

Partial Belief, Partial Intention

Richard Holton

Does intending to do something entail that you believe you will succeed in doing it? Or can you form an intention lacking any such belief? The question is important, shedding light on the nature of intention and of its relation to other mental states. Unsurprisingly it has been the subject of much debate.¹ I shall have something to say about some of the recent contributions, but for the most part my aim in this paper is to take a rather different course. My contention is that, as it stands, the question is not well posed. It is couched in terms of all-out belief, with no mention of partial beliefs. Once these are brought into the picture, things look very different.

For a start things are different because new questions become available: not just 'Does intention entail belief?' but also 'Does intention entail partial belief?' More radically though, the thought that belief need not be all-or-nothing opens the door to the thought that same might be true of intentions: that there might be *partial intentions*, standing to all-out intentions much as partial beliefs stand to all-out beliefs.

I explore the notion of partial belief, introduce the idea of partial intention, and then explore the relation between intention and belief in the light of these accounts. My conclusion will be that there are very few entailments here: all-out intention does not entail all-out belief, and partial intention does not entail partial belief. This is not surprising: intentions and beliefs are different states, so we should not expect necessary connections between them. Instead of looking for entailments we should be looking for normative requirements: the ways in which partial intentions rationally require partial belief in success. In pursuing these issues I shall have to say something about the relation of partial belief to all-out belief; and about the relation of partial belief to credence, two notions that, surprisingly enough, I want to distinguish. But that is to get ahead of myself. Let me start by taking, and modifying, some examples from Michael Bratman.

Two Examples

(i) Last night's storm brought down a tree across your driveway, penning in your car, which you urgently need to use this evening. You are not sure whether you will be able to move the tree yourself, having never confronted something quite this big before. Three possibilities have occurred to you. You might lever it out of the way with a crowbar (though you're not sure how much force you can exert with that). You might saw it into manageable pieces

¹ For the recent debate see Velleman 2007 and Bratman forthcoming. For earlier pieces arguing that intention requires a corresponding belief see Hampshire and Hart 1958, Grice 1971, Harman 1976, and Velleman 1985; for argument that it does not see Bratman 1987 and plausibly, Anscombe 1963; see her examples, at p. 94, of St. Peter, and of the man who knows he is to be tortured. I discuss these below.

with a chainsaw (though you're not sure that you will be able to get the chainsaw started). Or you might put a rope round it and drag it out of the way with your car (though you're not sure, given the way that the car is penned in, that you will be able to manoeuvre it into a position from which this would work). Alternatively, and at considerable cost, you could get the local tree company to move it; but the storm brought down a lot of trees, so to stand any chance of getting them out in time you would certainly need to phone them first thing. In the end you phone the tree company, making a (cancellable) arrangement for them to come late in the afternoon; then you walk down to the shed and load up the wheelbarrow with your biggest crowbar; the chainsaw, oil, and fuel; and a rope: all in preparation for a morning attempt to move it yourself.

(ii) You have some library books that are badly overdue; in fact so badly overdue that your borrowing privileges are about to be suspended (a major inconvenience) if you do not return or renew them by the end of the day. Since you have finished with them, the best thing would be to drop them off at the library on your way home; but that is after the departmental seminar, and you know that, once you get on your bike with your head full of ideas from the discussion, you are all too likely to cycle straight home. Alternatively, you could renew them on-line; but that would require your library password, which is scribbled on a piece of paper lying somewhere on your desk at home. If you renewed them on-line you would not need to take the books home with you, but you would need to take your laptop, which you would otherwise leave at work. In the end you head for the seminar with your bag weighed down by both the library books *and* your laptop.²

I take it that in the first case you do not *believe* that you will move the tree on your own. But, contrary to what some authors have assumed, you are not wholly agnostic on the matter; clearly you think that there is some chance you will.³ Similarly in the second you do not *believe* that you will take the library books back, though again you think that there is some chance that you will. What should we say about your intentions?

Linguistic Intuitions & Their Limits

We can start, in time honored analytic fashion, by consulting our linguistic intuitions. Here, it seems to me, the two cases are rather different. In the first case we would naturally say, not that you intend to move the tree by yourself, but only that you intend

² Bratman 1987 pp. 38–9. Bratman's examples concern a log and a bookshop, not a tree and a library; but since I have modified the examples, I thought it best to modify their characteristic objects to keep them distinct.

³ Velleman writes as though in such cases, lacking all-out belief, the agent simply has no opinion about whether they will succeed or not; see Velleman 2007. I discuss this below.

to *try*; and that datum might be taken to lend support to the thesis that intention requires belief in success.

But what of the second case? We surely would not say that you intend to *try* to take the library books back; it is not as though the library is currently under siege, and you are sceptical about your abilities to get through the defences. No—if you get home and find to your annoyance that the books are still in your bag, it will not be true that you had even tried to take them back.⁴

So what do we say of such cases? One perfectly natural description here is that you intend to take the books back, but are unsure whether you will remember to do so. So long as we add the qualification, the direct statement of intention is quite acceptable.

If what I have said about the library case is right, this should make us reconsider the tree case, for the two are very similar. In each you can see different ways of achieving the same end; and in each, whilst you prefer one of these ways (moving the tree yourself, returning the books on the way home) you are unsure that you will be successful in it. (To make the cases closely parallel, we can assume that the probability that you assign to success is much the same in each.) So in each case, whilst making preparations for acting in the preferred way, you simultaneously make preparations to achieve the end in other ways. (Again, to keep the cases parallel, we can assume that your assessment of the chances of achieving the end in these other ways is the same in each.) But if the two cases are so similar, and we would naturally say that in the second you have an intention to return the book coupled with doubt that you will succeed, why should we restrict our description of the first to talk of trying? Why not say straight-out that you have an intention to move the tree yourself, coupled with doubt that you will succeed?

