

Partial belief and flat-out belief *

Keith Frankish

1. Introduction: a duality

There is a duality in our everyday view of belief. On the one hand, we sometimes speak of credence as a matter of degree. We talk of having some level of *confidence* in a claim (that a certain course of action is safe, for example, or that a desired event will occur) and explain our actions by reference to these degrees of confidence – tacitly appealing, it seems, to a probabilistic calculus such as that formalized in Bayesian decision theory. On the other hand, we also speak of belief as an unqualified, or flat-out, state (‘plain belief’ as it is sometimes called), which is either categorically present or categorically absent. We talk of simply *believing* or *thinking* that something is the case, and we cite these flat-out attitudes in explanation of our actions – appealing to classical practical reasoning of the sort formalized in the so-called ‘practical syllogism’.¹

This tension in everyday discourse is reflected in the theoretical literature on belief. In formal epistemology there is a division between those in the Bayesian tradition, who treat credence as graded, and those who think of it as a categorical attitude of some kind. The Bayesian perspective also contrasts with the dominant view in philosophy of mind, where belief is widely regarded as a categorical state (a token sentence of a mental language, inscribed in a functionally defined ‘belief box’, according to one popular account). A parallel duality is present in our everyday view of desire. Sometimes we talk of having degrees of preference or desirability; sometimes we speak simply of *wanting* or *desiring* something tout court, and, again, this tension is reflected in the theoretical literature.

What should we make of these dualities? Are there two different types of belief and desire – *partial* and *flat-out*, as they are sometimes called? If so, how are they related? And how could both have a role in guiding rational action, as the everyday view has it? The last question poses a particular challenge in relation to flat-out belief and desire. Bayesian decision theory teaches us that the rational way to make decisions is to assign degrees of probability and desirability to the various possible outcomes of candidate actions and then choose the one that offers the best trade-off of desirability and likely success. Flat-out belief and desire just do not come into the picture. The everyday view thus faces a challenge: how can flat-out belief and desire have the psychological importance they seem to have, given their apparent irrelevance to

* NOTICE: this is the author’s version of a chapter published in F. Huber and C. Schmidt-Petri (eds), *Degrees of Belief* (pp. 75-93), Springer, 2009. It may not reflect changes resulting from the publishing process.

¹ It is worth stressing that it is the *attitude* that is unqualified in flat-out belief, not necessarily the *content*. It is possible for a flat-out belief to have a qualified content – for example, that there is a 75% chance of rain tomorrow.

rational action? Borrowing a term from Mark Kaplan, I shall refer to this as the *Bayesian Challenge* (Kaplan 1996).²

The challenge has an especially strong bite if one adopts an *interpretivist* view of the mind of the sort advocated by, among others, Donald Davidson and Daniel Dennett (Davidson 1975; Dennett 1981). On such a view, adherence to the norms of practical rationality is a precondition for the possession of mental states: we credit an agent with just those degrees of belief and desire that best rationalize their intentional actions, and we regard an item of behaviour *as* an intentional action just in case we can interpret it as the rational expression of a coherent set of subjective probabilities and desirabilities. Given a Bayesian view of practical rationality, it follows that habitual violation of Bayesian norms is conceptually impossible; the behaviour in question simply would not qualify as intentional. From this perspective, then, flat-out belief and desire are irrelevant to intentional action tout court. Indeed, it is hard to see what sense we could make of a person's *acting upon* a flat-out belief or desire. For any action they performed would, as a matter of conceptual necessity, already have a complete intentional explanation in terms of their partial beliefs and desires. How, then, could flat-out states get into the picture?

In this chapter I shall offer an explanation of the duality in our everyday view of belief and desire. I shall begin by arguing that some widespread views of the relation between partial and flat-out belief are inadequate, then go on to outline a better account, which not only does justice to our intuitions about these states and their cognitive role, but also repulses the Bayesian challenge. I shall focus primarily on belief but shall occasionally mention desire, too, and shall treat it as a condition for a successful account of belief that it should extend in a natural way to desire. Note that my primary aim will be descriptive: I shall aim to articulate notions of partial and flat-out belief that reflect our everyday intuitions about these states. The discussion will nonetheless have a bearing on issues in formal epistemology, and I shall say something about this at the end of the chapter.

2. Partial belief as a derivative of flat-out belief

Although the duality in everyday belief talk is widely recognized, few writers claim that it points to the existence of two distinct kinds of psychological state. It is more common to suppose that it highlights two *aspects* of a single one. The general assumption is that one form of belief is the core state and the other a derivative, which can be defined in terms of the core one. There are two strategies here, depending on whether flat-out belief or partial belief is taken to be the core state. I shall consider the first strategy in this section and the second in the next.

Suppose that belief is fundamentally a flat-out state, which is either categorically present or categorically absent. What then might we mean when we talk of having

² Note, however, that I interpret the challenge in a broader sense than Kaplan. Kaplan is concerned with the question of how flat-out belief can have a role in theoretical inquiry; I am interested in how it can figure in *practical reasoning* as well.

degrees of belief? One option is to maintain that we are referring to flat-out beliefs in explicit probability claims. So, for example, when we say that someone is 50% confident that a coin toss will come up heads, what we mean is that they believe flat-out that the probability of its coming up heads is 0.5. Similarly, it might be suggested, degrees of preference can be identified with flat-out beliefs in the desirability of various situations.

Now it is true that we do sometimes form flat-out beliefs about probabilities and desirabilities, but it is implausible to identify our degrees of confidence and preference with such beliefs. For we attribute degrees of confidence and preference to individuals who lack the conceptual sophistication required to form beliefs of this kind. We speak of children and animals having more or less confidence in something and preferring one thing to another, even though they do not possess the concepts of probability and desirability. Indeed, even if a person is capable of forming beliefs about probabilities and desirabilities, we do not suppose that they have actually formed appropriate beliefs of this kind every time their behaviour manifests some degree of confidence or preference – certainly not that they have done so consciously, and I do not think we are committed to the view that they have done so non-consciously either.

