are often supposed to be *actively formed* (de Sousa, 1971; Bratman, 1992). The idea is that, whereas belief is generated by automatic subpersonal processes, acceptance results from the act of deliberately *endorsing* or *assenting* to a proposition. This would explain why acceptance is qualitative (for any given proposition, one either has, or has not, actively accepted it).

Thirdly, acceptance is often characterized as a *behavioural* state (de Sousa, 1971; Kaplan, 1981). It is easy to show that, if acceptance were both a qualitative state *and* a state of confidence of any degree lower than certainty, then a rational agent would be required to accept contradictory propositions (the paradoxes of the lottery and the preface again; see Kaplan, 1981). One way to avoid this conclusion is to suppose that acceptance, unlike belief, is not a state of *confidence* at all, but a *behavioural* state. It has been suggested, for example, that accepting a proposition involves being disposed to bet on its truth (de Sousa), or to defend it for epistemic purposes (Kaplan), or to use it as a premise in one's reasoning (Cohen).

The view that acceptances are behavioural states harmonizes nicely with the view that they are actively formed. That we can actively choose to display certain patterns of behaviour is uncontroversial. That we can choose to have a certain degree of confidence in a proposition is, on the other hand, dubious, if not downright incoherent.

2.3 *Virtual Belief*

We have, then, two strands of thought, one implicit in commonsense psychological discourse, the other arising from reflection on the nature of belief, each of which posits the existence of a distinct kind of actively formed cognitive state. And the states they posit have much in common. As well as being actively formed, both are qualitative (making up one's mind, like acceptance, seems to be all-or-nothing), both require a degree of cognitive sophistication, and both involve a kind of behavioural commitment. It is very tempting, then, to identify them. (The suggestion is first made, I believe, in Dennett, 1978, ch.16, though it may be implicit in some earlier work on acceptance, particularly de Sousa, 1971). We could then appeal to acceptance-theory to provide an account of the commitment involved in theoretical change of mind. Dennett takes this line, suggesting that we complement Baier's account of change of mind with de Sousa's account of acceptance as a *bet on the truth of a sentence*. The states that result from such bets he calls *opinions*.

This suggestion is, I think, on the right lines. However, truth-betting (or 'sentence-collecting' as Dennett sometimes puts it) seems inadequate as a model for change and making up of mind.⁴ There must, for example, be more to making up one's mind than simply betting on a *sentence*. A monoglot English speaker cannot make up his or her mind by betting on a suitably vouched-for *Russian* sentence. In any case, it would be a tactical mistake for me to begin by identifying makings up and changes of mind with attitudes to linguistic items. For I want to *argue* that language does *in fact* play a constitutive role in these processes. (This conclusion really does have to be argued for – Baier, for example, rejects it; 1979 pp.166-7.) It would be unpersuasive, then, simply to *define* them as linguistic – unless, of course, it could be shown that no other characterization was available.

And, in fact, alternative characterizations are ready-to-hand. In particular, I want to make use of Cohen's (1992) account of acceptance as *premissing*. States of acceptance, Cohen argues, are not only actively initiated, but involve an extended 'pattern, system, or policy' of mental action (Cohen 1992, p.12). To accept a proposition, Cohen explains, is to commit oneself to taking it as a premise or inference-licence in one's conscious deliberations. Such deliberations, he writes, will be

guided, implicitly or explicitly, by the premisses or inference-rules that you have accepted previously. So in such cases you will be deliberately schooling your present thoughts to fit such premisses or rules, and you will evaluate your hypotheses as correct or incorrect, probable or improbable, in accordance with those criteria (1992, p.23).

⁴ For an analysis of Dennett's views on opinion, see my 1996.

For Cohen, then, accepting a proposition involves committing oneself to a series of further personal actions – to deliberately regulating or 'schooling' one's thoughts in order to keep them in line with the premise. Presumably, this means calculating what conclusions the premise entails or excludes and then making appropriate further acts of acceptance and intention-formation. In due course we will need to consider just how these calculations might be made. (Cohen suggests that it involves applying learned inference-rules, though, as we shall see, this view is too restrictive.) But for the moment this rough characterization will suffice.

I want to suggest, then, that making up one's mind about a matter of fact, p , involves accepting that p in Cohen's sense – that is, forming the intention to take p as a premise in one's conscious reasoning. This suggestion harmonizes well with Baier's analysis of a change of mind as a kind of *commitment*. Deciding to take p as a premise means choosing to impose certain normative constraints upon one's future deliberations. The proposal also nicely integrates our accounts of change of mind in practical and theoretical reasoning: both involve committing oneself to a plan of some kind – in one case to a plan of action, in the other to a strategy of reasoning. (Indeed, since deciding to adopt a reasoning strategy is just forming a certain kind of intention to action, the one is a species of the other.)

It may be objected that I have ignored one feature of acceptance that is often stressed by writers on the subject, and by Cohen in particular. This is its *context relativity* (Bratman, 1992; Cohen, 1992 pp.12-13). A person may decide to accept a proposition in one context (because their job requires them to treat it as true, say), but not in another (for example, where their aim is to ascertain the truth). And this fact seems to rule out a straightforward identification of acceptance with making up of mind. (The lawyer who decides, for professional reasons, to adopt the premise that a client is innocent, does not thereby *make up* his or her mind that the individual in question is innocent.) This objection is not fatal, however. It does not show that makings up of mind are not acceptances – only that not all acceptances are makings up of mind. We can still think of makings up of mind as a *subset* of acceptances – those, perhaps, that are motivated in a suitable way. For people have general epistemic ends, as well as short-term pragmatic ones. They desire to possess truths (or, perhaps, if they are epistemologists, to maximize *cognitive utility*). And it is such general ends that motivate them when they make up their minds or decide what to think about something. So I suggest that we identify makings up of mind with

the class of acceptances that are motivated in this way by general epistemic ends.⁵