An obvious suggestion here is that the explanation is *pragmatic*: it is not that it would be false to say that you intend to move the tree yourself, but rather that it would be misleading to say so when there is something more informative that we could say, namely that you intend to try to move it. The difference between the two cases is just in what is available to be said. In the tree case there are actions that you are confident that you will perform, and these actions constitute trying to move the tree even in the absence of success in moving it; so these are the actions that you can mention. In contrast, in the library case there is no action of trying to take the books back that you are confident of performing but that falls short of actually taking them back; so there is no such alternative thing that we could say.⁵

⁴ Forgetfulness is not the only failing that can give rise to such cases. Weakness of will can generate them too. Suppose that you arrive at the library steps, and, seeing the glint of gold tooling on the books' smooth cloth-covered boards, you decide that, disciplinary measures notwithstanding, you will not give them up. An unlikely upshot, I admit, but possible. Then once again it wouldn't be true that you *tried* to take the books back; we are not supposing some strange compulsion forced you hand back as you endeavored to place them in the return chute. And once again this is a failure that you might foresee even as you form the intention to take the books back; you can form the intention in the knowledge that you might not succeed.

⁵ An alternative suggestion about what distinguishes the cases is this: only in the tree case are you doubtful of your *ability* to perform the action; in the library case you simply doubt that you *will* perform it. So it might be suggested that a necessary condition on intending to perform an action is confidence that one will be *able* to perform it. This strikes me as implausible. There are

This suggestion is very plausible, though it is a little complicated to say quite why the statement that you intend to try to move the tree is more informative than the statement that you intend to move it. Presumably it adds the information that you lack confidence that you will succeed; but how does it do this? The best explanation I can see draws a parallel between our statements about our own future actions and our statements about our intentions. I see you strolling purposefully towards the tree, crowbar in hand, and ask what you are doing. Suppose first that you do not make explicit reference to your *intentions*, but just talk about what you are *going to do*. You would be in no position to say that you were going to move the tree unless you were confident that you were going to be successful. Lacking that confidence you might rather say that you were going to *try* to move it; at least you could be sure of that.⁶ But then plausibly the same considerations come into play even when you do explicitly describe your intentions: rather than saying that you intend to move the tree, you say that you intend to try to move it, since this description of your intention conveys the information, by analogy with the earlier descriptions of your action, that you lack confidence that you will succeed.

As I said, I think that there is much plausibility to this pragmatic account, but it is clearly not conclusive. The main lesson that I want to draw is negative: when it comes to cases like these it is not obvious what we should say. Certainly there are no compelling grounds for saying that we lack intentions to do anything more than try; and there are some grounds for saying that we have intentions, provided that we add a rider to indicate our lack of confidence in success. What we need then, if we are to make any progress, is not just an examination of intuitions regarding these cases, but a developed account of them.

Such an account, I suggest, should contain two elements. It will need to say something about the attitude, an attitude falling somewhat short of all-out belief, that we have to the proposition that we will succeed. And it will need to say something about the attitude, an attitude falling somewhat short of all-out intention, that we have towards the action that we might perform. The first topic has been much discussed in philosophy, under the heading of *partial belief*. I want, however, to understand it in a novel way, and it is with this that I shall start. The account of partial belief will set things up for the second notion, that, to maintain the parallel, I shall call *partial intention*.⁷

alternative versions of the library case in which you would be unable to return the book, but in which I think our linguistic intuitions would be unchanged. Suppose that exciting seminars have a tendency to provoke epileptic seizures in you; if you were to suffer such a seizure you would be unable to return the book the same day. Couldn't we still say that you intend to return the book, but are uncertain whether you will do so, even though in this case you are uncertain of your ability?

⁶ Lacking confidence, one does not simply *say* that one will try; one actually *thinks* of one's action in a different way, typically by focusing on the proximate mechanics of the movement rather than on the final goal. For research on our tendency to focus in this way in the face of difficulty, see Vallacher and Wegner 1987.

⁷ Terminological point: Bratman also talks of intentions being partial, but he means by this the very different idea that their details are not filled in. I hope that my very different usage does not lead to confusion. I considered using another term (Graded intention? Uncertain intention?) but

The Surprising Absence of Credence

We have assumed that you lack the all-out beliefs that you will be successful in moving the tree or in dropping off the books. But in assuming this we haven't assumed that you have absolutely no clue about things will work out. On the contrary: you think that there is some chance that you will succeed, and some that you will not. Indeed, in order to make the cases parallel I added the assumption that you give equal weight to the chance of success in each case. The standard way of understanding such attitudes is as *credences*, attitudes assigning a value between one and zero to each proposition in ways that, normatively at least, conform to the axioms of the probability calculus (so that, for instance, the credence assigned to p should be equal to one minus the credence assigned to not- p).

The idea of using credences is so well-established that we might expect other authors writing about intentions to have framed their discussion in terms of them; yet they have not. We should pause to ask why. It is not an oversight. Admittedly, some of the early discussion of intention took place before the notion of credence had become so familiar; perhaps it did not occur to the initial authors to work with anything other than all-out belief. But in more recent work the restriction has been deliberate. The terms of the recent debate were largely set by Bratman's highly influential book *Intention, Plans, and Practical Reason*; and there Bratman explicitly rejected any discussion of credences, in favour of all-out belief (Bratman 1987, pp. 36ff).

Although Bratman does not say much about his motivation for this, it is easy to see the kinds of considerations that moved him. First, it looks as though, as a matter of descriptive fact, we typically work with all-out beliefs rather than credences. Certainly when we make ordinary common sense attributions of cognitive attitudes to ourselves and to others, what we attribute are all-out beliefs. Of course this may just be a mistake: it could be that we are working with a crude ungraded notion in our folk theory, when the reality is graded. But it would be surprising if our ordinary notions were so far adrift from what we actually experience.

More specifically, there is good reason for thinking that we do not normally ascribe specific numerical credences to propositions. Certainly the psychological findings confirm that we are far happier making qualitative probability judgements ('I think it pretty unlikely that p ' 'There is a good chance that q ') than giving our judgments numerical values, even vague ones (Budescu and Wallsten 1995). That is unsurprising. Maintaining large numbers of credences leads to enormous complexity, especially if we always update by conditionalization, since this requires us to have a prior conditional probability for every pair of propositions.⁸ And there is widely known evidence that

none seemed right, and the parallel with the well entrenched notion of partial belief was too strong for me to resist.