The problem comes into sharper focus when we consider how the relevant flat-out beliefs are to be construed. They cannot be beliefs about one's own degrees of confidence, on pain of regress. If the aim is to analyse talk of partial belief in other terms, then we must not make reference to partial beliefs in the analysis. Nor will it do to think of the beliefs in question as concerning so-called 'objective probabilities', understood as facts about the frequencies of certain types of event or about the propensities of objects. For we can have degrees of confidence in single events, which do not have objective probabilities in this sense. (Again, this is not to say that people do not form beliefs about objective probabilities, just that we cannot analyse degrees of confidence in terms of them.) This does not exhaust the options, of course. It might be suggested that we could identify partial beliefs with beliefs about the degree to which hypotheses are confirmed by the evidence (so-called 'logical probabilities') or with beliefs about one's own behavioural dispositions (say, how willing one would be to bet on various propositions). However, these suggestions succumb even more clearly to the earlier objection: one can have preferences and degrees of confidence without knowing anything about confirmation relations or betting odds. A similar objection will hold, I suspect, for any other proposal along these lines.

On the view just considered, talk of partial belief serves to characterize an aspect of the *content* of our flat-out beliefs. Another option, advocated by Gilbert Harman, among others, is to see it as characterizing an aspect of the *attitude*. When we talk of our degrees of belief, Harman claims, we are referring to how *strongly held* our flat-out beliefs are, where this is a matter of how hard it would be for us to give them up (Harman 1986, ch.3). So, one either believes a proposition or does not believe it, but if one does, then one does so with a certain degree of strength or attachment. (Harman stresses that this need not involve making an explicit assessment of how important the belief is to us; our degree of attachment to it may be implicit in the way we reason –

the more attached to it we are, the more powerful the reasons needed to get us to abandon it.) I suspect that this view is held, at least tacitly, by many philosophers of mind, and it can be easily extended to desire. When we talk of our degrees of preference we can be understood as referring to the strength of our flat-out desires, where this is a matter of their tenacity or their power to override competing desires.

This view is also unsatisfactory, however. The main problem is that it means that one would have to believe a proposition flat-out in order to have a degree of belief in it. And this is surely wrong. I have some degree of confidence (less than 50%) in the proposition that it will rain tomorrow, but I do not believe flat-out that it will rain – not, at least, by the everyday standards for flat-out belief. (If anything, I believe the opposite.) Indeed, according to Bayesian principles, a rational agent will entertain some degree of confidence in every proposition of whose falsity they are not certain, including pairs that are contradictory. Yet it would be absurd to say that a rational agent will have a flat-out belief in every proposition to which they ascribe non-zero probability. Similar considerations apply to desire. I can prefer to be maimed rather than killed without desiring flat-out to be maimed. (This is not to say, of course, that we do not have degrees of attachment to our flat-out beliefs and desires, just that our degrees of confidence cannot be identified with them.)

Finally, note that both proposals considered still face the Bayesian challenge. On the former view, partial beliefs form only a subset of our flat-out beliefs, and it is hard to see how it could be rational to act on the remaining ones. And on the second proposal it is unclear how reasoning involving flat-out beliefs and desires could be sensitive to our degrees of attachment to those states in the way required by Bayesian decision theory.

It is implausible, then, to regard partial belief as a derivative of flat-out belief, and in what follows I shall treat it as a psychological primitive. This is not the same as an ontological primitive, of course; there will be some story to be told about how partial beliefs are physically realized. But I shall assume that this story will not appeal to psychological notions.

3. Flat-out belief as a derivative of partial belief

Turn now to the opposite view – that partial belief is the core state and flat-out belief the derivative. The obvious strategy here is to identify flat-out belief with some level of confidence, either the maximum (1 on the scale of 0-1 on which degrees of confidence are standardly measured) or a level exceeding some threshold short of the maximum. This view is assumed either explicitly or tacitly by many epistemologists and philosophers of science working within the Bayesian tradition, and it has the advantage of dispelling the Bayesian challenge. If talk of a person's flat-out beliefs is simply an alternative way of characterizing their degrees of confidence, then Bayesian theory poses no threat to it. The view faces serious difficulties, however.

Consider first the suggestion that flat-out belief is maximum confidence – a view reflected in the frequent use of the term 'full belief' for flat-out belief. The problem

here is that one can believe something, in the everyday sense, without being certain of it. I believe that my grandmother was born on the 3rd of August, but I am not absolutely certain of it. I may have misremembered or been misinformed. Nor is this lack of certainty necessarily a bad thing; a fallibilist attitude to one's own beliefs has much to recommend it. Another difficulty for the full-belief view arises in connection with practical reasoning. On Bayesian principles, to assign a probability of 1 to a proposition is to cease to contemplate the possibility that it is false and, consequently, to ignore the undesirability of any outcome contingent upon its falsity.³ One consequence of this is that if one is certain of something, then one should be prepared, on pain of irrationality, to bet everything one has on its truth for no return at all. For one will simply discount the possibility of losing the bet. (This is the problem of the 'all-for-nothing bet'; see Kaplan 1996; Maher 1986, 1993.) Yet we can believe something, in the everyday sense, without being prepared to stake everything, or even very much, on its truth. (I would bet something, but not a great deal, on the truth of my belief about my grandmother's birth date.) So flat-out belief is not the same thing as maximum probability. A third problem for the full-belief view is that it does not extend to desire. One can desire something, in the everyday sense, without assigning it maximum desirability. I want a new car, but I do not regard a new car as the most desirable thing in the world.

Suppose, then, that flat-out belief corresponds to a high level of confidence, albeit one that falls short of the maximum. For example, we might say that a person has a flat-out belief in a proposition if their confidence in it is greater than their confidence in its negation – that is, if it exceeds 0.5. Richard Foley dubs this *the Lockean thesis* (Foley 1993, ch.4).⁴ Flat-out desire might similarly be identified with a certain level of desirability.