Suppose, then, that acts of change of mind, making up of mind and suitably motivated acceptance all introduce essentially the same kind of state – a behavioural state involving commitment to a policy of premising. I suggested that folk psychology does not explicitly distinguish this state from passively formed (or, as I shall sometimes put it, ‘low-level’) belief. So what exactly is the relation between the two states? Well, from a behavioural point of view, they will be almost indistinguishable. (This would explain, of course, why they are often conflated.) A person who premises that p will make much the same inferences and perform much the same actions as one who has the low-level belief that p . The principal difference between them will be that the premiser deliberately *guides* their inferential processes in order to keep them in line with the premise, whereas the ordinary believer leaves them to subpersonal control. Borrowing a term from computer science, we might say that the premiser intentionally *emulates* the inferential processes of a low-level believer. Or, to use another computing term, we might say they have a *virtual* belief.⁶

The term is quite appropriate. A virtual machine is one, in Daniel Dennett's phrase, 'made of rules rather than wires' – that is to say, formed by programming a flexible low-level physical system so as to cause it to display high-level functional states characteristic of a different type of machine (that is, to *emulate* that machine). Thus by executing an appropriate program, a personal computer can mimic the behaviour of a card-index, a typewriter, or even another computer. Similarly, by executing an appropriate reasoning strategy, a person can mimic the inferential behaviour of someone with a certain belief. Just as a computer simulation is driven by the system's low-level programming instructions, so premising behaviour is driven by the premiser's low-level beliefs and desires. Thus, having decided to adopt a certain premise, the premiser *wants* to stick to this policy, *believes* that doing so requires accepting certain conclusions and rejecting others, and acts accordingly. Note that I am not suggesting that premisers will entertain these beliefs and desires *at a conscious level*. Typically, they will not. When one makes up one's mind about the truth or falsity of a proposition, all one consciously thinks about is the proposition and the evidence for and against it. I do claim, however, that making up one's mind involves manipulating premises (or representations of

⁵ I think this deliberately conservative policy on the role of pragmatic factors in making up of mind best captures our ordinary way of speaking; but the formulation can easily be revised so as to allow pragmatic factors a more substantial role. (In fact, all I need to claim is that makings up of mind are a *subset* of acceptances; discerning the exact boundaries of this subset is a matter for another time.)

⁶ I owe the term 'virtual belief' to Chris Hookway.

them) in ways which can best be explained by reference to one's low-level beliefs and desires about those premises.

I suggest then, that premising behaviour creates a kind of *virtual inference engine* which processes virtual beliefs. Note that a machine may be 'virtual' in this sense without lacking any essential properties of the real thing. To say that a machine is *virtual* is not to say that it is not *real*, only that it does not have a dedicated physical architecture. Thus in saying that virtual believers *emulate* or *mimic* belief, I do not mean to suggest that they do not really *have* the emulated beliefs. Rather, I want to suggest that emulating a certain belief is a distinct way of having that belief.⁷

Now it may seem that I am begging an important question here. For since premising is under voluntary control, it follows that, in suggesting that some premisings constitute a kind of *belief*, I am endorsing a form of *voluntarism* about belief. And there is a large philosophical literature devoted to showing that voluntarism is false, even incoherent. (For some anti-voluntarist arguments, see Williams, 1973; Winters, 1979; Pojman, 1985, and Bennett, 1990.) For example, voluntarism is often taken to involve the claim that we can induce beliefs in ourselves for purely practical reasons and regardless of the evidence for their truth. And this claim does, indeed, look very implausible. (It is hard to see how one could consistently think of a state induced in this way *as* a belief; see Williams, 1973 and Pojman, 1985). But nothing in my account commits me to the claim that virtual beliefs can be induced without regard to the evidence. I identified virtual belief formation with acts of acceptance that are motivated by general epistemic ends. And for an act of acceptance to be motivated in this way is just for it to be motivated *by* the evidence. (I cannot believe that I will advance my general epistemic ends by accepting that *p*, unless I believe that there is reasonably good evidence for the truth of *p*.) So while I accept that some kinds of belief formation are active, I deny that they are unconstrained by evidence. I am thus committed only to a *weak* form of voluntarism.⁸

⁷ The distinction between belief and virtual belief has some affinities with Dan Sperber's distinction between *intuitive* and *non-intuitive* beliefs (Sperber, 1997). Intuitive beliefs, Sperber argues, are those generated by spontaneous and unreflective processes – those acquired through perception, or the verbal communication of information which could have been presented perceptually, or by means of unconscious inferential processes. Non-intuitive beliefs, by contrast, are those that are acquired by conscious and deliberate inference or through the communication of ideas which could not have been presented perceptually. (Complex theoretical propositions are typical examples). Non-intuitive beliefs are thus formed by the sort of deliberate premising that is involved in virtual-belief formation.

⁸ Another objection centres on the voluntarist claim that beliefs can be induced *directly*, simply by *willing* to acquire them. Acquiring a belief, the objector points out, typically seems to be something that *happens to* one, rather than something one *does*. This is especially true of beliefs arising from perception, memory and certain basic inferential processes (see, for

2.4 Levels of mind

In this section I want to draw out one consequence of my story for theories of mind and mental processing. The consequence is salutary and constitutes, I think, another consideration in favour of the account.

If there are two distinct levels of mentality, then there may be distinct kinds of *processing* underlying them. Cohen suggests that this is indeed the case. The distinction between acceptance and belief, he claims, corresponds neatly to that between *digital* and *connectionist* processing. Because acceptance is a discrete, sequential, rule-governed process, it can be adequately modelled by digital computer programs. Belief states, on the other hand, being graded and overlapping, are better modelled by connectionist systems (Cohen, 1992, pp.56).

Cohen does not develop the suggestion, and, indeed, a lot of work would be required to turn it into a serious proposal. (In particular, it would be important to think hard about the relation between the two processing levels). But I think that it points in the right direction, and that, properly construed, a two-level theory offers the best hope for reconciling a connectionist cognitive science with a commonsense view of the mind. Let me explain briefly.

