Psychological Inference, Constitutive Rationality and Logical Closure

Ian Pratt

Published in Hanson, P. (ed.): *Information, Language and Cognition*, Vancouver Studies in Cognitive Science, vol.1, University of British Columbia Press, 1990, pp. 366-389

1 Synopsis

In this paper I wish to examine an argument that has come to command wide acceptance within the contemporary philosophical community. So plausible has this argument seemed that it has sometimes served as a foundation for quite amazingly counterintuitive claims. I hope to show that the argument is flawed, and thus save us from the counterintuitive claims.

Here is the argument. Whenever we describe someone as being in a certain state of mind—as doubting, perceiving, affirming, denying, willing, not willing various things—our description has this curious property: it implies, in the first instance, not how that state of mind *will* change, but, rather, how it *ought to* change. We can say of a student pondering a mathematics problem what solution he *ought* to arrive at, because we can say whether the proposition he is trying to establish really does follow from the axioms he is assuming. Or we can say of a detective engaged on a murder hunt which of his suspects he *should* regard as the most likely culprit, because we can say whether the detective's evidence confirms—according to our chosen inductive method—the proposition that this or that individual is guilty. But whether we think the student or the detective will *actually* arrive at the proper conclusions is another and less certain matter, one which depends on our crediting them with the propensity to reason as they ought. More generally, we can say of a person characterised in psychological terms how he ought to reason, by tracing the logical (or confirmation-theoretic) interrelations between the contents of the propositional attitudes—beliefs, desires, intentions and so forth—attributed to him. Thus, propositional attitude attributions determine, in the first instance, not how their subjects will reason, but only how their subjects will reason if they reason aright.

It follows—so goes the argument—that if psychological descriptions are to have any explanatory or predictive power, they must be accompanied by an assumption of general rationality—an assumption which implies that the subjects being described will, probably at least, reason as they ought. For, without such an assumption, there will simply be no observable predictions to be made. Add to this conclusion the view that the content of psychological descriptions is exhausted by the observable predictions which they facilitate, and you arrive at the following result: to ascribe psychological states is to ascribe rationality. In other words, it is literally nonsense to say of people (or any other creatures or entities) that they have beliefs, desires, intentions and the rest of it, unless you are also prepared to credit them with a propensity to revise those psychological attitudes in a generally rational way.

This basic argument, plus or minus various details, underlies many current philosophical accounts of the psychological attitudes: since psychological descriptions only determine, in the first instance, what their subjects will do under some assumption of optimal or rational functioning, some such assumption must be part and parcel of our willingness to engage in that kind of description at all.[1] Not that the argument lacks detractors. Cognitive psychologists have accumulated a large body of experimental evidence purporting to show that errors in human reasoning, far from being occasional blots on an basically commendable copy-book, are, on the contrary, many and systematic.[2] And these psychologists have been joined by a number of philosophers who are less than sanguine about an a priori argument for human rationality.[3]

The aim of this paper is to oppose the argument for general rationality given above, not, as many others have done, by citing evidence to the contrary, but rather, by challenging the premise on which it rests—that psychological descriptions without an accompanying assumption of general rationality are devoid of

[1] See, for example, Dennett, Daniel C: "Intentional Systems", *Journal of Philosophy*, LXVIII, 4 (1971), reprinted in *Brainstorms*, Montgomery, Vermont: Bradford Books (1980); Davidson, Donald: "Mental Events" especially pp.221–223, in *Essays on Actions and Events*, Oxford: Clarendon Press (1980); Lewis, D.K: "Radical Interpretation" in *Philosophical Papers* vol.1, New York: Oxford University Press (1983), and *On the Plurality of Worlds*, Oxford: Basil Blackwell (1986), section 1.4; Stalnaker, Robert C: *Inquiry*, Cambridge, Massachusetts: MIT Press (1984).

[2] For example, Johnson-Laird, P.N. and P.C. Wason: "A Theoretical Analysis of Insight into a Reasoning Task" in their collection *Thinking: Readings in Cognitive Science*, Cambridge: Cambridge University Press (1977); also, Tversky, Amos and Daniel Kahneman: "Judgment under Uncertainty: Heuristics and Biases", *Science* 185 (1974) pp.1124–31, reprinted in the same Johnson-Laird and Wason collection.

[3] For example, Stich, Stephen P: "Could Man be an Irrational Animal? Some Notes on the Epistemology of Rationality", *Synthese* (1984). See also: Thagard, Paul and Richard Nisbett: "Rationality and Charity", *Philosophy of Science* 50, 1983, where only a very limited form of constitutive rationality is endorsed.

explanatory or predictive power. The challenge takes the form of drawing attention to a rival view of how psychological descriptions figure in reasoning, a view which can be found in Hobbes, and which has since surfaced periodically in the philosophical, psychological and, more recently, artificial intelligence literatures. Indeed, it is of the essence of this theory of psychological inference that such inference exploits the weakness of human rationality rather than its strength. In the next section, I describe an aspect of human cognition which I believe underlies psychological inference, and which I call *virtual reasoning.* The account of psychological inference proper occurs in the following section, and in the remainder of the paper, I compare the approach developed here with that taken by the defenders of the argument for constitutive rationality.

One point of terminology. I call the concepts of belief, desire, hope, fear, intention, etc. *psychological concepts.* The psychological concepts are just the ones which we use distinctively to reason about other people's (or our own) minds. Reasoning involving such concepts I call *psychological reasoning* or *psychological inference.*

2 Virtual Reasoning

Let us begin with an experiment. Imagine that, upon unlocking the door of your office one morning, you are confronted by a scene of chilling gore: the bloodied corpse of a colleague lying face up on your desk, a dagger driven into his chest up to the hilt.

What is your reaction? Surprise. Horror. Evidently, the man has been murdered. The police must be summoned; nothing must be touched. Who could have done it? Only someone with a key to the office. The secretary? The department chairman? Speaking of which, the door was locked just now, which means that whoever did it locked the door after him. And why do that?...

Thus one's mind runs on. Or approximately thus. What people think of when asked to imagine a dead body greeting them as they arrive at work varies from person to person and from time to time; but something along the lines sketched in the previous paragraph is presumably what goes through most people's heads when asked to contemplate such a scenario. Imagine finding a body on your desk, and you find yourself imagining your reactions—inferring that the man was murdered, resolving to call the police and not to disturb any evidence, wondering who did it (and making a few immediate deductions), puzzling over the locked door. Of course, you do not really infer anything or resolve to do anything. It is simply that, when you imagine a situation such as the above, you complete and extend the situation by imagining that you infer and resolve to do certain things.