Although popular, this view also faces serious difficulties. First, the norms of flat-out belief are different from those of high confidence. Flat-out belief, as we commonly conceive of it, is subject to a conjunctive constraint: we accept that we ought to believe the conjunction of any propositions we believe (provided, at least, that we recognize it as such; it is plausible to think that rational norms are tailored to our needs and capacities as finite agents). This principle – *conjunctive closure* – underlies some

³ In assessing a course of action, the Bayesian calculates the desirabilities of the various possible outcomes it might have, weighting each by the probability of its occurring. Now suppose that a particular action, *A*, would have a very bad outcome if condition *C* obtained. That is to say, suppose that the desirability of performing *A* in *C* – symbolized as *des(A(C))* – is strongly negative. (Attractive outcomes have positive desirabilities, unattractive ones negative desirabilities, and neutral ones a desirability of zero.) Normally, this would count against performing *A*, even if one were fairly confident that *C* did not obtain. But now suppose one is *certain* that *C* does not obtain – i.e. one assigns *prob(C)* a value of 0. Then the weighted desirability of performing *A* when *C* obtains – *prob(C) × des(A(C))* – will also be zero, no matter how low *des(A(C))* is. That is, one should be indifferent between performing *A* and maintaining the status quo and should be willing to perform it for no return at all. The possibility of the bad outcome will be completely ignored.

⁴ For Locke's endorsement of the view, see his *Essay Concerning Human Understanding*, Book IV, chs 15 and 16. Modern adherents of the Lockean thesis include Chisholm (1957) and Sellars (1964). See also Foley's paper in the present volume.

important deliberative practices. Reasoning often involves conjoining propositions – putting together things one believes and then deriving a conclusion from them. A person who rejected the principle would not regard such inferences as compelling. Nor, by the same token, would they be troubled by the discovery that their beliefs conjointly entail a contradiction (Kaplan 1981, p.133). Yet if flat-out belief is high confidence then it will not be subject to conjunctive closure. For on Bayesian principles it will frequently be rational to assign a lower probability to a conjunction than to any of its conjuncts individually. And given the right numbers, this might make the difference between meeting and failing to meet a threshold for belief, even when the number of conjuncts is small.

A second problem for the Lockean thesis concerns the influence of belief on action. The common-sense conception of flat-out belief is that of a state which makes a qualitative difference to one's behavioural dispositions. Of course, just how the addition of any particular flat-out belief changes one's dispositions will depend on what other mental states one has – what other beliefs, what desires and intentions, and so on. But whatever the background, the addition of a new flat-out belief ought to make *some* qualitative difference. (This is true even if we restrict our attention to the context of theoretical inquiry. If a scientist forms the flat-out belief that a certain theory is true, then we should expect this to make a qualitative difference to what they are disposed to say and do in their professional capacity.) But if the Lockean thesis is true, the acquisition of a new flat-out belief may not involve any significant change at all. The move from a degree of confidence that just falls short of a threshold to one that just exceeds it may or may not make a qualitative difference to how one is disposed to act. It will all depend on one's background probabilities and desirabilities. Given one background, a threshold-crossing change in one's confidence assignment to a particular proposition may make a great difference; given another, it may make none. At any rate, it will not, as a rule, make a *greater* difference than a change of similar extent anywhere else along the confidence scale. Similar considerations apply to desire.

We can make the same point in a slightly different way. The folk notion of flat-out belief is that of an *explanatorily salient* psychological state. That is to say, the folk view is that our flat-out beliefs can be cited in explanation of our actions – or of some of them, at least. But this will be the case only if there are robust, counterfactual-supporting generalizations linking flat-out belief with action. And if the Lockean thesis is correct, there will be no such generalizations, since, as we have just seen, the acquisition of a new flat-out belief will often make no significant difference to how one is disposed to act. And, given this, it is hard to see why the state should hold any theoretical interest for us.⁵ It is worth stressing that the objection here is not that flat-out beliefs will be causally idle – they will, after all, possess just as much causal power as the states of confidence in which they consist. Rather, it is that they will not possess their causal powers *in virtue of* being states of flat-out belief; there will be no

⁵ Stalnaker makes a similar point (1984, p.91).

psychological laws defined over flat-out beliefs. The Lockean thesis thus avoids the Bayesian challenge only at the cost of making flat-out belief *qua* flat-out belief explanatorily idle. I conclude, then, that the Lockean thesis does not capture the everyday conception of flat-out belief. Again, the same goes for desire.

4. Flat-out belief as an intentional disposition

If we reject the full-belief view and the Lockean thesis, what options are left? Must we conclude that flat-out belief is also a psychological primitive? That would be an unwelcome conclusion. It would leave us with a mystery as to the relation between the two forms of belief and with seemingly no hope of fending off the Bayesian challenge. Indeed, if that were the only option, one might be inclined to doubt whether there really was such a thing as flat-out belief at all. There is another option, however. This agrees with the views just considered in treating partial belief as fundamental but postulates a more complex relation between it and flat-out belief. The idea is to think of flat-out beliefs as behavioural dispositions *arising from* the agent's partial beliefs and desires – *intentional dispositions*, we might call them. Thus, to say that a person has a flat-out belief with content *p* is to say that they have partial beliefs and desires such that they are disposed to act in a certain way – a way characterized by reference to the proposition *p* (examples will make this clear shortly). I shall refer to this generic position as the *behavioural view* of flat-out belief.

On this view, acquiring a new flat-out belief will make a qualitative difference to one's behavioural dispositions; indeed, it simply *is* acquiring a certain sort of behavioural disposition. And flat-out belief will accordingly have some explanatory salience; there will be counterfactual-supporting generalizations linking flat-out belief with actions that manifest the disposition. Yet because it treats partial belief as fundamental, the behavioural view also has the resources to answer the Bayesian challenge. The important point is that, on this view, flat-out beliefs are not something *over and above* partial beliefs and desires. An intentional disposition (say, to save money) exists in virtue of a set of underlying partial beliefs and desires which are, given a normal cognitive background, sufficient for it. In a widely used phrase, the disposition is *realized in* the partial states. Thus, on the behavioural view, flat-out beliefs are realized in partial beliefs and desires, and they will have an influence on action precisely equal to that of the realizing states. As we shall see, this offers a powerful response to the Bayesian challenge.