Many cognitive scientists believe that connectionist systems of some variety will eventually afford our best models of the human mind. Certainly, they seem to be the most neurologically plausible. But there is a persisting difficulty in seeing how connectionist systems can exhibit the sort of psychological states and processes which we ordinarily attribute to each other. The problem is that commonsense psychology seems to be strongly *realist*. It supposes that beliefs and other psychological states are real, functionally discrete, internal states, which can be individually formed, activated and erased. One of the key pieces of evidence for this claim is the fact that we often pick out individual beliefs and desires as causally responsible for particular actions, even when we know that the agent possessed other beliefs and desires which would equally have justified the action. For example, I may have several long-standing beliefs, each of which would independently justify my going into town (I need to buy food, want to go to the bank, and have a message to deliver), yet only one of these reasons (the need to buy food, say) may be *the* causally effective reason for my actually going into town this afternoon. But if beliefs

example, Pojman, 1985). Again, I can concede the point. I contend only that *some* types of belief can be induced voluntarily – the kind of reflective or theoretical beliefs that result from changing or making up one's mind. Nor do I claim that we can form virtual beliefs *simply* by willing to do so. Virtual belief formation is, we may say, *intentionally indirect*. That is to say, one forms a virtual belief *by* doing something else – by intentionally committing oneself to a policy of premising.

can be individually effective in this way, then they are functionally discrete; and if they are functionally discrete, then, it seems, they must have discrete *internal* representations (see Ramsey *et al.*, 1991).

Many of the most effective connectionist systems, however, appear to lack discrete internal states of this kind. These systems store information in a distributed, heavily overlapping way, which seems to preclude the selective activation of individual items of knowledge. Thus there is an apparent tension between connectionism and common sense – a tension which both sides have an interest in diffusing. Realists about folk psychology would welcome the news that connectionism poses no threat to them, while connectionists for their part would be reassured to know that they are not going to have to deny that people really do have beliefs and other folk psychological states.

There are two popular strategies for diffusing the tension. One is to argue that connectionist systems do, in fact, despite appearances, possess functionally discrete internal representations, at least of *occurrent* thoughts (see Botterill, 1994; Clark, 1990; O'Brien, 1991). The second, more radical, response, is to adopt a weak, quasi-behaviouristic, reading of folk psychological talk (see, for example, Clark, 1993). Neither strategy, I think, is wholly satisfactory. (The first offers, at best, a partial solution, while the second requires us to accept a significant weakening of our commonsense intuitions.)

But if our realist intuitions relate principally to *virtual belief*, rather than to its low-level counterpart, then we have, I think, another option. Since virtual beliefs are policy adoptions, their processing takes place entirely at a *personal* level. It involves a person *doing* various things: *endorsing* a premise, *keeping track* of it, *working out* what conclusions it entails, and so on. So virtual beliefs are, in a sense, *behavioural* states (stretching the term 'behaviour' a little to include certain kinds of mental behaviour such as silent acceptance and conscious deliberation). Now the existence of various kinds of human *behaviour* is not an issue between the folk psychological realist and the connectionist. The debate between them is over the nature of the processing required to support that behaviour. In so far as realists are concerned only with virtual belief, then, they can have no quarrel with connectionists.

Moreover, virtual beliefs will be functionally discrete in the way that realists suppose beliefs to be. Premising policies can be selectively adopted, executed and abandoned. Consider, for example, the role of virtual belief in guiding action. A particular premise will get to influence behaviour only if the premiser deliberately employs it in their practical reasoning. (At least, that is the only way it will get to influence behaviour *as a premise*; it may have unintended side-effects.) And it will be deliberately used in reasoning only if it is consciously *recalled* at appropriate moments. But relevant premises will not always be recalled at such moments. Suppose, for example, that I have separately

endorsed the premises *I need to buy food* and *I need to go to the bank*, each of which would independently warrant my going into town this afternoon. And suppose, too, that right now I recollect just one of these premises – the need to buy food, say. (I need not explicitly think of it *as* a premise; it is enough that I recollect it with an appropriate degree of commitment and am disposed to use it in inference.) I recognise that this requires me to go into town, and decide to go. Here we can say that, although I had endorsed both premises, and would, if prompted, have acknowledged that I was committed to both of them, nevertheless I *acted upon* only one of them. So, we can say that on this occasion only one of my two virtual beliefs was causally active. Moreover, we can say all of this without saying anything about my subpersonal neural states. (Thus we can explain my failure to act upon my belief about the bank as due to the fact that I temporarily *lost track* of the relevant premise, failing to recollect it and use it in inference.) So we can explain how virtual beliefs can have a selective causal influence on behaviour without having to suppose that it involves the selective activation of discrete neural subcomponents.

Of course, the proposed strategy for reconciling connectionism and common sense will work only if our commonsense commitment to realism does not extend to low-level belief.⁹ And it is not clear that this is so. For the story requires us to suppose that the personal-level actions and events involved in the processing of virtual beliefs have fairly determinate psychological characterizations – that they can properly be described as acts of *deciding to adopt a premise p*, *recalling that one has accepted p*, *working out that p entails q*, and so on. Only if this is so will we be able to speak of a subject selectively recollecting and employing in inference one of a number of semantically relevant premises. And the worry, of course, is that the characterizations just given are in terms of the agent's *low-level* psychological states. So it might seem that we will, after all, have to buy into a fairly strong form of realism about low-level belief.