I say "imagine inferring", "imagine resolving", but that's not exactly right. The sort of thing I have in mind is not at all like, for example, imagining taking a drug which causes one to infer that taxis are made of newspaper and to wonder why the sky is the colour of marmalade. The kind of imagination I tried to elicit in asking you to imagine finding a corpse in your office is quite different: it is not a matter of imagining engaging in certain pieces of reasoning, but rather, of *reasoning within the scope of an imagined cognitive predicament.* As such, the reasoning purports to be, in some sense, rational. It would be rational, if one *really did* see a corpse, to wonder how it got there, to resolve to call the police, and—given one's other beliefs, perhaps—to make certain specific deductions as to who was responsible. Likewise, when one *imagines* the same thing, one's mind runs on in a correspondingly rational way. Hume tells us (though with more than a hint of irony) that nothing is more free than the imagination of man. Possibly so. But the cases I want to draw attention to here are those where imagination is a reasoned, rather than a rhapsodic affair.

We need a name for this kind of thinking. I propose to call it *virtual reasoning.* (I should have preferred the phrase *imaginative reasoning*, but it already has another well-established English meaning.) Virtual reasoning is reasoning within the scope of an imagined cognitive predicament. Like ordinary reasoning, it includes adopting new beliefs, rejecting old ones, forming intentions, asking questions, and so on; except that all these changes are confined within the scope of an imaginary scenario. In order to distinguish these activities from their counterparts in thinking for real, I shall dub the latter, *agentive reasoning.* Agentive reasoning standardly leads to action in a regular and straightforward way. If you really see a corpse, and thereupon form the intention of calling the police, then you will reach for the telephone. If, in addition, you suspect the secretary, you may inform the authorities of that suspicion. Not so with virtual reasoning, needless to say. Another difference: agentive thinking has regular and straightforward lasting effects on one's psychological state, particularly in respect of adopting or retracting beliefs. Virtual reasoning, by contrast, is a relatively transitory affair: all 'inferences' and 'changes of mind' are confined to the reverie, and vanish when it vanishes, leaving only a memory of the thinking's having occurred. If you really discover the body of a colleague, you permanently modify your beliefs about him, permanently cancel your plans to meet with him, and so on; but when you *imagine* discovering the same thing, your long-term beliefs and intentions will not be affected in the same way. It is as if, in virtual reasoning, one makes a temporary copy of one's beliefs and intentions, to which changes can be made without affecting the original.

Of course, the difference between what I have been calling agentive and virtual reasoning cannot be enshrined in any sentence-long definition. And it would be a simplification to say, for example, that agentive reasoning both leads to action and affects one's long-term beliefs, whereas virtual reasoning does neither of these things. For one thing, imagining seeing a corpse in one's office might lead to the formation of contingency plans capable of affecting one's future actions. For another, merely imagining seeing a corpse might colour one's future reasoning in case one really does see a corpse—for example, by making one more receptive to suggestions that conform to what one earlier imagined. Nevertheless, the way in which reasoning about imagined situations guides our actions and changes our beliefs has an subtlety and indirectness about it which sets it apart from the more regular and straightforward way in which ordinary (what I called agentive) reasoning operates. The difference between agentive and virtual reasoning is one which I can only hope to capture by appeal to the reader's experience: the rough-and-ready characterisations given above are intended as an aid to such an appeal.

Virtual and agentive thinking are not, however, totally dissimilar. On the contrary, there are phenomenological correspondences between the two. The surprise and horror one imagines upon imagining seeing a corpse are *somehow* like real surprise and horror—except that they are less intense and immediate. The puzzlement arising from the locked door is *somehow* like the puzzlement that arises from a real conundrum—except that it is less vexing and urgent. The make-believe inference that the secretary might have been responsible is *somehow* like a real inference to the same effect—except that it has less force and vivacity. As Hume put it:

> I say, then, that belief is nothing but a more vivid, lively, forcible, firm, steady conception of an object, than what the imagination alone is ever able to attain. This variety of terms, which may seem so unphilosophical, is intended only to express that act of the mind, which renders realities, or what is taken for such, more present to us than fictions, causes them to weigh more in the thought, and gives them a superior influence on the passions and imagination.[4]

The themes of virtual reasoning can be as mundane as they can be sensational. For example, one can imagine replying to a certain chess gambit with which one is experimenting, observing how things look from the opponent's viewpoint as the gambit unfolds. Or one can imagine being faced with difficult questions at a meeting, or an unruly student in class, and do on. Virtual reasoning, or reasoning within the scope of a make-believe situation, is a faculty which normal people employ naturally and frequently.

I wish to put forward two hypotheses concerning virtual reasoning. Both are empirically testable; both are (I claim) intuitively plausible; in neither case, however, can I support them with hard—that is to say, with non-introspective—evidence. They are:

- Accuracy:

  When one engages in virtual reasoning within the scope of some imagined cognitive predicament, one reasons very much as one *would*, or at least, *might well* reason if one really *were* in that cognitive predicament.

[4] Hume, David: *Enquiry Concerning Human Understanding*, 3rd ed., Oxford: Clarendon Press (1975), p. 49.

- Efficiency:

  > Many of the cognitive processes responsible for virtual reasoning are also those responsible for agentive reasoning.

According to the accuracy hypothesis, the sorts of (virtual) thoughts that run through your mind when you imagine finding a corpse in your office are similar to the (agentive) thoughts which would (or at least might well) run through your mind if you really did see a corpse in your office. Thus, should such a situation befall you, you probably would be surprised, horrified, curious as to who did it, puzzled about the locked door; you probably would resolve to call the police, not to disturb any evidence, and so on. Similarly, *mutatis mutandis*, for the correspondence between, on the one hand, imagining replying to a chess gambit, or imagining answering a difficult question in a meeting, and, on the other, what you would have been thinking had you really been countering the gambit or fielding the question. In short, the accuracy hypothesis states that virtual reasoning can be a source—though not an infallible one, to be sure—of counterfactual knowledge about one's own agentive reasoning.