Given this general approach, what sort of intentional disposition might flat-out belief consist in? One option would be to identify it with a disposition to act as if the proposition believed were true – that is, to prefer options that would be more attractive if it were true over ones that would be more attractive if it were false. Following Kaplan, I shall refer to this as the *act view*.⁶ This view is unattractive, however, as Kaplan points out (Kaplan 1996, pp.104-6). For unless they are certain

⁶ Kaplan ascribes versions of the act view to Braithwaite (1932-3), Churchman (1956), Rudner (1953), and Teller (1980).

that p is true, rational agents will be disposed to act as if p is true in some circumstances and disposed to act as if it is false in others – it will all depend on the options available. For example, consider two scenarios. In (1) you are offered a choice between the status quo and betting £1 on p for a return of £10; in (2) you are offered a choice between the status quo and betting £10 on p for a return of £1. And suppose that your confidence in p is about 0.5. Then, if rational, you will be disposed to take the bet in scenario (1), thereby acting as if p is true, but disposed to reject the bet in scenario (2), thereby acting as if *Not- p* is true. It follows that a defender of the act view must either insist that flat-out belief requires certainty or accept that flat-out belief can vary with context, and neither option is attractive. As we saw earlier, it is implausible to claim that flat-out belief requires certainty. And flat-out belief, as commonly understood, is unqualified as to context, as well as to attitude. If we believe a proposition, under a certain mode of presentation, then we believe it under that mode of presentation in all contexts, not just when certain options are presented to us.

A more plausible option is to think of flat-out belief as a specific intentional disposition, linked to a limited range of activities. One suggestion is that it consists in a linguistic disposition – a disposition sincerely to assert the proposition believed. This view has its roots in an influential 1971 paper by Ronald de Sousa (de Sousa 1971). De Sousa suggests that we harbour an ‘epistemic lust’ – a hankering for objects of unqualified epistemic virtue – which prompts us to make all-out epistemic commitments to propositions (‘bets on truth’). And flat-out beliefs, de Sousa proposes, are simply dispositions to make or manifest such commitments. More recently, Mark Kaplan has developed a detailed version of this proposal, which he calls the *assertion view* of belief (Kaplan 1996; see also Maher 1993). According to Kaplan:

You count as believing [i.e. believing flat-out] P just if, were your sole aim to assert the truth (as it pertains to P), and your only options were to assert that P , assert that $\sim P$ or make neither assertion, you would prefer to assert that P .
(Kaplan 1996, p.109)

On this view, Kaplan points out, there is no simple relationship between flat-out belief and confidence. ‘The truth’ can be thought of as a comprehensive error-free account of the world, and in deciding what to assert when one’s aim is to assert the truth one must strike a balance between the aims of attaining comprehensiveness and of avoiding error. The assertability of a proposition will thus be a function, not only of one’s confidence in its truth, but also of one’s estimate of its informativeness, together with the relative strengths of one’s desires to shun error and to attain comprehensiveness.

This view avoids the pitfalls of the act view. Understood in this way, flat-out belief does not require certainty and is not context-dependent. There is only one context that is criterial for belief possession – that of disinterested inquiry, where one’s sole aim is to assert the truth – and provided one has a constant disposition to assert that p in this context, one counts as believing p flat-out, even if one does not always act as if p were true. Another attraction of the assertion view is that it does not conflict with

our common-sense commitment to the conjunctive constraint, as the Lockean thesis does. The conflict arose because confidence is not preserved over conjunction. But once we abandon the view that flat-out belief is a level of confidence, this no longer presents a problem. We can insist that we should be prepared to conjoin any pair of claims we believe, even if this means believing propositions to which we assign a relatively low probability – the increased risk of error in such cases being offset by the gain in comprehensiveness.

This view of flat-out belief is, then, closer to the everyday one than the others we have considered. Nevertheless it is not the everyday one. For although it accords flat-out belief some psychological salience, it does not accord it the sort of salience we commonly take that state to have. We think of flat-out belief as a state that enters into practical reasoning and has an open-ended role in the guidance of action. (This is true even if we restrict our attention to beliefs formed in the context of theoretical inquiry; a scientist's theoretical beliefs will have an open-ended influence on the conduct of their research.) Yet, as Kaplan himself emphasizes, on the assertion view flat-out belief is linked to only one action – assertion.⁷ A related difference is that, on the folk view, the influence of belief on action is *desire-mediated* – beliefs influence action only in combination with desires. On the assertion view, however, flat-out beliefs dispose us to act directly, without the involvement of desires. Indeed, the assertion view offers no account of what flat-out desires might be. I see no reason for denying that we can form linguistic dispositions of the sort Kaplan describes. We do sometimes give verbal endorsement to a proposition without taking it to heart and acting upon it. And we might, at a pinch, describe such a disposition as a form of belief. But it is an etiolated kind, not the everyday action-guiding variety.

5. Flat-out belief as a premising policy

I want to turn now to another way of developing the behavioural view of flat-out belief, which offers a more robust account of the cognitive role of the state. In recent years a number of writers have drawn a distinction between *belief* and *acceptance*. This has links with the distinction between partial and flat-out belief (the latter, too, is often referred to as 'acceptance'), but there are significant differences; indeed, writers on acceptance typically insist that it is not a form of belief at all, strictly speaking. There are a number of independent versions of the belief/acceptance distinction, each addressing different concerns and fostering different conceptions of the two states (for a survey, see Engel 2000). I want to focus on a version developed by Jonathan Cohen

⁷ Kaplan writes:

It is, I think, a mistake to suppose that we need recourse to talk about belief [i.e. flat-out belief] in order adequately to describe the doxastic input into rational decision making. ... That task has been taken over, without residue, by our talk of confidence rankings. (Kaplan 1996, p.107)

See also Maher 1993, pp.149-52.

(Cohen 1989, 1992; for similar accounts, see Bratman 1992; Engel 1998; Stalnaker 1984).

According to Cohen, belief is a disposition to entertain 'credal feelings'; to believe that p is to be disposed to *feel it true* that p when you consider the matter (Cohen 1992, p.5). Like other affective dispositions, Cohen explains, belief is involuntary and varies in intensity. Acceptance, on the other hand, is a policy, which can be actively adopted in response to pragmatic considerations. To accept a proposition is to decide to *treat it as true* for the purposes of certain kinds of reasoning and decision-making. It is, in Cohen's words:

to have or adopt a policy of deeming, positing, or postulating that p – i.e. of including that proposition or rule among one's premisses for deciding what to do or think in a particular context, whether or not one feels it to be true that p . (1992, p.4)

We are able to adopt such policies, Cohen implies, because our conscious reasoning is, to some extent, under our personal control. Acceptance-based mental processes are, he says, 'consciously guided by voluntarily accepted rules' (1992, p.56) – the rules in question being logical, conceptual, or mathematical ones of the sort we acquire in the course of learning to engage in reasoned argument. Cohen also identifies a parallel conative state – *goal adoption* – which is related to acceptance as desire is to belief (1992, p.44-45). (For convenience I shall use the term 'premissing policy' for both acceptance and goal adoption.)