I think this worry is misplaced, however. For although the story requires us to ascribe certain low-level beliefs to virtual believers, it does not require us to give semantic characterizations to neural components or subpersonal processes. In fact, it requires only two assumptions: (1) that premising policies, like other policies, can be individually formed, recalled and executed, and (2) that the actions involved in forming and executing a premising policy can be given fairly determinate semantic characterizations – sufficient, at least, to individuate the policies concerned. It is possible to accept both of these assumptions without endorsing a strong form of realism about low-level belief. (Note that if we typically give our premises linguistic form, as I shall argue we

⁹ The strategy does not, of course, assume the *falsity* of realism about low-level belief; but it does assume the absence of any commonsense commitment to its *truth*.

do, then condition (2) above will be relatively uncontroversial: premising policies will be adequately individuated by their associated linguistic representations.)¹⁰

3 Natural language

3.1 Premising and language

I avoided characterizing virtual belief as an essentially linguistic state, preferring a more neutral characterization in terms of premising. Of course, this leaves open the possibility that premising is *itself* a language-dependent activity. Cohen takes this view, and I am going to argue for a slightly weaker version of the same claim. Premising, I shall argue, does not *have* to involve language, but will in fact typically do so. Given that premising is constitutive of virtual belief formation, this conclusion will give us our limited vindication of the cognitive conception of language.

Note that this defence will extend only to *virtual* belief. For simplicity's sake I shall assume that low-level belief does *not* constitutively involve natural language. I shall assume also that there is no limit to the range of contents that can be entertained as objects of low-level belief. To deny either of these assumptions would be to embark on a much wider defence of the cognitive conception.

Given these assumptions, however, it looks as if it will be difficult to show that premising is linguistic. If one can *entertain* the content p without putting it into words, then there is no obvious reason why one could not decide to take p as a *premise* without putting it into words. One might weigh up the various bits of evidence one has for p , think about how it coheres with one's other

¹⁰ The reader may ask why, given the two assumptions just mentioned, single-level theorists cannot say everything I say about the functional discreteness of beliefs. They can say, of the case mentioned in the text, that I recalled that I needed to buy food (a personal episode with a determinate semantic characterization) and that this recollection caused me to go into town. I failed to recall that I needed to go to the bank and so was not moved to act by that need. This position is, in effect, that adopted by advocates of the first conciliatory strategy mentioned in the text. My reasons for rejecting it are two-fold. First, although it accounts for the functional discreteness of *occurrent* thoughts, it does not do the same for *standing-state* beliefs. In particular, I suspect that it cannot give a satisfactory account of the status of the temporarily dormant beliefs, such as my belief about the bank. Secondly, it is not clear from this story why a belief should have to be *explicitly recalled* in order for it to influence behaviour (indeed, it would be very odd to say that it has). At any rate, the story is incomplete until supplemented with an account of the cognitive role of explicitly recalled beliefs. This, I suggest, is just what virtual belief theory provides.

premises and beliefs, and finally decide to accept it as a premise in one's reasoning. Of course, adopting any sort of premising policy would probably require a degree of conceptual sophistication not to be found in creatures without language. But even so, this would not establish the conclusion we want. To say that premising would be beyond the reach of languageless creatures is not to say that it is *itself* a linguistic activity.

It would be rash, however, to conclude from this that language is inessential to premising. For even if there is no need to give a proposition linguistic form in order to decide to *adopt* the policy of treating it as a premise, it might still be necessary to do so in order to *execute* that policy. We will have to look more closely at just what premising involves.

Premising that p means taking p as a given in one's conscious explicit reasoning – *schooling one's thoughts* to fit p , as Cohen puts it. This, I suggested, involves calculating what conclusions p entails and excludes and then making appropriate further acts of acceptance and intention-formation.¹¹ Premisers, then, will need some way of making these calculations – some way of *evaluating* their premising commitments. (They would also need to make such evaluations in the course of deciding whether or not to accept, or continue to accept, a premise – since it would be important to know exactly what a positive decision would commit them to.) Of course, if they *believed* their premises (believed them, that is, in a passive, low-level way), then such conclusions might occur to them spontaneously. The required inferential operations would be performed by automatic subpersonal mechanisms. But premisers undertake intentionally to *emulate* the operation of those mechanisms. They need not, of course, employ the very same algorithms as those employed at a subpersonal level, though they will need to use ones that generate the same results.¹² And for this they must have *personal* mastery of some technique, or set of techniques, for deriving normatively warranted conclusions from their premises. Of course, premisers need not work out *all* the entailments of their premises (a tedious task, given that there will be an infinite number of them). Rather, they will need to employ techniques which reliably generate some of their more useful and informative entailments (working out that from p and q one can infer $p \& q$ might be useful; working out that one can infer $r \rightarrow p$, $s \rightarrow p$, $t \rightarrow p$... probably would not).

The obvious strategy here would be to apply learned rules of inference. This is what Cohen envisages. Acceptance, he says, is 'consciously guided by

¹¹ Unless, of course, these conclusions are so unpalatable as to induce one to repudiate the original premise.

¹² Or, at any rate, which they *take* to generate the same results; some people may, after all, be rather bad at emulating belief.

voluntarily accepted rules' (1992, p.56). And these rules, he argues, together with the premises upon which they operate, will necessarily be linguistic:

Premises and rules of inference have to be conceived in linguistic terms ...
That is how logic can get to grips with inference and formulate its principles
as rules for linguistic transformation (1992, p.12).

Thus, there is, he claims, an 'a priori conceptual requirement' that the objects of acceptance must be linguistically expressed – though not necessarily overtly vocalized.

One can see why Cohen thinks this. Inference-rules will need to have a degree of generality (rules for making *individual*, content-specific, inferences would be of little use). That is to say, they will have to specify *formal* operations. So, for example, one useful inference-rule might be *modus ponens*, which tells us that a pair of premises of the form 'If *p* then *q*' and '*p*', entail a conclusion of the form '*q*'. Another useful rule might be one to the effect that premises of the form '*x* is a dog' normatively warrant conclusions of the form '*x* is an animal'. (Although this reflects a semantical principle rather than a logical one, it is still specified in formal terms.) But if premising involves applying rules which, like these, specify formal operations, then it does seem to follow that premises must be linguistically expressed. For to apply formal inference-rules, one would need to have access to the *form* of one's premises. And this means that those premises must be represented in a medium to whose formal properties one has access. The obvious candidate for such a medium is natural language.