The efficiency hypothesis is harder to justify. But it would at least be unsurprising if many of the same cognitive processes that are engaged when one imagines seeing a corpse in one's office—those processes, that is, whereby one makes the various virtual inferences and resolutions as outlined above—are those very cognitive processes which would be responsible for the corresponding real (i.e. agentive) inferences and resolutions if one really were to see a corpse in one's office. Unsurprising also, if the cognitive faculties that are engaged in countering imaginary chess positions are those very cognitive processes that are engaged in countering real chess positions, i.e., the cognitive processes underlying one's ability to play chess. The efficiency hypothesis is thus an explanation for the accuracy hypothesis: virtual thinking is a reliable source of knowledge as to how one would think in the imagined situations, simply because the same mechanisms are involved in reasoning within the scope of an imagined situation as would have been involved if that situation had been for real. The efficiency hypothesis would also explain the phenomenological correspondences between virtual and agentive thinking, as chronicled by Hume.

But I have no desire to press these points: I am willing to let the efficiency hypothesis remain a hypothesis open to empirical confirmation or disconfirmation. The hypothesis will, however, figure crucially in the account of psychological reasoning given in the next section.

The following observations concerning virtual reasoning will stand us in good stead for the remainder of this paper. First, virtual reasoning, like agentive reasoning, has to do with more than deciding whether given propositions are true or false. For example, when one imagines seeing a corpse, any ensuing bafflement (virtual bafflement) about the locked door can only be as a result of having spontaneously inferred (virtual inference) first that there was a murder, and second that the door had been locked, probably by the murderer. One does not need to be asked these things; nor need one imagine being asked: the inferences can come unprompted. Likewise, the processes whereby one resolves (virtual resolution) to telephone the police and not to disturb any evidence are *reasoning processes*, in the wide sense in which I use the term, but they are not processes for deciding whether a given proposition is true or false. There is much more to reasoning—virtual and agentive—than trying to settle given questions: unprompted inference, the raising of questions, the resolving of contradictions, the formation and revision of intentions, are also important.

Secondly, as is well-illustrated by the case of the imagined chess gambit, one has a sense, in virtual reasoning (as in agentive reasoning), of which inferences are easy and obvious, and which are difficult and obscure. The pawn threat to the queen may be such that any fool could have seen it; but seeing the long-term strategic threat in the center may take some inspired thinking. All this one experiences for oneself in the process of imagining facing a given board position. Thereby, one can know which inferences, had the imagined situation been real, one would have made easily, and which, by contrast, would have been difficult. Thereby also, one can draw conclusions as to which inferences one would have made for sure (the easy ones), and which might well have eluded one (the hard ones).

Thirdly, as the chess example also illustrates, one may, in imagining replying to the gambit, imagine doing so without knowing beforehand how it is intended to unfold. Of course, since it is a gambit of one's own devising (hence the need to test it in this way), one is aware from the start what it involves, what its main threats are designed to be, what defensive value it has, and so on. Nevertheless, one can, when imagining replying to it, suspend that knowledge temporarily—exclude it from one's imagined cognitive predicament. Thus, in imagining replying to the gambit, one confronts the problem of figuring out (virtual figuring-out) the intentions of one's imaginary opponent as one 'sees' his position develop. And this exercise of imagining that one does not know the intended threats in advance can, I claim, be a source of information as to how one would have—or, at least, might well have—responded had one really been ignorant in the way imagined. For example, one might conclude that one would have been unable, in such circumstances, to divine the plan behind the gambit before it was too late. (More on the utility of such knowledge anon.)

Presumably, the more radical the ignorance one imagines labouring under, the less reliable, as a source of counterfactual knowledge, will be any virtual inference that ensues. Imagine that you barely know the rules of chess, have not a smattering of elemental strategy, but are nevertheless contemplating a given board position. What you imagine, if anything, is unlikely to be a reliable guide to how you would have reasoned in such circumstances. Nevertheless, virtual reasoning can proceed in some situations where one imagines oneself lacking beliefs one in fact has (as well as having beliefs one in fact lacks), provided that the adjustments one's imagination is called on to make are sufficiently small.

Fourthly, one thing that virtual inference is not. It is sometimes important to engage in contingency planning. By considering what one should do should one see a corpse in one's office or discover a fire in a labyrinthine building, one can rehearse thinking which one might not have the time to produce in a real emergency—the most appropriate course of action, the nearest escape. The results of such rehearsed thinking can then be distilled and preserved for later use as and when required. In the same way, devising replies to standard chess openings or answers to anticipated questions can be of use should one ever encounter then in the future. Thus, contingency planning increases the efficiency of thinking by spreading the cognitive load more evenly over time.

None of this is quite what I have in mind when I speak of virtual reasoning. Consider again the corpse-discovery example. When I imagine being confronted with the body, I can perform the reasoning in which I might well engage if I were in that situation—and this is reasoning *on the spot*, reasoning which must be performed quickly in (make-believe) distressing and possibly dangerous circumstances. Performing such virtual reasoning is not a matter of reflecting at leisure on what would be the best thing to do in such circumstances. Of course, when reasoning in the imagined situation, one is *trying* to decide on the best course of action; but to try in difficult circumstances under pressure of time—even if those circumstances and that pressure are imaginary—is not necessarily to succeed. Indeed, it may be as a result of doing badly in such a piece of virtual reasoning that one decides to devise an optimum plan. Similarly for the case of imagining discovering a fire: if you want to work out the best course of action in such circumstances, I do not recommend simply imagining being in them and observing your reactions. *Supposing* that a situation obtains is one thing; *imagining* being in it, quite another.

The point emerges even more clearly in the chess gambit example. When I imagine encountering the gambit as one who has never seen it before, I can consider what I might well do if deprived of the knowledge of how the gambit is to develop. It is quite possible that I should fail to notice (virtual noticing) the threat which it is the whole point of the gambit to pose. It is quite possible, then, that I should imagine responding with an inappropriate move. Clearly, this is not contingency planning. Contingency planning makes sense only if one uses all the information at one's disposal, and only when one is interested in inappropriate actions or erroneous conclusions so as to avoid them. Thus, virtual reasoning incorporates limitations, especially of time and knowledge, which contingency planning should have every interest in assuaging. Of course, virtual reasoning *is* reasoning, and, as such, involves the same attempts at accuracy and correctness that infuses all reasoning; it is just that virtual reasoning can be carried out in imagined situations which recognise features likely to be detrimental to these goals.