Cohen does not say much about what is involved in executing premissing policies, but the general idea is clear enough: in accepting a premise or adopting a goal, we commit ourselves to a policy of taking it as an input to our conscious reasoning and decision-making, manipulating it in accordance with whatever inference rules we accept. Typically, these will include the rules of classical deductive logic (or informal versions of them), together with simple, non-probabilistic rules of practical and inductive reasoning. I shall assume that we also commit ourselves to acting upon the results of these calculations – adopting any conclusions as further premises or goals, and performing, or forming intentions to perform, any dictated actions. (These commitments will, of course, be contingent upon the continued acceptance of the premise or goal itself; we can repudiate a policy if we find its dictates unacceptable.) I take it that acts of policy adoption themselves are performative ones, like promises, which both announce and create the commitment in question. In the case of premissing policies, these actions will of course typically be silent and internalized. When I speak of a person being *committed* to an action, or course of action, I mean that they regard themselves as obliged to perform it – often, though not always, in virtue of some prior performative of this kind.⁸

⁸ For detailed discussion of the nature of premissing policies and the procedures involved in their execution, see my 2004, ch.4.

Now, Cohen's conception of belief is somewhat idiosyncratic, and I shall set it aside, but his account of acceptance is very relevant to our concerns. For acceptance in Cohen's sense has many of the characteristics of flat-out belief, as we commonly conceive of it. First, it is an all-or-nothing state: for any proposition, *p*, and any deliberative context, *C*, one either has or has not adopted a policy of premising that *p* in *C*. (It is true that we may have varying degrees of *attachment* to our premising policies – finding some easier to give up than others. But so long as we hold on to a given set of policies, our commitment to each of them will be the same.) Secondly, acceptance is subject to the same inferential norms as flat-out belief. This is definitional: accepting a proposition involves committing oneself to manipulating it in accordance with whatever inferential rules one recognizes. Thirdly, acceptance will be salient in the explanation of action and inference in the way we take flat-out belief to be. In accepting a proposition we undertake, not just to *speak* as if the proposition accepted were true, but to *reason* and *act* as if it were. And, fourthly, this account of flat-out belief extends smoothly to flat-out desire, which can be identified with goal adoption.

Moreover, although Cohen does not present it in this way, it is natural to view acceptances and goal adoptions as intentional dispositions. If one has a policy of performing some action then one will be motivated to perform it on appropriate occasions – that is, one will have an intentional disposition to perform it. Not every intentional disposition counts as a policy, however; one can be intentionally disposed to eat chocolate biscuits at regular intervals without having a policy of doing so. The difference lies in the nature of the motivation. The motivation for a policy-related action is not simply the desirability of the action itself or of its immediate outcome, but rather the desirability of adhering to the policy. We adopt policies because we think that they will bring some long-term benefit, and we adhere to them because we want to secure this benefit. This sort of motivation is, I suggest, the defining characteristic of policy-based action. That is, to have a policy of *A*-ing is to be disposed to *A* precisely because one regards oneself as having embarked upon a course of *A*-ing and attaches a high desirability to sticking to it.⁹ The present suggestion thus affords a way of developing the behavioural view of flat-out belief outlined earlier. If flat-out beliefs are policies of premising, then one will count as having a flat-out belief with content *p* if one is highly confident that one has embarked upon a course of premising that *p* and attaches a high desirability to adhering to it. These partial states will be sufficient for the existence of the flat-out state, which can be thought of as being realized in them.

⁹ This is slightly simplified. Typically, I think, the *immediate* motive for a policy-related action will be a general desire to adhere to whatever policies one has adopted – to honour one's policy commitments, as it were. We cannot always keep in mind the reasons for pursuing particular policies, so a general desire of this kind will be useful, helping to promote consistency in our activities and underwriting long-term planning. On this view, then, the desire to obtain the benefits of following a particular policy will be an *indirect* motive for the ensuing actions, having served to motivate the adoption of the policy in the first place.

It is worth stressing that the partial beliefs and desires that sustain a premising policy need not be consciously entertained. Actions can be consciously performed even if the beliefs and desires that motivate them are not themselves conscious. For example, I consciously press various keys on my computer keyboard because I desire to type certain words and believe that pressing them will achieve that. But I do not consciously entertain those beliefs and desires as I hit the keys; I just think about the *content* of what I am typing. Much the same, I assume, would go for the beliefs and desires that support a premising policy. Typically, all I would think about at a conscious level is the content of my premises and goals. The mental states that motivate my premising activities would reveal themselves in my attitude *to* these contents – in the fact that I regard myself as committed to my premises and goals and bound to act upon their consequences.

Despite all this, we cannot simply identify flat-out belief with acceptance. Indeed, writers on acceptance typically deny that it is a form of belief and stress the apparent differences between it and belief. I shall mention two of these. First, acceptance is responsive to prudential considerations in a way that belief appears not to be. We can decide to treat a proposition as true for a variety of reasons – not only evidential, but also professional, ethical, religious, and so on (Cohen 1992, p.12, p.20). For example, professional ethics may oblige a lawyer to accept that their client is innocent for the purposes of defending them. Secondly, acceptance, unlike belief, can be context-relative. We can accept something when reasoning on certain topics while rejecting it, or suspending judgement, when reasoning on others. The lawyer may accept their client's innocence when planning their defence, but not when reasoning about personal matters. Belief, on the other hand, does not seem to be compartmentalized in this way.