This is too swift, however. It is true that useful inference-rules must specify formal operations, and true, too, that we do not have direct *perceptual* access to the formal properties of our non-linguistic thoughts. But we might nonetheless have a kind of *reflective* access to them. For example, if I know that I accept the premise *If the butler did not murder the colonel, then the rector's wife must have done so*, then I can tell that I accept a conditional (assuming I have the concept *conditional*, of course). And if I also know that I accept the premise *The butler did not murder the colonel*, then I can tell that I accept the antecedent of this same conditional. And if I am familiar with *modus ponens*, then I can go ahead and apply it to these premises. If I can entertain all of these thoughts non-linguistically, then I can apply formal inference-rules non-linguistically too.

Cohen is wrong, then, to say that there is an *a priori conceptual* requirement for inference-rules to be linguistically expressed. Nevertheless, it seems likely that they will *in practice* be so. Semantical principles, in particular (that is, information about the inferential roles of particular concepts) will almost certainly be derived from reflection on our linguistic practices. It is just this sort of information that is codified in dictionaries. Logical principles, too, are usually acquired and applied as rules for performing linguistic transformations.

Besides, there may be another, more direct, argument for the involvement of language in premising. Most people are able to construct and evaluate arguments without applying explicit rules of inference. Instead, they make use of *practical* inferential skills, which can be thought of as embedding implicit or procedural knowledge of such rules. These skills are typically acquired in the course of linguistic interaction, and consist in being able to spot certain formal patterns in people's utterances, and to impose similar patterns upon one's own. Learning them is, in effect, learning how to engage in reasoned argument.

Now people with these skills could easily employ them to evaluate their premising commitments. Suppose, for example, that I have learned to spot patterns of inference which instantiate *modus ponens* and to classify them as valid. This skill could be entirely procedural; I do not need to be able to articulate *modus ponens*, but simply to recognize inferences of *that*, demonstratively identified, form as valid. Suppose, too, that I have learnt to regulate my own argumentative utterances in accordance with *modus ponens* – again without articulating the rule. So if I have asserted sentences of the form 'If *p* then *q*' and '*p*', I regard myself as licensed to assert a sentence of the form '*q*', too, and as obliged to refrain from asserting any sentence incompatible with one of that form. Then, if I were to *accept* sentences of those forms as premises, it would be natural for me to regard myself as committed to accepting the corresponding conclusion as a further premise, and to rejecting any premises that conflict with it. (Just *saying* the premises over to myself might prompt me to supply the dictated conclusion.) In this way it would be possible for me to execute a premising strategy without drawing on explicit theory at all, but relying entirely on my pre-theoretical argumentative skills.

Moreover, these skills will, I think, be essentially linguistic. It is very hard to think of any non-linguistic (or at any rate non-symbolic) personal routines which could embed implicit inference-rules. We have no perceptual access to the form of our thoughts, and cannot directly manipulate propositions in the way that we can manipulate their representations. And although we can have reflective knowledge of the formal properties of our thoughts, such knowledge would be useless unless accompanied by an equally reflective grasp of inferential rules. (For example, the knowledge that I accept a conditional and its antecedent will be of little use to me, unless I can remember what I am licensed to infer from such a conjunction.) Of course, we can think of our *subpersonal* processes of belief formation and revision as embedding implicit rules of inference. But these processes cannot be directly controlled and exploited by the premiser as part of a deliberate reasoning strategy (though, as we shall see, they can be *indirectly* exploited).

I think we can conclude, then, that premising strategies which exploit procedural knowledge of inferential rules will be language-based. Moreover,

such strategies are likely to be rather more widespread than ones exploiting explicit knowledge. People tend to acquire skills procedurally before they begin to theorize them. Certainly, many people can classify presented patterns of inference as good or bad without being able to say precisely what their goodness or badness consists in. Cohen seems to overlook this, supposing that inference-rules will always be acquired explicitly in the first place, even if they subsequently become second-nature (1992, pp.12, 23, 56). Nothing hinges on this claim, however, and dropping it would only strengthen his case for the linguistic dependency of premising.

3.2 *Simulation routines*

It looks, then, as if language-based forms of premising will be both simpler and easier to acquire than non-linguistic ones. We must not be too hasty, however. For there is another kind of inferential technique which seems to be both powerful and yet not essentially linguistic. This is *simulation*. Many philosophers, and some psychologists, believe that our skill in ascribing psychological states to other people depends on our ability to run cognitive simulations (see, for example, Goldman, 1989, 1992; Gordon, 1986; Heal 1986, 1994, Harris, 1989, 1992). The idea is this. In order to work out what another person is likely to think or do, one pretends to share their beliefs, lets one's inferential system run 'off-line' (so that its outputs are not passed to memory or motor control), and waits to see what conclusions one comes to. One then ascribes belief in these conclusions to the other person. What I want to suggest here is that one could also run *self-simulations* in order to evaluate one's premising commitments. That is to say, one could pretend that one *believes* one's premises, run an off-line simulation, see what conclusions one comes to, and then accept these conclusions as new premises.

For example, suppose that Miss Marple accepts the premise *The butler was at the public house all evening*. She runs an off-line simulation upon this premise and finds that the proposition *The butler did not murder the colonel* is generated as output. (This reflects her belief that the colonel was murdered at home, shortly after dinner.) Miss Marple now knows that if she *believed* the premise, then she would infer that the butler did not murder the colonel. Since she is committed to making just those inferences that she would make if she believed her premises, she concludes that she is committed to accepting that the butler did not murder the colonel.