3 Psychological Inference

Preliminaries having been dealt with, I now address the topic of psychological inference. The central insight I want to develop is this. If virtual reasoning can be a reliable guide to how one would have thought had one been in a certain cognitive predicament, then it can be a reliable guide to how *another person* would have thought had *he* been in such a cognitive predicament. In which case, it can be a reliable guide to how another person *will* think if he *is* in such a cognitive predicament.

Let me illustrate with an example. Suppose that, while playing chess, you see an opportunity to capture the white queen. The plan involves advancing pawns so as to wall in the queen, enabling a piece of lower value such as a knight to move in for the kill. As things are, however, it will take several moves to close the trap, during which time white could still move to safety. Accordingly, you decide to move the pawns furthest from the white queen first, so that your intentions will be less obvious; in addition, you leave some other pieces vulnerable to a development which your opponent has set in train, and for which he must keep his queen where it is. This kind of thinking should be familiar to any chess player: in chess, it is no good posing threats to which obvious responses exist, but only those which are difficult to see until it is too late. Thus, one way to get the better of an opponent is to be able to reason about which things will be easy for him to see, and which difficult. Of course, this is not reasoning about chess *per se;* it is reasoning about your opponent's likely thoughts. That is, it is psychological reasoning.

I claim the following. Sometimes, we can tell whether a threat is obvious or obscure for our opponent because we imagine facing the situation he is facing and, when we do so, experience either ease or difficulty in seeing (virtual seeing) that the threat exists. Thus, you can tell that your plan to capture your opponent's queen might prove difficult for your opponent to see, because you can imagine yourself in his situation—i.e. faced with a knowledge of the dispositions of the pieces on the board but not of the black player's intentions—and, reasoning within the scope of this imagined situation, you yourself find the threat unobvious. My suggestion, then, is that we can tell the obvious from the obscure by employing our faculty of virtual thinking.

Again, consider the position of a murderer contemplating fleeing the country for fear of being caught by the police. He may know what evidence the investigating detective has, and he may know that that evidence implicates him. But that does not answer his question: the important thing is whether the detective is likely to use the clues at his disposal to full effect. And the murderer may, on considering the detective's cognitive predicament (by imagining being in it), decide that the inference in question would never occur to him in such circumstances. On this basis, he may conclude that a suitable chain of reasoning is unlikely ever to occur to the detective either.

The same kind of reasoning occurs in everyday situations. Why might I worry that my friend may doubt that I like him? Because I have, unavoidably, missed several recent appointments with him, uttered comments which, on reflection, might have been misinterpreted, and so forth. I know that the comments might have been misinterpreted because I can imagine hearing them (and can imagine having the particular concerns and interests with which I know my friend to be currently preoccupied), whereupon the offence-giving interpretations, which had not occurred to me before, now become obvious. I know that these interpretations, together with the missed appointments and so forth, might cause my friend to doubt my loyalty because, when I imagine having experienced the things I believe him to have experienced, that thought occurs (virtual occurrence) to me. Notice that the fact that I can devise a clever argument to demonstrate my loyalty, or the fact that I am assured of it anyway, are nothing to the point. It is enough that, in contemplating the situation from my friend's point of view, I can well imagine coming to that conclusion.

These examples illustrate a faculty that I believe is universal among normal persons. One can know how other people will—or at least may—think, and hence how they will or may behave, because one can imagine oneself in their cognitive predicament and do the relevant reasoning for oneself. In other words, one can use one's own cognitive processes to simulate those of others. This, I suggested above, is how one can know what is or is not an obvious threat in chess, whether the police are likely to be onto one, whether one's friend doubts one's fidelity. I also think that this is how one sometimes decides whether a lecture will be comprehensible to one's students, a journal article to its intended audience, and a remark in a conversation to one's interlocutor. To quote Hobbes:

> But there is another saying not of late understood, by which they might learn truly to read one another, if they would take the pains; and that is, *Nosce teipsum*, *Read thy self*... [which is meant] to teach us, that for the similitude of the thoughts, and Passions of one man, to the thoughts, and Passions of another, whosoever looketh into himself, and considereth what he doth, when he does *think*, *opine*, *reason*, *hope*, *feare* &c, and upon what grounds; he shall thereby read and know, what are the thoughts, and Passions of all other men, upon the like occasions.[5]

[5] Thomas Hobbes: *Leviathan*, (ed. C.B. Macpherson) Harmondsworth, Middlesex: Penguin Books (1968), p.82. For a discussion of Hobbes' idea and its relation to consciousness, see Humphrey, Nicholas: *Consciousness Regained*, Oxford: Oxford University Press (1983), chapters 1 and 3; see also Lyons, William: *The Disappearance of Introspection*, Cambridge, Massachusetts: MIT Press (1986). The most influential modern champion of reasoning-as-simulation is Kenneth Craik, whose views on this subject are expounded in his book, *The Nature of Explanation*, Cambridge: Cambridge University Press, 1943; see especially, pp. 50–61. Following in Craik's footsteps, Kahneman and Tversky develop a similar idea in their essay, "The Simulation Heuristic", in Kahneman, Daniel, Paul Slovic and Amos Tversky: *Judgment*

Predicting the course of another's reasoning by simulating it with your own will succeed only if certain uniformities obtain between your cognitive processes and his. If you want to predict the likely course of your opponent's thinking in a game of chess, and thence his likely moves, it is no good your deciding whether *you* would have seen the threat implicit in a certain chess position, if that is unlikely to bear any relation to whether *your opponent* would have seen it. Indeed, anyone who has played a pocket computer at chess may have noticed how, especially in the end-game, it can be defeated by ploys which we know would be obvious to (and therefore useless against) good human players. And, as for chess, so for anything else: if you try to predict someone's likely thinking by means of virtual reasoning, you will succeed only insofar as his cognitive processes function as yours do.

Not that you have to be soul-mates, you and he. For virtual reasoning is a powerful and flexible cognitive tool. As we have seen, we can, in imagining ourselves in another's cognitive predicament, imagine believing things which we in fact doubt, imagine also lacking knowledge to which we are in fact privy. We can imagine that we, like he whose thinking we are out to simulate, have a distaste for long, involved calculations (or a passion for them), or that we, like he, are of an impetuous disposition (or of a cautious one), and so on. If we think our chess

---

Under Uncertainty: Heuristics and Biases, Cambridge, Cambridge University Press, 1982. Another writer who touches on the idea of modelling the thoughts of others by means of one's own thoughts is Quine:

> ...[I]n indirect quotation we project ourselves into what, from his remarks and other indications, we imagine the speaker's state of mind to have been... Correspondingly for the other propositional attitudes, for all of them can be thought of as involving something like quotation of one's own imagined verbal response to an imagined situation.