⁸ For this criticism see Harman 1988, pp. 25–7. Put as I have put it, the complaint is overstated. Rather than assigning prior conditional probabilities to all pairs of propositions we might use Bayesian Networks, representations that implicitly assume that many factors are independent, and

when we do try to make probabilistic calculations we often make terrible mistakes, ignoring base rates, reacting differently to the same information presented in different ways, and so on.⁹

This is not to deny that there are areas of action and belief formation in which our behaviour seems to be well described by assuming that we act as Bayesian agents that form credences and update by conditionalization; and the same is true of other animals.¹⁰ But such areas typically involve unconscious judgments, or unconscious manipulation of conscious judgments, about the kinds of things for which we might well expect specific competences to have evolved, such as causal dependencies. Such competences no more suggest that we are able to engage consciously in Bayesian reasoning than the fact that some aspects of vision is well modeled by Fourier series shows that creatures who can see are competent with harmonic analysis.¹¹ However much we might want to consciously entertain sets of credences that conform to the axioms of the probability calculus, and manipulate them by methods like conditionalization, it looks as though we simply can't do it, any more than wanting to swim like a dolphin enables us to do that. And the situation doesn't change if we understand what would be involved in entertaining sets of credences, and can even perform some of the appropriate manipulations in isolation, just as understanding the biomechanics of dolphin swimming, and even learning to ripple parts of one's body in the appropriate manner, doesn't enable one to swim like a dolphin.

A second problem concerns the issue of how to relate credence to all-out belief. A natural first thought is that all-out belief is just high credence. But, for familiar reasons, such an approach is hopeless. How high does an agent's credence need to be before we ascribe them an all-out belief? We might require complete certainty, but then almost all of our attitudes to contingent propositions (and to some necessary propositions) would not count as belief. Alternatively we might try to establish a threshold which credence must pass to count as belief: 0.9 say. That does no better. One problem is that any particular threshold will seem arbitrary. Another is that we normally expect belief to agglomerate: if an agent believes *A* and believes *B*, then they will believe *A and B*. But if credences obey the rules of the probability calculus then we will not in general get this result. If *A* and *B* are independent beliefs that are just on the threshold of 0.9, then the agent's credence in *A and B* will be the product of the two individual credences: in this case 0.81, which will fall below the threshold, and so not count as belief.¹² A related problem, the Lottery Paradox, confirms that the threshold view and a rule of agglomeration do not fit together. Suppose that there are 1000 tickets in a fair lottery.

so drastically reduce the number of priors needed. For an introduction to Bayesian Networks, with clear discussion of this point, see Charniak, 1991. Nevertheless, the computational complexity remains outside our conscious ability.

⁹ The classic study here is Kahneman & Tversky, 1973; there has been much more since. See the articles collected in Kahneman & Tversky, 2000.

¹⁰ For review see Chater *et al.*, 2006, and the articles that follow it in the same journal.

¹¹ The comparison is from Chater *et al* 2006, p. 288.

¹² For these two worries see Stalnaker, 1984, pp. 91–2

We reasonably assign a credence of .999 to any given ticket that it will not win. Since such a credence takes us over the threshold for belief, it follows, on the threshold view, that we believe of each ticket that it will not win. In itself that sounds wrong. Add an agglomeration rule and it would follow that we would believe that no ticket will win. Yet we don't surely don't believe that: we know that one of them will win.

Of course problems in relating credence to all-out belief do not by themselves tell against credences; in fact they have more often been used by proponents of a credence account for explaining why they should have no truck with all-out belief. But if we accept that we do have all-out beliefs, then these considerations provide further reason for not just adding credences in as well.

So there are good grounds for being suspicious that real agents work with credences. But they should not lead us to try to work exclusively with all-out belief. Clearly the tree and library examples can only be handled by introducing some attitude that falls short of all-out belief. The most that we can conclude is that this attitude is not well modeled by the classical credence account. I think that in effect Bratman concedes this. Bratman thinks that one can rationally intend to F iff *either* one believes that one will F, *or* one neither believes nor disbelieves that one will F. All that is ruled out is that one disbelieves that one will F (Bratman 1987, p. 38). In proposing this I think that he is not just working with two attitudes, believing that p and disbelieving that p. He is implicitly working with a third: neither-believing-nor-disbelieving that p. This is not simply the state that agents are in when they do not believe that p and do not disbelieve that p. One could be in that state as a result of never having raised the question whether p. Yet if disbelieving that one will F is incompatible with intending to F, then surely failing even to have considered whether one will F is incompatible with intending it. Rather the attitude of neither-believing-nor-disbelieving that p looks to be a distinct attitude that one can take to p, an attitude that assigns it a degree of belief less than that assigned by believing it, and greater than that assigned by disbelieving it. But once we admit three attitudes we no longer have an all-out account of belief; we have a very coarse grained version of a graded account.

Belief, Partial Belief, and Credence

So can we give an account of partial belief that is psychologically more realistic than that provided by the credence account, and that can be integrated into an account of all-out belief? I think that we can. In this section I want to sketch the outlines of such an account, and to say something about how it relates to the idea of credence.

We saw how hard it is to arrive at all-out belief by starting with credences: the threshold account is clearly hopeless. So let us try to move in the other direction. Let us start by getting clearer on the nature of all-out belief, which we can take to be primitive, and then inquire how it might be made more partial.

A number of recent authors have argued that the central feature of all-out belief concerns its relation to practical deliberation. For instance, in distinguishing belief from credence, Timothy Williamson, has suggested that one all-out believes a proposition just in case one is willing to use it as a premise in practical deliberation

(Williamson, 2000, p. 99). The proposal echoes some earlier comments of Bratman's who, in discussing his original log case, wrote:

To believe something is not merely to assign a high probability to its occurrence. I might assign a high probability to my failing to move the log without believing that I will fail ... what seems distinctive about believing that I will fail is that it puts me in a position to plan on the assumption of failure.¹³

As it stands the proposal is over restrictive. Mightn't I have all-out beliefs about certain things (things in the distant past, for instance) where the question of action does not arise? And, conversely, seeing no other possibility of getting what I want, mightn't I act on a supposition even though I do not fully believe it? Moreover, all-out belief surely has connections with other attitudes that are independent of action. To take a nice example from David Owens: my feelings of guilt, or pride, or anger only really come in when I all-out believe that I have done badly, or well, or have been badly done to (Owens, forthcoming).