Assuming it can be done at will, simulation seems to offer a powerful tool for evaluating premising commitments. Indeed, it would facilitate certain sorts of inference which the other procedures we have considered do not encompass. For example, there are no simple rules or procedures for deriving sound abductive inferences (that is, for working out which of various possible

hypotheses is the *best explanation* of a set of data). Making such inferences involves testing candidate hypotheses for coherence with one's network of background beliefs. Simulation, being an holistic process, would facilitate this. For the same reason, it would afford an excellent means of assessing candidate premises for acceptability prior to formal acceptance. Conscientious premisers would doubtless make full use of it. But there seems to be no crucial role here for language. If one can entertain a thought non-linguistically, then there is no obvious reason why one could not run a simulation from it non-linguistically, too.

There are some problems with simulation, however. First, premising policies guided by simulation will be rather more extensive than those guided by explicit rules or localized inferential procedures. In the latter cases, when one adopts a premise, *p*, one commits oneself to accepting whatever conclusions one can derive from *p* by applying appropriate rules or procedures. In simulation-based premising, by contrast, one commits oneself to accepting *whatever* conclusions one would draw if one *believed* that *p*. And this is a much wider commitment. For the outcome of simulation routines will be determined, not only by one's explicitly accepted premises, but also by all of one's background beliefs. (Simulating the belief that *p* does not involve pretending that one believes *only* that *p*, but that one believes that *p in addition to* all one's other beliefs.) For example, suppose that when Miss Marple runs a simulation on the premise *The butler was at the public house all evening*, she derives the simulated conclusion *The butler is morally depraved*. (Suppose she tacitly believes that all people who frequent public houses are morally depraved.) And suppose, too, that she wishes to reject this conclusion. (She also thinks the butler is a cat-lover and believes that no cat-lover is morally depraved.) Now if a premising policy just *is* a policy of accepting the results of simulation, then a refusal to accept one of these results will mark an *abandonment* of the policy – and thus, given our identification of premising with making up of mind – a change of mind. But it would be absurd to say that in refusing to accept that the butler was morally depraved Miss Marple thereby *changed her mind* about his being at the public house.

So if making up one's mind involves committing oneself to a policy of reasoning, then this policy cannot be one of accepting without exception all the results of simulation. Nor is it clear how simulation could be restricted in scope so as to exclude background beliefs irrelevant to the premising process. Simulation, then, on its own, would not offer a satisfactory means of evaluating one's premising commitments. It would have to be *supplemented* with rules or procedures for assessing simulated conclusions for normative warrant, relative to one's current premises. And, as we have seen, such rules and procedures will generally be linguistic.

A second problem for simulation arises from the fact that premisers may know they have instinctive inferential *defects*. For example, I may know that I have a deep-rooted tendency to make certain mistakes in reasoning with conditionals, or to be influenced by wishful thinking on certain topics. Now simulation-based premising would reflect these tendencies. (Simulation, remember, involves feeding a premise to your instinctive inferential processes.) But it would be odd to suppose that in committing myself to a premising policy, I commit myself to replicating my known inferential defects. If anything, one would think, I commit myself to trying to *rectify* them. But simulation will offer no guidance as to how to do this. It will tell me what I *would* infer if I believed that *p*, not what I *should* infer. Again, it seems, simulation will have to be supplemented by more specific inferential procedures or principles. And, again, these will generally be language-based. It appears, then, that premisers will inevitably have to draw at some point on specific inferential principles, either explicitly represented or embedded in practical inferential skills.

It may be objected that to suppose that we adhere to principles of inference is to suppose that some inferential transitions are non-negotiable (analytic) and thus that there are canonical procedures for revising and updating one's belief system in the light of conflicts with new evidence. (If it is a semantical rule that all bachelors are unmarried men, then the inference 'Fred is bachelor \rightarrow Fred is an unmarried man' is non-negotiable: were I to uncover evidence that Fred is not an unmarried man, then I would have to repudiate belief in his bachelorhood). But, as Quine famously argued, there are no non-negotiable inferential transitions. Any inference may be revised, provided one makes sufficiently drastic changes to other elements of one's belief system. And in updating our beliefs, what matters is not that we respect local semantical rules, but that we maintain the most stable global configuration in our belief system.

Now this may be true enough of low-level belief. But if the above argument is sound, it is not true of virtual belief. Forming a virtual belief involves adopting inferential strategies that are relatively *insulated* from one's background beliefs. We should not, I think, find this conclusion surprising, given what we know about virtual belief. Forming a virtual belief involves deliberately schooling one's thoughts. And it will be easier to do this by applying discrete local principles than by making assessments of global coherence. This conclusion also reflects what we know about the function of virtual belief formation. To make up one's mind about a matter is, in effect, to *foreclose* on deliberation about it – to give it, as it were, the status of cognitive *trumps*. And it would make sense to keep items with this status relatively insulated from mundane changes in one's background beliefs.

3.3 Some conclusions

Let me sum up the state of play. Cohen is wrong to suppose that premising *must* be linguistic. In theory, it could exploit explicit inference-rules that are not formulated in language. In practice, however, such rules will generally be linguistically formulated, coming to us via dictionaries and logic primers. Moreover, the simplest and most easily acquired forms of premising (ones employing practical inferential skills, rather than explicit theory) will be language-based. Simulation, too, although apparently not requiring language, would have to be supplemented either by explicit theory or by practical skills. Given this, I think it is safe to conclude that non-linguistic premising, if it occurs at all, will be a late and sophisticated development – probably involving *suppression* of a linguistic component – rather like the ability to read in silence.

We can now fit this conclusion into the wider argument of this chapter. So far, our defence of the cognitive conception of language has followed something of a self-denying ordinance, ignoring the introspective evidence for the role of language in thought, and concentrating instead on relatively a-prioristic considerations. Thus, I have suggested that we can actively form some of our mental states (that we can *accept* propositions, and *make up* and *change* our minds) and have argued that we do so by engaging in policies of premising. And I have outlined some reasons for thinking that premising will generally take a linguistic form. On this basis, one could predict that, whenever we find creatures who make up and change their minds, we will find them talking and reasoning with themselves, either silently or out loud. And this of course is just what we *do* find in our own case.