(Quine, W.V.O: *Word and Object*, Cambridge, Massachusetts: MIT Press (1960), p.219). And this theme is developed (though in a different direction from that taken in this paper) by Stephen Stich in "Dennett on Intentional Systems", *Philosophical Topics*, 12, 1 (1981), and in *From Folk Psychology to Cognitive Science*, Cambridge, Massachusetts: MIT Press (1983), chapter 5. See also David Lewis' "Prisoners' Dilemma is a Newcomb Problem", in *Philosophical Papers* vol.2, New York: Oxford University Press (1986) pp.299–304.

Well-developed views on psychological inference by means of psychological simulation are to be found in the artificial intelligence literature: see Creary, Lewis G: "Propositional Attitudes: Fregean Representation and Simulative Reasoning", *Proceedings, Sixth International Joint Conference on Artificial Intelligence*, Tokyo (1979), and Haas, Andrew R: "A Syntactic Theory of Belief and Action", *Artificial Intelligence*, vol 28, 3 (1986). More generally, the study of simulation as a technique in artificial intelligence has recently been gathering momentum: the journal *Simulation* now has a regular AI feature, often devoted to some aspect of reasoning-as-simulation.

Some people have commented that there are points of contact between the view taken in this paper and the writings of the *Verstehen* theorists, or—less kindly—that I have "rediscovered hermeneutics". In fact, I think the connection superficial. However, I prefer to leave the question to the reader's judgement, rather than to engage in an amateurish exposition of material I cannot understand.

opponent stupid and dull, we discount, in imagining ourselves in his situation, our more inspired assessments of the state of play, and take into account only routine and immediate thoughts. If we think the detective on our tail is a Clusot, rather than a Poirot, we can allow for that too. Of course, the more we have to stretch our imagination to accommodate the differences between ourselves and our subject, the less reliable—or so I assume—will be the results of the process of psychological simulation I have just described. But total certainty is, needless to say, not what we need be after. A generally correct idea of the likely possibilities will do.

The view I am putting forward, then, is that psychological inference is sometimes psychological simulation. In order to know how someone will—or may well—reason, we imagine ourselves in his position and simulate that thinking by means of our faculty of virtual reasoning.

Psychological simulation can be an efficient method for predicting the behaviour of others. This is because, in simulating another's thought, we need no *theory* of how human reasoning—either ours or anyone else's—is done; we simply need to be able to do it. And if the efficiency hypothesis of section 2 is true, then many of the same cognitive mechanisms that we use in simulating another's thought are the cognitive mechanisms used in ordinary (non-psychological) agentive thinking. For instance, when I predict my opponent's likely responses to a chess gambit, I use many of those cognitive mechanisms which enable me to reason non-psychologically about chess. Whatever cognitive mechanisms that allow me to scan a board for threats and opportunities, compute the values of exchanges of pieces and so forth, now find alternative employment: they enable me to make inferences about how my opponent will or may think in a given situation. And this is so because my ability to engage in virtual thinking means that I can use those cognitive mechanisms to simulate the likely workings of their counterparts in my opponent's head. And, in this way, I can experience his likely difficulties first hand.

This is the crux of the matter: when one predicts that another person might well fail to see a threat in chess, or fail to make an inference, or otherwise engage in some seductive but erroneous piece of thinking, one need not have any sort of *theory* of how that person thinks. Such knowledge is unnecessary because psychological inference can simply trade on the fact that one person's cognitive mechanisms are really very much like another's. An inference that is obvious to one person is likely to be obvious to all; an inference that will require one person to search through and evaluate a large set of complicated possibilities will impose the same requirement on all. Not that people all think exactly alike: it is just that people think in sufficiently similar ways to allow one to guess at another's reasoning by imagining himself in that other's cognitive predicament, and by doing the reasoning for himself.

I say again: none of the above implies that psychological inference requires us to *know how* anyone's cognitive mechanisms work; i.e. none of the above implies that psychological inference requires us to possess a theory of reasoning. On the contrary, it is of the essence of the view of psychological inference as psychological simulation that such knowledge is unnecessary. Whilst a theory of reasoning is not required, however, neither is one precluded. It is possible that sometimes, psychological simulation is overridden by explicit theorizing; possible also that our psychological simulations are somehow 'infiltrated' by—and thus affected by—our explicit beliefs about the way people reason. Exactly when and how such overriding or infiltration takes place is a good, and partly empirical, question. But it is one I cannot delve into here.[6]

Let us review the chess example to get a fuller picture of what is going on. From your knowledge of your opponent—that he has an unobstructed view of the board, that he has been playing as if to win, that he has made no illegal moves—you are able partially to characterise his state of mind—that he *knows* the locations of the pieces, *desires* to win, *is acquainted with* the rules of chess. Perhaps you have also inferred from his previous moves that he is a good player, and that he is following a certain plan. The point of this psychological characterisation of your opponent, i.e. of the ascription of beliefs, desires and intentions to him, is that it enables you to know, in the relevant respects, what it is like to be he. It enables you to put yourself in his cognitive predicament, whence you can predict, by means of virtual reasoning, the likely evolution of that predicament. Having thus reached a conclusion about the various things that he may or may not think, you can then predict his likely behaviour. For instance, if you think he is likely to intend to move his queen so as to threaten your bishop, then you can conclude that he may well execute such a move.

The following diagram illustrates the process:

[6]Paul Thagard has pointed out to me that examples of this phenomenon are provided by cases where our psychological predictions appear to be affected by false beliefs about the way people reason. An example would be the Milgram experiments on obedience to authority. What do you think most people would do if asked to administer lethal electric shocks to helpless victims? Reflecting momentarily on what you would do, you conclude (by virtual reasoning, maybe), that you, and hence presumably they, would at the very least vigorously decline. But that is not the way many people react (they throw the switch after only minimal persuasion), and it is very possibly not the way you would react. If so, that suggests (but does not imply) that your inference to the contrary is not the result of pure virtual reasoning as described above, but is conditioned by a widespread and fallacious belief about people's moral sensibilities.

$$
\begin{tikzcd}
\text{Initial Physical description} \atop \text{+ General Knowledge} \arrow[r] & \text{Initial Psychological description} \arrow[d] \\
\text{Inferred Physical description} & \text{Inferred Psychological description} \arrow[l]
\end{tikzcd}
$$

From behavioural description plus general knowledge, infer psychological characterisation; from psychological characterisation employ virtual thinking to predict likely future psychological characterisations; from likely future psychological characterisations, infer likely behaviour.