Nevertheless, I think that the central idea, that all-out belief is at heart a practical attitude, is right.¹⁴ We are cognitively limited creatures. Maintaining and manipulating large numbers of credences would overload our capacities, and our reasoning would get nowhere. So as a practical matter it makes sense to accept certain premises as given, and to do our reasoning on the basis of them, even though we acknowledge that there is some chance that they will be wrong. In making my plans this morning I took it as a premise that the philosophy department would be functioning, that my route to it would be open, that my computer would not have been stolen, and so on; and this is so even though I would give to each of these propositions a credence of less than one. I even took it as a premise that the lifts would be working, pressing the button and waiting patiently for one to come, though the credence I would assign to that proposition is quite considerably less than one.

On this picture, the role of all-out beliefs is in many ways parallel to the role of intentions. Just as intentions enable us to resolve deliberative uncertainty in order to facilitate action, so all-out beliefs enable us to resolve epistemic uncertainty in order to facilitate action. They allow us to reduce an unmanageable amount of information to a manageable amount by excluding certain possibilities from our practical reasoning. They provide us with a relatively simple description of what the world is like, to which straightforward non-probabilistic reasoning can be applied, around which plans can easily be built, and which can easily be communicated to others.¹⁵

¹³ Bratman 1987, p. 40. Elsewhere Bratman is keen to distinguish the idea of acceptance, which is what is needed for reliance, from belief; see Bratman, 1990. For other interesting accounts of all-out belief that tie it to the idea of action see Frankish, forthcoming and Weatherson, 2005. Weatherson is also adamant that all-out belief is not a distinct state from partial belief, but that the distinction stems only from our different descriptions of it. I am inclined to think that the distinction is psychologically real; see below.

¹⁴ Here I am especially indebted to conversations with Agustin Rayo.

¹⁵ There is a further way in which beliefs are similar to intentions, in that they both have a certain inherent stability. This can be best understood, I think, in terms of differing thresholds for formation and for revision. Some considerations that would have been enough to stop me forming

Note how this approach differs from that suggested by Williamson and Bratman. Deliberating with all-out beliefs is an attitude dictated by practical considerations. But the deliberation that we do is not then limited to practical matters. I can deliberate about what happened in the distant past, or what will happen in the distant future, or about what I would have done if things were very different to how they are, even though none of this will affect my behaviour. The approach offered here thus plausibly entails that if someone believes a proposition they will use it in their practical deliberation; but it does not identify belief with such a use.¹⁶

So it makes sense to deliberate with all-out beliefs wherever possible. But of course we cannot always do so. Sometimes it is clear to us that we do not know what did or will happen; this is true of the tree case and of the library case. In such situations we have to make use of a more partial notion. But this does not take us all the way to credences, for two reasons. First, as we have seen, we do not normally assign anything like numerical values to the possibilities.¹⁷ We just talk of one outcome being likely or unlikely, or being more likely than, or equally likely as, another. Sometimes we do not even this, but simply think of a number of possibilities as being live, without trying to rank them.

Second, and more importantly, even when we do assign non-zero credences to a range of propositions, we need not accept each of them as what we might call a *live possibility*. In the tree case I attributed to you a partial belief that you would succeed in moving the tree using one of three methods, and a further belief that you would fail with all three and that the tree company would move it. But unless your view was remarkably blinkered, this would not exhaust your credences; or at least, it would not exhaust the credences that you would be inclined to entertain, given the right prompting. You would surely give a non-zero credence to the tree being moved by a helpful gang of passing weightlifters, or to it being blown out of the way by a freak local tornado; and you would give a much higher credence to the tree not being moved at all, either because both you and the tree company tried and failed, or, more likely still, because you failed

an intention will not be enough to lead me to revise it once it has been formed. Similarly, some evidence that would have been enough to stop me forming a belief will not be enough to get me to revise it once it has been formed. Though obviously such tendencies can lead to irrationality, I think that it is quite rational to have them to some degree. I suggest that they work by placing different thresholds for initial consideration and for reconsideration. For discussion in the case of intention see Holton 2004.

¹⁶ The proposal has a similar subsuming relation to various other proposals in the literature. Thus, for instance, Mark Kaplan has argued that the defining characteristic of all-out belief in P is, roughly, the agent's preparedness to assert that P if they are restricted to asserting either P or \sim P. See Kaplan, 1996, Ch. 4 for details. This feature plausibly follows from the account proposed here, but I do not think that it captures the essence of belief.

¹⁷ Exception: cases in which the uncertainty is in not primarily epistemic, but can be thought of as deriving from a chance that is in some sense objective: the chance of rolling a six on a fair die for instance. Here people are quite prepared to assign numerical probabilities. But exactly because of this objective construal, it is plausible to think that the probability assigned is part of the *content* of an all-out belief, rather than a measure of the *attitude*: one all-out believes that there is a one sixth chance of rolling a six. Certainly the standard linguistic formulations suggest this. No one who is not a theorist talks of having a one sixth degree of belief.

and the tree company did not show up. So on reasonable assumptions, your credences over the four options would not sum to one. Nevertheless, it is a very normal case in which you narrow your focus to these four, as we have assumed that you do. These are the only four upon which you base your reasoning and your plans.¹⁸ You do not consider what to do if the tree is not moved. Or at least, if you did initially consider this possibility, you do not continue to do so, and you certainly do not make plans around it. It is not live. You thus do not make plans around the possibility that you will be unable to move the car: you do not telephone those you should meet this evening to warn them that you might not make it, do not ask to see if someone can pick you up, do not start finding out about taxi services. So there is something special about the four options that you believe partition the ways in which the tree will be moved. You have an all-out belief in their disjunction, even though your credence in that disjunction remains less than one. It follows that the attitude you take to each of the disjuncts is not simply credence.¹⁹

I conclude then that there is a practical stance that one can take towards a proposition that is rather like all-out belief, but is partial. We might call it *partial all-out belief*; but, since that has an air of oxymoron, I shall somewhat stipulatively call it *partial belief*.²⁰ Although I have been keen to emphasize the ways in which partial belief differs from credence, there is an important way in which they are alike. Like credence, partial belief involves a distinctive (partial) attitude to a normal content, rather than a normal attitude to a distinctive (partial) content (one partially believes that p, rather than all-out believing that p is probable). Given that partial belief and credence share this structural feature, we might ask whether credence can be accommodated within the account. I suggest that we think of credences as the partial beliefs that an agent would have if they were quite unconstrained by cognitive limitations; or, more plausibly, as the partial beliefs that they do have when their cognitive limitations are irrelevant. One obvious circumstance in which these limitations are largely irrelevant is in simple betting behaviour. When you are asked the odds that you would give to a certain outcome—the odds that a gang of passing weightlifters will move the tree—this doesn't require much from you. It doesn't require you to formulate complex plans contingent on that outcome, since it doesn't require further action on your part; if you win the bet the