It remains possible, however, that flat-out beliefs form a subset of acceptances, and this is the view I want to propose. The question of how the subset is delimited is a complex one, and here I can only outline the view I favour (for extended defence of this view, see my 2004, ch.5). Given the points just made, the obvious suggestion would be that the distinguishing factors are motivation and context – that flat-out beliefs are those acceptances that are epistemically motivated and unrestricted as to context. This is too swift, however. Motivation is, I think, irrelevant here. In general, aetiological considerations are not decisive in determining beliefhood; we would not deny a mental state the title of belief simply because it had been formed in an anomalous way. Moreover, while it is true that we cannot choose to believe anything we like, it is arguable that there is a sense in which belief can be responsive to pragmatic reasons (see Frankish 2007). Context, on the other hand, is important. Belief, in the everyday sense, plays an open-ended role in deliberation, and acceptances that are restricted to particular deliberative contexts do not count as beliefs in this sense.

Not all unrestricted acceptances are beliefs, however. A person might accept a proposition for general reasoning purposes without actually believing it. Think, for example, of a highly insecure person who, for therapeutic reasons, accepts it as a

general premise that they are capable and successful, even though they do not really believe that they are. For an acceptance to count as a belief, I suggest, it must not only be unrestricted, but must also extend to an important class of deliberations which I shall call *truth-critical with respect to premises*, or *TCP* for short. A TCP deliberation is one where we attach an overridingly high desirability to taking only truths as premises – for example, theoretical deliberations where there are large penalties for coming to false conclusions. The insecure person who accepts for therapeutic reasons that they are capable and successful does not count as believing that proposition, since they would not be willing to rely on it in TCP deliberations. (Therapeutic considerations may influence *which* deliberations they treat as TCP – perhaps determining that they treat few as such – but they will not affect their willingness to use a proposition as a premise in those deliberations they *do* treat as TCP.) Note that, if rational, one will be willing to take a proposition as a premise in TCP deliberations only if one has a high degree of confidence in it (higher at least than in any of the relevant alternatives), so it is a corollary of this view that high confidence is necessary for flat-out belief – though not, of course, sufficient for it.

I shall return to this suggestion at the end of the chapter, but for present purposes nothing turns on it. The important claim is that flat-out beliefs form a subset of acceptances; the exact boundaries of the subset do not matter. Note that no similar complication arises in the identification of flat-out desires with goal pursuits, since folk psychology recognizes the existence of instrumental desires that are formed for pragmatic reasons and effective only within restricted contexts.

6. Flat-out belief and action

I turn now to the role of flat-out belief and desire in the guidance of action and to my response to the Bayesian challenge. Here the view just sketched has an important consequence. It is that our premises and goals influence our behaviour *in virtue of* our partial beliefs and desires concerning them. In adopting premises and goals we commit ourselves to performing any actions they dictate. So if we believe that our premises and goals mandate a certain action, and want to honour our commitment to them, then we shall be motivated to perform the action for that very reason. In other words, when we perform an action in response to our flat-out beliefs and desires we do so because we are highly confident that we have premising policies that dictate the action and attach a high desirability to sticking to them. (Again, these partial attitudes need not, and typically will not, be conscious ones.)

Any such action will thus have two different intentional explanations – one citing the flat-out beliefs and desires involved, and another citing the partial beliefs and desires in which these flat-out attitudes are realized. Take a simple case. Suppose I am ordering in my favourite Chinese restaurant. I have just requested spring rolls, when I consciously recall that these items are high in fat and that I want to cut down my fat intake – these thoughts being recollections of a pair of previously formed flat-out beliefs and desires. I instantly calculate that these attitudes dictate a change of order

and request something else. Here it would be natural to explain my action simply by citing the flat-out belief and desire just mentioned. However, since these attitudes influence my behaviour in virtue of my partial beliefs and desires concerning them, there will be another – much less obvious – intentional explanation of my action, which represents it as the outcome of Bayesian decision-making involving non-conscious partial attitudes. Thus, upon consciously recalling my flat-out belief and desire, and calculating that they warranted a change of order, I became highly confident that I had premising policies that warranted a change of order. Since I strongly desired to adhere to my premising policies and to perform the actions they dictated, the option of changing my order now had the highest overall estimated desirability, so I took it. A similar explanation will be available whenever an agent acts upon flat-out beliefs and desires. In each case, the explanation citing such attitudes will be compatible with, and underpinned by, another explanation, which represents the action as the rational product of the agent's non-conscious partial beliefs and desires. Of course, not all actions will have dual explanations of this kind, but only those that are the product of conscious deliberation involving flat-out beliefs and desires. Much of our behaviour is of a less reflective kind and will (on this view) be explicable only in terms of partial attitudes. But any action that *does* have an explanation in terms of flat-out attitudes will also have an explanation in terms of partial ones.

This view gives us a response to the Bayesian challenge. The challenge was to explain how it could be rational to act upon flat-out beliefs and desires and to ignore degrees of confidence and desirability (indeed, how it could be *possible*, given an interpretivist view of the mind). Now, the challenge assumes that motivation by flat-out beliefs and desires is incompatible with motivation by partial ones – that actions cannot be simultaneously motivated both by flat-out states and by partial ones, and thus cannot be simultaneously responsive both to the norms of classical practical reasoning and to those of Bayesian decision-making. But if flat-out beliefs and desires are premising policies of the kind described, then this assumption is false. For, as we have seen, on this view our flat-out attitudes influence our behaviour in virtue of our partial beliefs and desires concerning them. We act on our premises and goals because we attach a high desirability to adhering to our premising policies and to performing the actions they dictate. Thus, in acting upon our flat-out beliefs and desires we are not departing *from* Bayesian norms but adhering *to* them, and the resulting actions can be justified both on classical grounds, as dictated by our flat-out beliefs and desires, and on Bayesian grounds, as warranted by our probabilities and desirabilities. The Bayesian challenge is thus neutralized.

For the same reason, belief in the efficacy of flat-out belief and desire is compatible with a broadly interpretivist view of the mind. Given a Bayesian conception of practical rationality, interpretivism dictates that all intentional actions will, as a matter of conceptual necessity, have explanations in terms of partial beliefs and desires. But, again, on the present view, this is compatible with some of them also having classical explanations in terms of flat-out beliefs and desires.