This raises the question of how we obtain the psychological characterisation of our subject in the first place, and of how, once having obtained it and having decided (by means of virtual inference) on some of the likely ways in which that psychological state might develop, we draw the appropriate conclusions about the subject's likely behaviour. I suggest that we can get by with fairly crude pieces of common knowledge. For example, it is common knowledge that, if someone obeys the rules of a board game, he almost certainly *believes* that the rules are thus-and-so; common knowledge that persons in our culture who play board games desire seriously—though not maniacally—to win; common knowledge that people staring at chess-piece-sized objects twenty-five centimetres away tend to know where those objects are (under suitable descriptions); and so on. And of course, if somebody says something, that can also be an indication that they believe it. Common knowledge about what beliefs people have is, on the present view, no different in principle from common knowledge about anything else.

One can also ascribe beliefs and other psychological attitudes to a person on the basis that such attitudes would explain his behaviour. Thus, if I know that the investigating detective has been asking lots of questions about me, and has sent his sergeant to tail me, I can infer that he probably suspects me. For when he asked the questions about me and had me tailed, he presumably did so *intentionally* (it being common knowledge that these are among the sorts of things that are done intentionally). Having attributed these intentions to the detective, I can then use my faculty of virtual thinking to decide which possible beliefs (compatible with what I already know about the detective) might, if *I* had had these beliefs, have led *me* to form such intentions. The same beliefs, I can then conclude, might be responsible for the detective's actions. Thus virtual reasoning can be important in constructing the initial psychological characterisation, as well as in predicting its evolution.

As for the prediction of behaviour from psychological descriptions, I suggest that similarly mundane bits of common knowledge come into play. Thus there are certain things which we know—be it from the cradle, or be it by bitter experience—that people are liable to do if they fully intend to: moving chessmen according to a plan is one of them; meeting submission deadlines for papers is not. But again, these are just pieces of common knowledge like any other.

The essence of psychological concepts, I claim, lies the primary function these concepts perform in human thought. And that primary function is to facilitate the prediction of behaviour by means of virtual inference. A psychological description just is a description which tells us, in some sense, what it is like to be in the position of the subject.[7] It is a description which enables us immediately and unproblematically to imagine being in the described state and which therefore enables us, so I have been arguing, to use our own cognitive faculties to predict its likely development, whence we can predict the behaviour it is likely finally to induce. Of course, a physicalistic (or information-theoretic) description of our subject lacks this property: to describe someone's physical constitution and dispositions (or to detail the flow of information contained in his internal states) is not to be able to imagine being in his situation, unless we first infer a psychological characterisation from that description.

4 Constitutive Rationality and Logical Closure

Virtual inference lets us use the limitations of our own cognitive abilities to probe the likely limitations of those of others. Paradigmatically, our difficulty in seeing the long-term threat implicit in a chess position tells us that others might experience the same difficulty. In this respect, our ability to perform psychological inference trades on the fact, if it is a fact, that people think very much alike; it does not invoke any assumption about how rational the common inferential dispositions are. In particular, according to the view of psychological inference as psychological simulation, there is no assumption of general rationality to mediate any inference from how the subject ought to reason to how he will reason; on the contrary, psychological simulation is as good for predicting error as it is for predicting sound thinking.

In this respect, the view of psychological inference propounded above stands diametrically opposed to what now almost constitutes an orthodoxy in contemporary philosophy of mind, an orthodoxy maintained by such figures as Dennett, Davidson, Lewis, and Stalnaker. Writes Dennett:

> Suppose we travel to a distant planet and find it inhabited by things moving about its surface, multiplying, decaying, apparently reacting to events in the environment, but otherwise as unlike human beings as you please... Once we have tentatively identified the perils

[7] Cf. Nagel, Thomas: "What is it like to be a bat?", *Philosophical Review* 83 (1974).

> and succors of the environment..., we shall be able to estimate which goals and which weighting of goals will be optimal relative to the creatures' *needs...* . Having doped out these conditions (which will always be subject to revision) we can proceed at once to ascribe beliefs and desires to the creatures. Their behavior will "manifest" their beliefs by being seen as the actions which, given the creatures' desires, would be appropriate to such beliefs as would be appropriate to the environmental situation. Desires, in turn, will be "manifested" in behavior as those appropriate desires (given the needs of the creature) to which the actions of the creature would be appropriate, given the creature's beliefs.[8]

In other words, the attribution of beliefs and desires is constrained by the requirement that the beliefs and desires so attributed render the actions of the attributees rational—and therefore explicable—given those beliefs and desires.

To be sure, all who claim that assumptions of general rationality are constitutive of psychological concepts confess the possibility of occasional irrationality or non-rationality: of shoddy reasoning, of failing to believe logical consequences of what we believe, of forgetfulness, of destroying the things we love. Up to a point, that is. That such failings are possible means that the constitutive principles of rationality must be hedged with provisos that they only apply for the most part and in certain circumstances, and must be bolstered by principles (perhaps specific to an individual or culture) describing the more systematic ways in which subjects are likely to err. But there are limits to which constitutive rationality can be usefully etiolated in this way. Again, to quote Dennett:

> Not surprisingly, as we discover more and more imperfections (as we banish more and more logical truths from the creature's beliefs) our efforts at intentional prediction become more and more cumbersome and undecidable, for we can no longer count on the beliefs, desires and actions going together that *ought* to go together.[9]

Or: what is the point of attributing beliefs and desires when we, the attributors, are not going to make significant use of the logical (or confirmation-theoretic) interrelations between their contents, and how can we make use of those interrelations without some assumptions of general (if hedged) rationality?

(To which I answer: the point of attributing beliefs and desires is that the resulting descriptions are ones we can imagine to be true of ourselves, thus enabling us to predict their likely evolution by simulation.)

This brings us to the further question of logical closure. Common sense has it that, if someone believes that p, and if p implies q, then that person does not

[8] Dennett, "Intentional Systems", p.8. For relevant publications of the other authors just mentioned, see note 1.