¹⁸ It is an interesting question, one that will not be properly pursued here, how we determine which possibilities are live and which are not. We certainly need not think that this is under the immediate control of the will. An obvious place to start would rather be by looking at the rules that contextualist accounts of knowledge use to determine relevant epistemic possibilities. See, for instance, Lewis, 1996. One plausible difference concerns Lewis's rule of attention, by which the mere mention of a possibility makes it epistemically relevant; though it is debatable whether this rule is correct even for the case of knowledge.

¹⁹ It remains open whether or not even in this case our inability to manipulate these states in the way that credences should is so great that they do not deserve to be called credences; even here we could still be trying to be dolphins.

²⁰ Williamson discuss the idea that all-out belief comes in degrees whilst remaining distinct from credences, but he means to measure something rather different by it, namely how readily one would abandon one's reliance if the stakes were to go up (Williamson 2000, p. 99).

winnings will be delivered to you.²¹ So in such circumstances there is typically no need to exclude the possibility from your reasoning: you just come up with an answer. It typically isn't a very good one (we tend, for instance, to overestimate small possibilities and underestimate large ones), and the kind of work that would be needed to get it better would require serious work; but nonetheless we are unconcerned about saying something. So we can concede to the orthodox approach that betting behaviour provides the touchstone of credence, whilst insisting that it is really a very particular phenomenon, one provides a bad model for ascertaining our cognitive attitudes more generally.

The model that I have been sketching thus sees all-out belief as the norm, with partial belief involving as limited a divergence as possible. Partial beliefs are generated in response to cases where it is clear that we lack knowledge, cases where we recognize that more than one possibility is a live option. Credences are generated in even more special circumstances. We do not start with credences and then build our beliefs out of them.²² And even when we do have a set of credences, the relation between credences and beliefs are complex. It is the role of the good insurance salesman to generate credences; and then to move them from being mere credences to being partial beliefs, i.e. attitudes on which we will act by spending money on an insurance policy. Even a good salesman will be more effective on some potential customers than others. If beliefs are to play the role that I have assigned them without leading us into disaster, then, supposing we do have a credence in a proposition, it must be sufficiently high before we treat it as a live possibility, and even higher before we treat it as the only live possibility. However, there is no particular threshold that is required. It will often happen that people will differ in their attitude to a proposition even when they assign it the same credence. Some (the self-confident? the foolhardy?) will all-out believe it—i.e. will treat it as the only live possibility—while others (the cautious? the paranoid?) will not. Moreover, there are factors other than just level of credence that will be relevant in whether one forms a belief. Obviously what is at stake will matter, so that as the cost of false belief rise, the tendency to believe will fall. But there are other less obvious factors. Since there is no particular credence threshold, the rule of agglomeration may be retained. But since all-out beliefs plausibly must be consistent, we do not take the attitude of all-out belief where it is clear that agglomeration will lead to inconsistency.

An aside: This last consideration provides, I suggest, our best hope in describing what happens in the lottery paradox. When we know that these are lottery tickets, and so do tend to assign credences to each proposition that a particular ticket will win, and hence notice that an inconsistent result will follow if we all-out believe of each that it will not win, we tend not form such an all-out belief for any of them, no matter how

²¹ Things would be very different if you were required to ascertain whether you had won, and then to claim the winnings yourself. For the extra cognitive load that this would impose would make it rational to avoid getting involved in much betting behaviour.

²² For an account sympathetic to this approach, one that 'works with no underlying psychoeconomic scales', see Stewart *et al.*, 2005.

small the odds.²³ I mean this simply as a descriptive claim about our belief forming practice. Of course it entails that people will frequently all-out believe propositions in which they have a lower credence than propositions that they do not all-out believe. Further, people often fail to recognize when the beliefs they hold will lead to inconsistency. I only claim that once they come to recognize that they will lead to inconsistency, or merely come to have the false belief that they will, they tend to refrain from all-out belief.

Partial Intention

Return once more to the tree example. Let us suppose that you acknowledge four live possibilities: that you will move it with the crowbar (which you will try first); that you will fail with the crowbar and manage with the chainsaw (which you will try next); that you will fail with the chainsaw and manage with the car (which you will try last); and that you will give up and let the tree company move it. And let us say that you think that success in each of the first three possibilities is rather unlikely, but that overall you think you are about as likely to move it yourself as to have to wait for the tree company to move it.

You have then a set of partial beliefs; but that is enough to start planning and acting on them. That is what you do when you phone the company and fill up your wheelbarrow with the various tools. In so doing you exhibit various attitudes that are intention-like. It is not clear that we would ordinarily call them 'intentions'; but equally we would not ordinarily call your partial beliefs 'beliefs'. The attitudes are certainly like all-out intentions in many respects. They play the same roles of curtailing deliberation, resolving indeterminacy, and enabling coordination that intentions play: you fix on a small number of plans from the many that occurred to you and that you might have pursued, and as a result of this you are coordinate with your other plans (the plan to go out in the evening), and with other people (the people at the tree company). What distinguishes the states you are in from normal intentions is simply that they are partial: they stand to all-out intentions much as partial beliefs stand to all-out beliefs. You do not all-out intend to move the tree by means of the crowbar; but you do *partially intend* to do so. Moving it with the crowbar is one component—a subplan—of your larger all-out intention of moving it before the evening.

There are two ways to understand the idea of a partial intention. We might say that an agent has a partial intention whenever they have merely a partial belief in its success. Or we might say that it is essential to partial intentions that they be only a proper part of an overall plan i.e. that they be accompanied by alternative partial intentions to achieve the same end. In the partial beliefs of rational agents the analogues of these two different features come together: if I have a merely partial belief in *p* I will automatically

²³ It helps, of course, that the tickets have an equal chance of winning, so that we have no grounds for restricting our beliefs to any proper subset.

have an accompanying partial belief in its complement $\sim p$.²⁴ With intentions the result does not come automatically. If I have only a partial belief that I will achieve my end by succeeding in a certain intention, it does not automatically follow that I have an alternative intention designed to achieve the same end. I might have no backup plan: I might simply have a partial belief that I will achieve the end at all.