We talk to ourselves a lot, and it seems that when we do so we are not just idling, but engaging in genuine ratiocination. This intuitive view, however, cries out for theoretical underpinning – for an explanation of how it is possible for language to have a cognitive role. And this is just what our story about virtual belief has provided. Sometimes, when we speak to ourselves, what we are doing is silently articulating premises for adoption or manipulation, thus forming and processing virtual beliefs. (This may not, of course, be quite how we *think* of our activity at such moments; but, if I am right, it is the best way of characterizing it.) It is important to note that this view assumes that inner speech is entertained as *interpreted speech*. That is to say, our inner verbalizations present themselves to us as expressions of various *propositional contents* which are available for acceptance or rejection. So in claiming that we think in language, I am *not* claiming that we think by manipulating *uninterpreted* natural language symbols. (That view would be rather unattractive, given that such symbols will often be ambiguous in a way that thought is not.) Indeed, it might be better to say that we think *with* language, rather than *in* it; language provides a kind of *scaffolding* for our premising activities.

Of course, as I noted earlier, there are other, and more widely accepted, explanations of how language can have a cognitive role. In the final section of this chapter I shall briefly compare these accounts with the one outlined here.

3.4 *A comparison*

All defenders of the cognitive conception of language agree that we can think by tokening¹³ natural language sentences. And all agree that what makes a tokened sentence a judgement, rather than a fantasy, say, or idle speculation, is, in a broad sense, its causal role: beliefs have certain systematic effects on subsequent reasoning and behaviour that fantasies and idle speculations do not. The theories differ, however, in the accounts they give of the *determinants* of causal role.

According to what I shall call the *standard* version of the cognitive conception, the determining factors are subpersonal. What makes a sentence-token a judgement rather than a fantasy is the fact that it has a certain computational role – that it would causally interact in a certain characteristic way with various other token states, linguistic and non-linguistic. These relational facts, the story goes, supervene upon intrinsic properties of the token state and the local environment in which it is tokened. There are various ways of developing this view. According to one, there exists a specialized cognitive processor dedicated to the manipulation of natural language sentences. Advocates of such a view include Carruthers (1996) and, possibly, Bickerton (1995). A weaker version of the standard view holds that the causal role of sentence-tokens is determined by associative mechanisms: activating a representation of a natural language sentence tends to lead to the activation of other, semantically related, sentences. This view is advocated by Smolensky (1988), and, in a more guarded form, by Dennett (1991).

According to the version of the cognitive conception defended here, by contrast, the determinants of causal role are certain *intentional actions*. One can *decide* whether a sentence is going to function as a judgement by deciding whether or not to endorse it as a premise. If one does choose to endorse it, then the sentence will acquire the causal role of a belief in virtue of subsequent *personal* events – remembering that one is committed to taking it as a premise, working out that it mandates certain further acts of acceptance, determining to honour one's premising commitments by performing these acts, and so on. Call this view the *intentionalist* version of the cognitive conception.

I think this view has some distinct advantages over its rivals. First, it does not require there to be a dedicated hardwired cognitive processor for natural

¹³ Either by overtly uttering them, or, more often, by forming auditory or articulatory *images* of them (see Carruthers, 1996).

language sentences. The processing of virtual beliefs is done by a kind of *virtual* processor, which is realized in the subject's low-level psychological processes. This is a consequence of the fact that the formation and processing of virtual beliefs is under intentional control. People accept the consequences of their premises, not because they are wired up do so, but because they *believe* the consequences are warranted and *want* to honour their premising commitments. (Let me emphasize again that these beliefs and desires will not, as a rule, be *consciously* entertained.) Evolutionary considerations, I think, favour this story over its rival. For virtual processors will be much easier and cheaper to install than hardwired ones; existing cognitive systems would not have to be rewired or reduplicated, simply *reprogrammed*. Indeed, given the apparent recency of language on the evolutionary timescale, it is not clear that there would have been time for the necessary rewiring to occur.

The intentionalist story also has advantages over the weaker, associationist, version of the standard defence. For associationism seems to offer no scope for active, one-off, making up of mind. Associationist processes are not under intentional control. One cannot *decide* that a certain sentence will have the sort of cognitive associations that a belief ought to have: one can only *hope* that it will do so. The intentionalist story, by contrast, is much more flexible. One can decide that a certain sentence will function as a premise by deciding to *treat* it as one, employing it in appropriate inferential strategies.

Why, then, is the intentionalist account not widely accepted? One reason may be that it seems to imply that inner speech is under *personal control*. And this looks wrong: words can just pop into our heads, as unbidden as the thoughts they carry. This objection, however, misrepresents the intentionalist position. The intentionalist need not claim that the *production* of sentence tokens is always under personal control – only that the act of *endorsing* them is. Note, too, that in saying that the cognitive mode of inner speech is under personal control, the intentionalist is *not* committed to claiming that there is a central executive which decides how sentence-tokens are to function. Nor is the intentionalist committed to denying that personal actions may issue from low-level *pandemonium* processes of the sort Marvin Minsky has described (processes, that is to say, of evolutionary competition among unintelligent neural subsystems or 'demons'; see Minsky 1985). The distinction between centralized and pandemonium accounts of the genesis of intentional action is *orthogonal* to the distinction between subpersonal and intentionalist versions of the cognitive conception. Advocates of pandemonium models are not committed to denying that there *is* a personal/subpersonal distinction, at least of a rough-and-ready kind; they merely seek to reject one account of what the distinction consists in. And this is all that the intentionalist defender of the cognitive conception needs. The intentionalist's thesis is simply that there is a distinction, at least of

degree, between those events and processes that are under personal control and those that are not, and that the events which determine the cognitive role of natural language sentence-tokens often fall into the former class.