[9] *Ibid.* p.11.

necessarily believe q. Not so, the advocates of constitutive rationality sometimes say: belief, contrary to what many suppose, is *closed* under logical consequence. At one point, Dennett writes:

> ...one gets nowhere with the assumption that entity x has beliefs p, q, r,... unless one also supposes that x believes what follows from p, q, r,...; otherwise there is no way of ruling out the prediction that x will, in the face of its beliefs p, q, r,... do something utterly stupid, and, if we cannot rule out *that* prediction, we will have acquired no predictive power at all.[10]

This argument is similar to the one rehearsed in the first section of this paper: that since the contents of psychological attitudes determine in the first instance, via the logical interrelations between their contents, how their possessors *ought* to use them, the explanatory and predictive utility of psychological characterisations depends upon an assumption of general rationality. Similar, but not quite identical: an argument for constitutive rationality is not by itself an argument for the logical closure of belief. For it is not obvious—as Dennett has since conceded—that simply trying to believe lots of logical consequences of one's current beliefs is a particularly rational cognitive policy. In a more recent article, Dennett qualifies his earlier position on the relationship between deductive closure and constitutive rationality:

> First, a few words on what rationality is *not*. It is not deductive closure.
>
> Nor is rationality perfect logical consistency, although the *discovery* of contradiction between propositions one is inclined to assent to is always, of course, an occasion for sounding the epistemic alarm... . Inconsistency, when discovered, is of course to be eliminated one way or another, but making the rooting out of inconsistency the preeminent goal of a cognizer would lead to swamping the cognitive system in bookkeeping and search operations to the exclusion of all other modes of activity... .[11]

Believing consequences of one's beliefs, and avoiding inconsistency among them are laudable activities, but only up to a point: if carried to excess they are likely to mar one's chances of doing anything useful.

Does that mean that Dennett thinks that we need not presume general rationality among believers after all? No, it just means that rationality cannot be equated with maintenance of consistency and striving for completeness. It may

[10] *Ibid.*

[11] *Ibid:* "Making Sense of Ourselves", in Dennett, Daniel, C.: *The Intentional Stance*, Cambridge, Massachusetts: MIT Press, 1987, pp. 94, 95.

be rational, says Dennett, quoting Cherniak with approval,[12] to be sometimes inconsistent because memory must be compartmentalized in order to make retrieval processes tollerably efficient, and because a compartmentalized memory will give rise to inconsistencies. Rational also, not to spend one's every waking hour making pointless logical deductions; rational to satisfice rather than maximize, and so on. Thus Dennett is denying the necessity of deductive closure and consistency without retracting the necessity of general rationality.

This move, Dennett thinks, blunts much of the criticism (notaby, that levelled by Stich) that his views have received. If the thesis of constitutive rationality seemed almost trivially false, that was because what had seemed to be paradigms of rationality—deductive consistency and logical closure—are in fact caricatures. If we adopt more realistic standards, Dennett assures us, the thesis of constitutive rationality is not so counterintuitive.

Not all constitutive rationalists take this view however. Stalnaker, for example, is explicit in arguing that if beliefs are attributed in such a way as to render the actions of their attributees explicable and predictable, then belief is logically closed:

> To say that an agent believes that P is to say something like this: the actions that are appropriate for that agent—those he is disposed to perform—are those that will tend to serve his interests and desires in situations in which P is true. But this is not quite right for the following reason:... [s]uppose I believe, as I do, that someone will be elected president of the United States in 1988. One way in which that proposition could be realized is for *me* to be the one elected, but I know that that is not the way my belief will come true. For my actions to be appropriate, given that I have this belief, it is surely not required that I take account of that possibility, since it is excluded by other of my beliefs. The actions that are appropriate for an agent who believes that P depend not only on what he wants but also on what else he believes. So it is necessary to define appropriateness relative to a total set of beliefs, or a belief state. And all that matters about such a belief state, as far as the appropriateness of actions or the agent's dispositions to act are concerned, are the entailments of the belief state. So there is no basis, on the dispositional account, for excluding from the set of an agent's beliefs any propositions that are entailed by those beliefs.[13]

Stalnaker argues from a constitutive principle of rationality (one that guarantees conformity of actions to beliefs and desires) to the conclusion that the objects of belief have a certain structure—a structure that takes a belief state to be a

[12] Cherniak, C: "Rationality and the Structure of Memory", *Synthese* 57, 1983.

[13] Stalnaker: *Inquiry,* p.82

logically closed set of propositions.[14] He is of course aware that this conclusion is apparently at odds with common usage, and he offers us an ingenious explanation of their compatibility by complicating the former and judiciously re-interpreting the latter.

But I do not wish to review these arguments here. Nor do I wish to adjudicate between Dennett's and Stalnaker's differing stances on the argument from constitutive rationality to logical closure, for my main argument is with the thesis of constitutive rationality itself. I mention Stalnaker's arguments on this question because his position contrasts so sharply with the view developed in this paper. On the psychological simulation view, it is of the essence of belief that one does not always believe everything one's beliefs entail. For it is, on this view, a paradigm of psychological inference to conclude that one's chess opponent might miss the point of one's new gambit because, putting oneself in his position, one does not find the threat it poses obvious; a paradigm also to conclude that the detective will not infer one's guilt, not because his evidence does not implicate one, but because one cannot imagine managing to make the required inferences oneself.

More generally, the contrast is this. For the constitutive rationality advocates, the happy hunting grounds of psychological inference are where we can assume optimal or rational functioning. Thus, we can ascribe beliefs, desires and intentions to a person or an alien or a computer with whom (which) we are playing chess, just as long as he plays well enough for the logical relations between the contents of these attitudes to be a useful guide to his behaviour. If we want to predict his mistakes, we have to resort to a theory of how our opponent functions (what Dennett calls the "design stance"); that is, we must leave behind psychological reasoning proper.

On the psychological simulation approach, by contrast, psychological inference only really comes into its own when we have to predict sub-optimal functioning. Thus, ascribing beliefs, desires and intentions to a person (though not perhaps to an alien or computer) with whom we are playing chess can be a useful guide to the threats and opportunities he is likely to miss, because such a psychological characterisation enables us to put ourselves in his cognitive predicament and experience his likely difficulties first hand. Prediction under assumptions of optimal functioning, say I, is a different matter, one which need make no interesting use of psychological notions at all. Such prediction is just a matter of the kind of look-ahead which is best done with a large memory and not much imagination; it does not require virtual reasoning.