So we need to make a choice of how to define partial intentions, in terms of partial belief in success or of the presence of alternative intentions to the same end. I take the second path:

Partial Intention

An intention to F is partial iff it is designed to achieve a given end E and it is accompanied by one or more alternative intentions also designed to achieve E. If an intention is not partial it is all-out.

After all, it does seem that if something is partial there should be other parts that make up the whole; that is the idea the definition captures. As defined here, the crucial aspect of partial intentions is that, like partial beliefs, they involve a division in the agent's consideration. An agent with partial beliefs considers two or more competing live possibilities. An agent with partial intentions is working with two or more competing plans.

Note though that since all-out intentions are in turn defined as those that are not partial, the lack of an alternative intention is enough to qualify an intention as all-out. The definition thus leaves open the possibility that there may be all-out intentions in whose success the agent has only a partial belief; in other words it does not follow from the definition that all-out intention entails all-out belief. Of course there may be other reasons for believing in that entailment; we shall examine some later. But it does not follow by fiat.

It might look as though this definition makes the partiality of an intention sensitive to the way in which the end is described. The worry is that if my end when wielding the crowbar is to *move the tree*, then clearly the intention is partial, since I have other intentions aiming for the same end; but if my end is rather to *move the tree with the crowbar*, then it is not partial. So let us be clear: given the definition, an intention is partial just in case there is *some* description relative to which it is partial. The intention to move the tree with the crowbar is thus partial *simpliciter*. This raises a worry from the other direction. Mightn't it be the case that, at some level, all of our intentions are directed at the same end: our *eudaimonia* for instance? The worry is misplaced. Even if all of our intentions were ultimately directed at our *eudaimonia*, it is very unlikely that they would represent *alternative* ways of achieving it as the definition requires; rather they are *contributory*. One intention is an *alternative* to another in achieving an end just in case, were it successful, there would be no possibility of using the other to achieve the end.

²⁴ The point is even clearer on the traditional account of credences: if I have a credence of less than one in p then, assuming my credences conform to the axioms of the probability calculus, I will automatically have an accompanying credence of greater than zero in $\sim p$. Credences come in bunches that sum to one.

My talk of partial intentions may be misleading in one respect, so let me flag it. A partial intention need not be half-hearted. When you set to work with the crow bar, it does not follow that you push more gently than you could, on the grounds that your intention is only partial. You may: you may want to hold some energy in reserve for the other options, or, knowing that you have them, you may decide not to risk your back. But that is not essential to the endeavour. Although your intention is, by my lights, partial, when you are executing it you may do so with everything at your disposal. A partial intention is only partial because of the presence of alternatives. Does this bring a disanalogy with partial belief? It is hard to know. It may seem that, in contrast to this whole-heartedness, the state of partially believing requires one's belief to be half-hearted. But I doubt that that is an apt parallel. It is the *action* that results from the partial intention that can be whole-hearted, and this is equally true of an action that results from a partial belief. My confidence that the chainsaw will start may be very low, but that doesn't mean I pull the start cord with any less vigour.

To sum up this section: I propose then that we should admit a notion of partial intentions, standing to all-out intentions much as partial beliefs stand to all-out beliefs. Whether we call them 'intentions' or not is of little importance: after all, ordinary usage is unclear on the parallel issue of whether partial-beliefs should be classed as a form of belief. In the light of the earlier considerations, I suggest that much of the reluctance we have to call them intentions is merely pragmatic, triggered by the availability of mentioning instead a distinct action that is a trying. Where no such action is available (as in the library case) we have less reluctance. But I am not too concerned if I am wrong about that; ordinary language is not an infallible guide to natural kinds.

Do We Need Partial Intentions?

Perhaps though we could live without partial intentions; perhaps we could account for the kinds of cases that I have discussed with just the normal framework of all-out intentions. I have already explained, by means of the library example, why we cannot in general think of them as all-out intentions to try. We might instead try to understand them as *compound* all-out intentions. Presumably they are not conjunctive: you do not all-out intend to move the tree by yourself *and* have it moved by the tree company, since you know that that is impossible. But mightn't they be disjunctive: couldn't we simply say that you all-out intend *either* to move the tree yourself, *or* have it moved by the tree company?

Of course we could say that; but to say only that would be to lose explanatory force. For we need to break compound intentions down into their elements if we are to understand quite what explains what. Consider a parallel example with ordinary all-out intentions. Here presumably conjunction is permissible: if I intend to hear a concert and intend to buy some whisky, then I intend to hear a concert and buy some whisky. But we wouldn't want to be constrained to use only the conjunctive sentence. It is my intention to hear the concert that explains why I buy a ticket; it is my intention to buy some whisky that explains why I divert to the off-licence. It is only if we break the conjunction down into its consistent atoms that these explanations become available.

The same is true when we try to give all-out disjunctive surrogates for partial intentions. It is my partial intention to get the tree company to move the tree that causes me to phone them; if we are limited just to all-out disjunctive intentions, we can give no explanation of this.