Let us look at a few objections to this account. First, doesn't one explanation undermine the other? If an action can be adequately explained by reference to the agent's partial beliefs and desires, doesn't that make it redundant to offer another explanation in terms of their flat-out ones? No – no more than the fact that an action can be explained in physical terms makes it redundant to offer an intentional explanation for it. The two are pitched at different levels. On the behavioural view, flat-out beliefs and desires are realized in partial ones and are causally effective in virtue of them. We might think of the different explanations as forming part of a layered framework of explanatory levels – physical, biological, intentional. Just as some events, as well as having physical and biological explanations, also have an intentional explanation in terms of partial beliefs and desires, so some have a second, higher-level intentional explanation in terms of flat-out beliefs and desires.

Secondly, is it not highly counter-intuitive to claim that some of our actions have two intentional explanations – and even more so to claim that one of these explanations cites beliefs about premising policies and their dictates? Surely the reason I change my order in the restaurant is simply that I believe that spring rolls are fatty and want to reduce my fat intake – not that I believe my premising policies *require* me to change my order? The claim that some actions have dual intentional explanations is novel, but not, I think, specially counter-intuitive. It is common to recognize the existence of two types of belief-desire explanation, classical and probabilistic; the novelty lies in suggesting that one underpins the other. And while it is certainly counter-intuitive to explain everyday actions by reference to meta-level beliefs and desires concerning premising policies, it is important to remember that the attitudes in question will typically be non-conscious. They will manifest themselves in the way we conduct our conscious reasoning, but they will not figure in it themselves as premises, and explanations citing them will be far from obvious. Moreover, when an action is performed in response to meta-level attitudes of this kind, there will always be a simpler and more intuitive explanation available citing the first-order flat-out beliefs and desires that featured in the agent's conscious reasoning. Since we do not expect actions to have dual intentional explanations, we shall therefore systematically overlook those explanations that cite partial attitudes concerning premising policies. It is not surprising, then, that we find the claim that such explanations exist counter-intuitive.

Thirdly, isn't the proposed view subject to the same objection as the full-belief view discussed earlier? The worry is that adopting a proposition as an unqualified premise is equivalent to treating it as certain, and therefore involves being willing to stake everything on its truth. And, as I pointed out, we can believe something flat-out without being willing to do this. My response here is to deny that adopting a proposition as a premise involves treating it as certain (or, in the case of a goal, treating it as maximally desirable). There will be situations where there is little to gain from acting on a given premise or goal and much to lose, and in such situations it would be prudent to refrain from doing so. For example, suppose I have accepted that the gun in my desk drawer is unloaded. And suppose I am offered a small sum of

money for taking the gun and, without checking the chamber, pointing it at the head of a loved one and pulling the trigger. In deciding whether or not to accept this offer, I might, quite reasonably, refrain from relying on the premise that the gun is unloaded. This presents no special difficulty for the present account, however. There is no need to suppose that in adopting a premise or goal we commit ourselves to acting on it come what may. In some cases the risks of adhering to a policy will outweigh the advantages, and we shall prefer to refrain and err on the side of caution. This would not necessarily mean abandoning the policy; we might continue to regard ourselves as committed to it and remain ready to act upon it in other situations where there is less at stake. Of course, it will be important not to make a habit of failing to act upon our premises and goals; doing so would erode our commitment to them. But occasional failures will not be too serious. Indeed, it is plausible to regard premising policies as having built-in exception clauses, excluding those deliberations where it would obviously be dangerous or absurd to rely on the premise or goal in question. What will typically happen in such cases, I suspect, is that we will make the reasons for caution explicit, qualifying or supplementing our premises and goals so that we can act upon them without undue risk.

Fourthly, what is the point of adopting premising policies? If I attach a sufficiently high probability to a proposition, then Bayesian theory will dictate I should, in most circumstances, act as if it were true. What do I gain by accepting it as a premise, too? The same goes for desire. What do we gain by adopting an outcome as a goal, in addition to attaching a high desirability to it? The short answer, I suggest, is that doing so affords us a measure of *personal control* over our thought processes. A lot of our mental life goes on below the surface. Much of our behaviour is the product of automatic, non-conscious mental processes of which we are unaware and over which we have no direct control. It is these processes that succumb to description in Bayesian terms (which is not to say that they involve actual Bayesian calculations, of course). Premising policies, on the other hand, are consciously formed and executed and are subject to personal supervision and control. By adopting premises and goals, and engaging in conscious reasoning, we can take manual control of our thought processes. We can decide what to think about, and when, and can direct our minds to theoretical problems with no immediate behavioural relevance. We can evaluate propositions and goals in the light of appropriate norms and choose which to endorse. And we can reflect on the inferential procedures we use, assess them, and enhance them. In effect, by forming and executing premising policies we create a highly flexible general-purpose reasoning system, whose procedures are open to continual refinement and extension.

To sum up, then: on this view flat-out beliefs and desires are premising policies, realized in non-conscious partial attitudes and effective in virtue of them. They form what amounts to a distinct level of mentality, which is conscious, reflective, and under personal control (elsewhere I have dubbed it the 'supermind'). The case for this view is essentially an inference to the best explanation: it offers the best account of how we could harbour attitudes with the profiles of flat-out belief and desire, and it does

justice to our intuition that it is not irrational to act upon such attitudes. I have argued elsewhere that this view also receives support from other sources – in particular, from reflection on the nature of occurrent thought, our ability to *make up* and *change* our minds, and the role of natural language in cognition. Again, many of our intuitions about these matters can be accounted for on the hypothesis that we are in the habit of forming and executing premising policies (see my 2004). The view also harmonizes well with recent work in the psychology of reasoning (see my 2009).

7. Flat-out belief and formal epistemology

My aims here have been descriptive. I have sketched accounts of partial belief and flat-out belief which reflect our everyday intuitions about these states and their roles, and have not offered rigorously defined notions that could be employed in a system for formalizing normative principles. But the descriptive task does have relevance for formal epistemology. We need to understand how our everyday doxastic concepts work in order to determine whether they, or refined versions of them, can be incorporated into a formal system and to see how to translate everyday claims into a formalized language. Investigation of these latter issues is a matter for another time, but I shall conclude this chapter with a brief comment on the utility of the everyday concept of flat-out belief for formal purposes.