In fact, Dennett briefly considers the idea that psychological inference might proceed by simulation, an idea he finds "lurking in Stich's view". But, says

[14]Such a conclusion does not make allowance for attitudes that are, to use David Lewis' phrase, irreducibly *de se*, though Stalnaker has made such allowances elsewhere. But this detail is unimportant for present purposes. David Lewis (see works cited in note 1) holds a very similar position to that of Stalnaker.

Dennett:

> How can [psychological simulation] work without being a kind of theorizing in the end?... If I make believe I am a suspension bridge and wonder what I will do when the wind blows, what "comes to me" in my make-believe state depends on how sophisticated my knowledge is of the physics and engineering of suspension bridges. Why should my making believe I have your beliefs be any different?[15]

Is reasoning-by-simulation a kind of theorizing, as Dennett suggests? In particular, when I reason about another's beliefs by means of virtual reasoning, does what "comes to me" in my make-believe state depend on my knowledge of the laws of human psychology? Well, yes and no. One could say that, in a sense, if people can predict the reasoning of others (even by simulation), they must *know* the laws of psychology, thought it would of course have to be admitted that this is knowledge so implicit that anyone who could articulate it as a theory (in the everyday sense) would have the profession of psychology at his feet. This way of talking, although not the one I have been using, I have no serious objection to: the interesting question then becomes that of how this 'knowledge' is stored in people's brains. And my answer is to refer the reader to the conjectures offered insection 2 of this paper.

But whether information thus encoded deserves to be called *knowledge of the laws of psychology* is beside the point. For Dennett's original argument had it that psychological ascriptions would be almost totally unconstrained were it not for some general principles mediating psychological inference, and that such principles could only be provided by some sort of general rationality assumption. Ergo, constitutive rationality. The point about psychological inference by simulation is that it enables psychological inference to trade on an assumption, not of general rationality, but of general uniformity. Whether or not that counts as having a psychological theory, rationality never gets into the picture.[16]

---

[15] Dennett, *ibid.* p. 100.

[16] In the same article, a few pages earlier, Dennett appears to endorse a line of argument according to which what people actually do think is evidence for how they ought to think, the reasoning being that our judgements about the latter rely on our "intuitions about what makes sense". Thus (p. 98): "When considering what we *ought to do*, our reflections lead us eventually to a consideration of what we *in fact do...*". So perhaps rationality is connected with psychological inference thus: psychological inference requires a theory (in a suitably broad sense of the word) of reasoning, and that theory of reasoning—i.e. our view of how we actually do reason—is ultimately the only evidence for our view about how we ought to reason.

Two points. First, even on the large assumption that Dennett is correct about the ultimate source of evidence about what we ought to do, it certainly does not follow that the principles of rationality which we arrive at will generally endorse human reasoning. Secondly, Dennett must beware of a strategic difficulty. He originally argued that psychological ascriptions would be unconstrained if we did not assume general rationality; now he claims that the ultimate evidence for our beliefs about rationality comes from how we think we (and others) *actually*

The situation is this, then. The constitutive rationality advocates think that psychological inferences and explanations are mediated by assumptions of general rationality, whence they argue a priori for (some sort of) general rationality, and whence some of them argue for logical closure; I, on the other hand, think that psychological inference proceeds by psychological simulation, whence I reject the argument for general rationality, and whence I argue against closure. Perhaps there is some useful notion which the constitutive rationality advocates are eliciting. The notions of 'belief', 'desire', 'intention' etc. which they characterise may indeed—though I have yet to be convinced—prove useful for predicting and explaining the behaviour of intelligent aliens, chess-playing computers and even humans. Furthermore, I can agree that these notions are useful for predicting the behaviour of *optimally functioning* creatures or entities, as Dennett suggests, and are subject to a form of logical closure, as Stalnaker argues. But I demur from calling these notions "belief", "desire", "intention" etc. (except within scare-quotes), and I demur from calling episodes of their employment "psychological inference". On the question of logical closure in particular, I prefer to side with the vulgar: of course people do not believe all the consequences of their beliefs. That, I have been arguing, is the whole point of the psychological notions, the feature without which the faculty of virtual reasoning would find no employment in psychological inference. The reason some constitutive rationality advocates depart from common usage in this respect is that they believe it impossible that psychological characterisations could have any predictive power—and hence psychological concepts any content—except under assumptions of rationality which ultimately violate this common sense. The burden of this paper has been to argue that this belief is untrue.

5 Conclusion

I trust I have not given the impression that it is *obvious* how psychological inference can proceed by means of virtual reasoning along the lines described. For I have not given an account of how it is that we imagine we have beliefs, desires, perceptions, plans, and predilections that we in fact do not have, or of how it is that we imagine we are oblivious of that which we know full well, while still managing to reason within such make-believe situations. Nevertheless, I have tried to make plausible the idea that we can reason effectively about the thinking of another person precisely because we can compare it with our own reasoning, the fact that he does not share all of our attitudes notwithstanding.

If so, I argued, then we can perform psychological reasoning even though we have no theory of how it is that people reason (why bother with a *theory* of reasoning when we can just *do* the reasoning?), and even though we make no pre-

think. But our having a view about how we (or others) actually think just *is* a matter of our making lots of psychological ascriptions. And what is to constrain those?

sumption of general rationality on the part of those we are reasoning about. So conceived, psychological reasoning represents an efficient use of our cognitive resources. The psychological processes involved in predicting our chess-opponent's likely responses are, in large measure, just those psychological processes we use to evaluate chess positions ourselves.

It is this kind of thinking that psychological notions are *for*. For it is of the essence of describing someone as having beliefs, desires, intentions and what have you, that one can immediately and unproblematically know what it is like to have those attitudes oneself. In so doing, one can immediately and unproblematically put oneself in a position to simulate the thinking that those attitudes are likely to induce. The same is not true of a description in terms of behaviour and dispositions to behave: one must first infer what psychological attitudes such a behavioural description is likely to correspond to before one can imagine being in the described situation oneself.

If this account of psychological inference is correct, it follows that the argument proffered by the constitutive rationality advocates for general rationality—an argument which assumes that psychological descriptions can wield explanatory or predictive power only under an assumption of general rationality—is undercut. So is the argument for logical closure. And the argument for constitutive rationality is undercut whether or not it is taken to imply logical closure, and whatever, exactly, the relevant benchmarks of rationality are taken to be.[17]

[17] I am indebted to Gilbert Harman, Paul Thagard, Dieter Münch and an anonymous referee for valuable help and suggestions.