A more plausible approach uses *conditional* all-out intentions: you all-out intend, *if* your attempt with the crowbar fails, to cut up the tree with the chainsaw; and so on. This avoids the problem raised for the disjunctive approach, since in place of each partial intention we get a separate conditional intention. Nonetheless, I do not believe that the proposal will work. Of course it is true in a trivial sense that you intend each subplan to work if the others fail; to that extent your partial intentions are conditional. That, however, is just to say that they are partial intentions; it does not explain how they are to be characterized as a class of all-out conditional intentions. What gives some support to this latter claim is the idea, made explicitly in the tree case, that each of the later stages in your plan will be triggered by the failure of earlier stages. But that does not provide a general recipe for reducing partial intentions to conditional intentions. First it is not true of the first stage: moving the tree with the crowbar is a partial intention but not a conditional one. Nor, in the library example, is your intention to return the books on the way home a conditional intention. But we want to register that these are not normal all-out intentions: they are already enmeshed in further intentions that would have no place if they were. Second, one may formulate a set of partial intentions without yet deciding how to integrate them; perhaps when you place the tools in the wheelbarrow you have not yet decided which to employ first. So there is not yet a set of conditional intentions in place. Of course in this case there is a compound intention: to start with either the crowbar or the chainsaw or the rope, and if that fails, to proceed to one of the remaining methods, and so on. But we have already seen that compound intentions of this form cannot do the work that is needed of them. Third, whilst it is true that the subplans in the tree moving example and the library example are basically sequential—I go through one subplan after another—not all cases will be like this. Again we can use an example from Bratman to illustrate the point.²⁵ I am faced with two identical video games. Success in either involves hitting a target with an electronic missile. This is a difficult feat that will probably require many attempts. Ambidextrous as I am, I decide to start shooting in both games at once, one per hand, despite my knowledge that the games are linked, so that were I to be about to hit both targets simultaneously (a chance I rate as very slim), both would shut down.²⁶ In this case I do not intend to hit both targets; I know that to be impossible. Rather, I partially intend to hit a target with each missile. But these are not conditional intentions; I do not fire the missiles sequentially, but simultaneously.

²⁵ Bratman 1987, pp. 114ff. Bratman denies that there are intentions involved here, for reasons that I discuss below.

²⁶ Perhaps being ambidextrous is not enough; perhaps I would need independent eyes if I were to aim each of the guns independently. Still, that is possible, though it would probably be best achieved by severing my corpus callosum (compare the case of Kim Peek, the model for the central character in the film *Rain Man*, who was born without one, and can reputedly read the left and right pages of a book simultaneously).

So all-out intentions to try will not do all of the work required of partial intentions, nor will all-out disjunctive intentions, nor will all-out conditional intentions. But I have responded to each with a different counter-example. So a worry remains that a patchwork of different strategies might do the work: replace some of the partial intentions with intentions to try, some with disjunctive intentions, and some with conditional intentions. Such a response certainly seems inelegant; but can it be refuted? I suggest that the initial library example provides a refutation. I have already argued that that cannot be seen as an intention to try. But nor can it be seen as simply a disjunctive intention to either return or renew the books (I need separate explanations of why I take the books with me, and why I take the laptop). Finally it is not a conditional intention: since it is the first in the series, it will not be triggered by the failure of some other intention. I conclude then that we need a separate notion of partial intentions.

Consistency Worries

Michael Bratman rejects the idea that there are partial intentions. But this is not on the grounds that they are unnecessary. Rather, he holds that they infringe a plausible consistency requirement on intention. My main aim in this section is to see how good the argument is.

There is a minimal consistency requirement that everyone can concede, so let us start by getting it out of the way. It should never be the case that your total set of intentions and partial intentions puts you in a practically impossible situation: it should never require you to do two inconsistent things (or, at least, two things that you believe inconsistent) at the same time. Call this the *execution requirement*. Though the intuitive idea is clear, the requirement is rather hard to formulate. Sometimes a perfectly rational plan will require you to try to do two things that are in a sense inconsistent at the same time: that is just what happens in the video game case. The point is rather that the demands on your immediate, proximate actions should be consistent: you should not have to stand up and sit down at the same time, to shoot the gun and not shoot the gun, to turn simultaneously to the left and the right, and so on. What is consistent here is a highly contingent matter: perhaps an intention to rub my stomach and pat my head is consistent for me, perhaps not; perhaps the only way to find out is to try.

The execution requirement poses no threat to the existence of partial intentions. But this is not the kind of requirement that interests Bratman. Rather, he is interested in the consistency of our intentions with our beliefs in their success. A weak requirement of this kind is:

Weak consistency

If an agent forms an intention, then they must not believe that they will fail to realize that intention.²⁷

²⁷ Bratman endorses such a principle (Bratman 1987 p. 38).

This is an *atomic* requirement: it places a constraint on each intention individually. We might want something stronger, a *holistic* consistency requirement that places a constraint on the agent's intentions and beliefs all taken together. The obvious candidate is the constraint that Bratman endorses, namely:

Strong consistency

If an agent forms an intention, then the realization of that intention must be consistent with their beliefs and with the realization of all their other intentions.²⁸

Clearly the weak consistency requirement follows from the strong, since if an agent believes that they will fail to realize their intention, then the realization of that intention is inconsistent with their beliefs. The strong requires, in addition, a holistic coherence. For those who endorse the intention-entails-belief thesis, the strong consistency follows from the requirement of consistency on beliefs. But for those, like Bratman, who deny this thesis, it is an independent constraint.

Strong consistency leads Bratman to deny partial intentions, or states like them, the status of intentions. Consider, for instance, the mental states involved in the video game example described above, where I simultaneously shoot at two targets, one with each hand, knowing that I cannot hit both (or the games would shut down). It follows, by both weak and strong consistency, that I cannot intend to hit both targets with both missiles for I believe that this is impossible. Further, since neither missile is privileged, any intention I have regarding one I must have regarding the other. So, by strong consistency, it follows that I cannot intend to hit either, for if I intended to hit one, I would also intend to hit the other, and the realization of both intentions is not compatible with my beliefs. In consequence, following Chisholm, Bratman classes them as mere *endeavourings*, states that resemble intention in being action guiding, but that are otherwise very different (Bratman 1987 pp. 129ff.).

I think that Bratman's response here is mistaken. As we have seen partial intentions are like intentions in that they enable agents to coordinate, to curtail deliberation, and to resolve indecision. Thus the only ground for denying them the status of intentions is that they fail the strong consistency requirement. Yet whilst simple consistency is a plausible constraint on all-out intention, as it is on all-out belief, it is not a plausible constraint on partial intention. Rather, the consistency constraint that we place on our partial intentions should be like the one that we place on our partial beliefs; otherwise we are not treating like with like. We do not require that partial beliefs be consistent, in the sense that everything that we partially believe must be compossible; so we should not require this of our partial intentions. Instead, we should require of our partial intentions at most the same kind of consistency that we require of our partial beliefs.

So how can we reformulate analogues to the consistency conditions on partial beliefs? A natural reformulation of the weaker requirement is:

²⁸ I.e. the proposition describing the realization of that intention must be consistent with the propositions that form the content of their beliefs, and the propositions that describe the realization of their other intentions (Bratman 1987, p. 31.)