I noted earlier that there is a strong case for thinking that rational flat-out belief is subject to conjunctive closure. And if it is a form of acceptance, as I have argued, then this will be the case. Accepting a proposition involves embarking on a policy of taking it as a premise in one's conscious reasoning, deductive and otherwise. And that involves being prepared to conjoin the proposition with others one has accepted. I suggested, however, that the concept of flat-out belief is that of a specific type of acceptance, defined by context. An acceptance counts as a flat-out belief, I suggested, only if its possessor is prepared to rely on it in deliberations where they attach a high desirability to taking only truths as premises ('TCP deliberations'). And, as I pointed out, this means that rational flat-out belief requires high confidence. From this perspective, then, rational flat-out belief is not subject to conjunctive closure, since rational high confidence is not preserved over conjunction. A rational agent might be prepared to rely on a series of propositions individually in TCP deliberations, but unwilling to rely on their conjunction in such deliberations.

If this is right, then the folk concept of flat-out belief is a mongrel one: that of a premising policy *coupled with* high confidence. And there is consequently a tension in it: *qua* premising policy flat-out belief is subject to conjunctive closure, *qua* state involving high confidence it isn't. Now this is not an objection to the descriptive account offered here – indeed it is a positive recommendation for it. For flat-out belief does seem to be subject to a tension of just this sort. We do feel both that we ought to adhere to conjunctive closure and that it is sometimes acceptable to violate it – for example, that it is not unreasonable to believe that one has some false beliefs, even though this claim is incompatible with belief in their conjunction. (The Lockean

thesis, by contrast, cannot explain this tension, since it offers no reason to think that flat-out belief should be subject to conjunctive closure in the first place.) The presence of the tension does mean, however, that the everyday concept is unlikely to be suitable for use in formal epistemology.

What sort of flat-out doxastic notion do we need for formal purposes? The question takes us beyond the scope of the present chapter, but the short answer, I suggest, is simply that of acceptance, in the sense described above, relativized to various purposes and contexts and without the requirement for high confidence. That is to say, the relevant attitude is a commitment to using a proposition as a premise in reasoning and decision-making for some purposes in some contexts, regardless of the degree of confidence we have in it. And, as with belief, this attitude will be motivated and sustained by the agent's partial beliefs and desires – epistemic and otherwise – and justified on decision-theoretical grounds. In theoretical contexts, for example, it might involve trading off informativeness against probability in the way Kaplan describes – some highly informative propositions being believable in this sense even though our confidence in them is low. From this perspective, then, flat-out belief may turn out to be a relatively uninteresting variant of a much more important attitude.

Acknowledgement

This chapter draws in part on material from my book *Mind and Supermind* (CUP, 2004) with thanks to Cambridge University Press. Some of this material has been substantially revised and rewritten. Thanks are also due to the editors of this volume for their comments and suggestions on earlier drafts of this chapter.

References

- Braithwaite, R. B. (1932-3). The nature of believing. *Proceedings of the Aristotelian Society*, 33, 129-146.
- Bratman, M. E. (1992). Practical reasoning and acceptance in a context. *Mind*, 101, 1-15.
- Chisholm, R. M. (1957). *Perceiving*. Ithaca: Cornell University Press.
- Churchman, C. W. (1956). Science and decision making. *Philosophy of Science*, 23, 248-249.
- Cohen, L. J. (1989). Belief and acceptance. *Mind*, 98, 367-389.
- Cohen, L. J. (1992). *An Essay on Belief and Acceptance*. Oxford: Oxford University Press.
- Davidson, D. (1975). Thought and talk. In S. Guttenplan (Ed.), *Mind and Language* (pp. 7-23). Oxford: Oxford University Press. Reprinted in Davidson 1984.
- Davidson, D. (1984). *Inquiries into Truth and Interpretation*. Oxford: Oxford University Press.
- de Sousa, R. B. (1971). How to give a piece of your mind: or, the logic of belief and assent. *Review of Metaphysics*, 25, 52-79.

- Dennett, D. C. (1981). True believers: the intentional strategy and why it works. In A. F. Heath (Ed.), *Scientific Explanation* (pp. 53-78). Oxford: Oxford University Press. Reprinted in Dennett 1987.
- Dennett, D. C. (1987). *The Intentional Stance*. Cambridge, Mass.: MIT Press.
- Engel, P. (1998). Belief, holding true, and accepting. *Philosophical Explorations*, 1, 140-151.
- Engel, P. (2000). Introduction: the varieties of belief and acceptance. In P. Engel (Ed.), *Believing and Accepting* (pp. 1-30). Dordrecht: Kluwer.
- Foley, R. (1993). *Working Without a Net*. Oxford: Oxford University Press.
- Frankish, K. (2004). *Mind and Supermind*. Cambridge: Cambridge University Press.
- Frankish, K. (2007). Deciding to believe again. *Mind*, 116, 523-547.
- Frankish, K. (2009). Systems and levels: dual-system theories and the personal-subpersonal distinction. In J. S. B. T. Evans and K. Frankish (Eds), *In Two Minds: Dual Processes and Beyond* (pp. 89-107). Oxford: Oxford University Press.
- Harman, G. (1986). *Change in View: Principles of Reasoning*. Cambridge, Mass.: MIT Press.
- Kaplan, M. (1981). Rational acceptance. *Philosophical Studies*, 40, 129-145.
- Kaplan, M. (1996). *Decision Theory as Philosophy*. Cambridge: Cambridge University Press.
- Maher, P. (1986). The irrelevance of belief to rational action. *Erkenntnis*, 24, 363-384.
- Maher, P. (1993). *Betting on Theories*. Cambridge: Cambridge University Press.
- Rudner, R. (1953). The scientist *qua* scientist makes value judgments. *Philosophy of Science*, 20, 1-6.
- Sellars, W. (1964). Induction as vindication. *Philosophy of Science*, 31, 197-231.
- Stalnaker, R. C. (1984). *Inquiry*. Cambridge, Mass.: MIT Press.
- Teller, P. (1980). Zealous acceptance. In L. J. Cohen & M. Hesse (Eds.), *Applications of Inductive Logic* (pp. 28-53). Oxford: Oxford University Press.