The bulk of this chapter, linking change of mind with premising and premising with language, has been devoted to outlining one line of argument for this conclusion. But I think that everyday experience also confirms it. It is wrong to suppose that inner speech always comes to us with its psychological mode (speech act, judgement, whimsy, whatever) pre-determined. Suppose I am listening to a speaker delivering a long and rather dense paper. As I struggle to make sense of his words, the sentence 'He's talking nonsense' suddenly pops into my head. Do these words constitute a serious thought or just a display of idle pique? Consider what may happen next. I may think about the words, decide, after brief reflection, that they cannot be justified, and put them down to irritation at the speaker's prolixity. Alternatively, the words may alert me to the fact that the paper really is very bad. In this case, I will endorse the sentence and start to reason with it as a premise. (I will begin to wonder why the speaker was invited, what I ought to say if he asks my opinion of his paper, and so on.)

Scenarios such as this suggest that it will not do to say that the cognitive role of sentence tokens is determined wholly by subpersonal processes occurring upstream of their tokening. Equally important are subsequent personal decisions to employ a sentence in one role or another. This is not to deny, of course, that some verbalizations spontaneously assume a certain role for us, without our needing to make any decision at all. Such cases need not cause a problem for the intentionalist. For omissions as well as actions can be intentional. The intentionalist can allow that premises may be endorsed *by default* – by failing to reject them – as well as by datable acts of endorsement.

The intentionalist may also, I think, be able to motivate a rather stronger claim. For in some cases it does seem possible to control the direction of one's inner speech, much as one can control the direction of a conversation. Think of a case in which one is deliberately trying to reason one's way to the solution of some problem. It is not implausible to suggest that in such cases our long-term argumentative goals can influence the direction and content of our silent soliloquizing. If so, then our personal control of inner speech may sometimes extend upstream as well as down.

Acknowledgements

I am grateful to Jill Boucher, Peter Carruthers, Jonathan Cohen, Gloria Origgi, Dan Sperber, Yorick Wilks and three anonymous referees for comments on earlier drafts of this paper, and to all the participants in the Hang Seng

conference on Language and Thought for some stimulating discussions. Thanks, too, to Gavin Boyce, Tom Dickins, Ewan McEachran and Jenny Saul.

References

- Baier, A. (1979). Mind and change of mind. *Midwest Studies in Philosophy* 4, 157-176.
- Bennett, J. (1990). Why is belief involuntary? *Analysis*, 50, 87-107.
- Bickerton, D. (1995). *Language and Human Behaviour*. London: UCL Press.
- Botterill, G. (1994). Beliefs, functionally discrete states, and connectionist networks: a comment on Ramsey, Stich and Garon. *British Journal for the Philosophy of Science*, 45, 899-906.
- Bratman, M. (1992). Practical reasoning and acceptance in a context. *Mind*, 101, 1-15.
- Carruthers, P. (1996). *Language, Thought and Consciousness*. Cambridge: Cambridge University Press.
- Clark, A. (1990). Connectionist minds. *Proceedings of the Aristotelian Society*, 90, 83-102.
- Clark, A. (1993). *Associative Engines*. Cambridge, Mass.: Bradford Books.
- Cohen L.J. (1992). *An Essay on Belief and Acceptance*. Oxford: Oxford University Press.
- Dennett, D.C. (1978). *Brainstorms: Philosophical Essays on Mind and Psychology*. Montgomery, VT: Bradford Books (English edition, Sussex: Harvester).
- Dennett, D.C. (1991). *Consciousness Explained*. Boston: Little. Brown.
- de Sousa, R. (1971). How to give a piece of your mind: or, the logic of belief and assent. *Review of Metaphysics*, 25 (1), 52-79.
- Frankish, K. (1996). Dennett's opinions. Unpublished typescript: University of Sheffield.
- Goldman, A. (1989). Interpretation psychologized. *Mind and Language*, 4, 161-85.
- Goldman, A. (1992). In defence of the simulation theory. *Mind and Language*, 7, 104-19.
- Gordon, R.M. (1986) Folk psychology as simulation. *Mind and Language*, 1 (2), 158-171.
- Harman, G. (1973). *Thought*. Princeton, New Jersey: Princeton University Press.
- Harris, P. (1989). *Children and Emotion: the development of psychological understanding*. Oxford: Blackwell.
- (1992). From simulation to folk psychology: the case for development. *Mind and Language*, 7, 120-44.
- Heal, J. (1986). Replication and functionalism. In J.Butterfield (ed.), *Language, Mind and Logic*, pp. 135-150. Cambridge: Cambridge University Press.
- (1994). How to think about thinking. In M.Davies and T.Stone (eds.) *Mental Simulation: philosophical and psychological studies*, pp. 33-52. Oxford: Blackwell.
- Kaplan, M. (1981). Rational acceptance. *Philosophical Studies*, 40, 129-45.

- Levi, I. (1967). *Gambling with Truth*. London: Routledge & Kegan Paul.
- Minsky, M. (1985). *The Society of Mind*. New York: Simon and Schuster.
- O'Brien, G. (1991). Is connectionism commonsense? *Philosophical Psychology*, 4, (2), 165-178.
- Peacocke, C. (1992). *A Study of Concepts*. Cambridge, Mass.: MIT Press.
- Pojman, L.P. (1985). Believing and willing. *Canadian Journal of Philosophy*, 15 (1), 37-55.
- Ramsey, W., Stich, S. Garon, J. (1991). Connectionism, eliminativism and the future of folk psychology'. In *The Future of Folk Psychology*, ed. J.D.Greenwood, pp. 93-119. Cambridge: Cambridge University Press.
- Smolensky, P. (1988). On the proper treatment of connectionism. *Behavioural and Brain Sciences*, 11, 1-74.
- Sperber, D. (1997). 'Intuitive and reflective beliefs. *Mind and Language*, 12 (1), 67-83.
- Williams, B. (1973). Deciding to believe. In his *Problems of the Self*, pp. 136-51. Cambridge: Cambridge University Press.
- Winters, B. (1979). Believing at will. *Journal of Philosophy*, 76, 243-256.