Very weak consistency for partial intentions

If an agent forms a partial intention, then they must not all-out believe that they will fail to realize that intention.

Or, slightly stronger, we might require them to have, not just an absence of all-out belief that will fail, but at least a partial belief that they will succeed:

Weak consistency for partial intentions

If an agent forms a partial intention, then they must have a partial belief that they will succeed in realizing that intention.

After all, if they are partially intending to do something, then they will have a subplan in which that partial intention figures. But if they lack a partial belief in success, and yet do not all-out believe that they will fail, this must be because they have failed to consider whether they will succeed; in which case it is surely irrational to have formed that subplan.

The move from very weak consistency to weak consistency look an innocuous one, but it brings an interesting consequence. Very weak consistency clearly places only an atomic requirement on our intentions. Weak consistency appears to do the same. But since it requires a partial belief for each partial intention, it follows that if we have a consistency condition on partial beliefs, it will be transformed into a global requirement on partial intentions. For if each partial intention must bring with it a partial belief, and there is a global consistency requirement on partial beliefs, that in turn will provide a global consistency requirement on the partial intentions.

However, placing a global consistency requirement on partial beliefs is no trivial task. If we were dealing with credences, it would be clear how to proceed. We have a notion of consistency for sets of credences that is provided by the axioms of the probability calculus: a set of credences is consistent just in case it conforms to those axioms. Can we make use of a similar approach for partial beliefs? The problem, of course, is that it was precisely one of the features of partial beliefs, part of what distinguished them from credences, that they did not get numerical values.

There are two ways to go. One is to retreat from the position that an agent will frequently only have partial beliefs. Let us suppose then that, in addition to their partial beliefs, an agent, at least a rational agent, will go on to establish precise credences in ways that correspond to the axioms of the probability calculus. Then we might say, plausibly enough, that the partial beliefs will inherit their values from the corresponding credences in a way that preserves the ratios between them. Thus if your credence that you'll do something yourself is .45, and your credence that you'll get someone else to do it is .45, and you all-out believe that you'll get it done in one of those two ways, then your partial belief in each of the two possibilities will be .5.

The alternative, far more in keeping with the account of partial belief that I originally proposed, is to try for a weaker test of consistency. It would be analogous to the requirements of the axioms of the probability calculus, but applying, where it does, in an ordinal way. Treating one possibility as likely would require treating its negation

as unlikely; the conjunction of two possibilities must be treated as no more likely than either of them, and so on. Such a requirement would be more easily satisfied—in particular, it would place no constraints at all in cases in which probabilities were not even ranked—but it would still provide a global constraint.

Still, it might seem too weak. Suppose, to vary our example, that you think that moving the tree by sawing it up is very likely to succeed, whereas moving it with the car or the crow bar is almost certain to fail. But suppose that, despite this, you plan to put almost all of your effort into the latter two plans. And suppose that there is no independent justification for this: it is not as if the latter two plans require more effort, or have to be tried first, or that there is something immoral or offensive or otherwise undesirable about sawing up the tree. Then there are some grounds for saying that you would be behaving in a practically irrational way. Yet the weak consistency requirement, even with a consistency requirement on partial belief, says nothing against it.

To rule out this kind of imbalance we would need a stronger global measure of coherence akin to that given by the strong consistency requirement. That requirement works by demanding consistency between all the agent's beliefs and the realization of all their intentions. We might try to do likewise. Suppose then that, in addition to placing a value on partial beliefs, we could place a value of between zero and one on each partial intention. Then we could impose a consistency requirement between the beliefs and the intentions by insisting that these states should, when taken *together*, conform to the axioms of the probability calculus. To see how to do this, imagine replacing each occurrence of 'has a partial intention of degree n that he will F' with 'has a partial belief of degree n that he will F' as these occur in the full description of the agent's intentions. Then one needs to check that putting the class of such statements together with the class of statements fully describing the agent's beliefs gives a class that is consistent with the axioms of the probability calculus. Calling this the 'extended probability calculus test' we could now define a strong consistency requirement for partial intentions (all-out intentions now becoming the special case where $n = 1$):

Strong consistency for partial intentions

If an agent forms a partial intention, then that intention must be consistent with the agent's other intentions, and with the agent's beliefs (partial and all-out), as this is established by the extended probability calculus test.

Or, slightly stronger (and following parallel reasoning to that which took us from very weak consistency to weak consistency):

Proportionality requirement

For each partial intention of degree n the agent should have a credence of degree n that it will succeed; and the agent's total set of credences, including those that result from this requirement, should conform to the axioms of the probability calculus.

How though should we place a measure on partial intentions? It is no good saying that a partial intention has degree of n just in case one has a credence of degree n that it

will be successful; that makes the proportionality requirement trivial and so does nothing to rule out the cases of imbalance. Instead we need a truly independent measure of the strength of a partial intention. At a first pass we might say that one partial intention is twice as strong as another just in case one relies on it twice as heavily. This in turn we can gloss by saying that one relies on one intention twice as heavily as on a second just in case one plans to invest twice the resources into the first as into the second.²⁹ That still hasn't got us a cardinal scale, but we can achieve this by adding that the values of the different partial intentions that one has for achieving the same all-out intention sum to one.

This is rough; how, for a start, do we measure the investment of resources? Still, it is hopefully clear enough to give some sense to the proportionality requirement; enough, I think, to see that it is not plausible. For it may be that one option will simply require more resources to be successful, however resources are measured; you are not irrational if you give it what it needs to have a chance of succeeding, even though you think it no more likely to succeed than the less demanding option. Alternatively you might simply prefer one of the options, even if you think it less likely to succeed; again you are not irrational if you put more resources into it than into the others.

Indeed, I suspect that the whole attempt to place a measure on partial intentions is wrong-headed. I do not mean just that they need not receive values, in the way that partial beliefs need not. More fundamentally, when it comes to partial intentions, there is no role for values to play. Our preferences can receive values, and so can our partial beliefs, and these are important for the intentions that we form. But once we have formed a partial intention, we add no explanatory advantage to our account of it if we also ascribe it a value—this is one important way in which partial intentions differ from partial beliefs. I conclude then that there is no place for anything like strong consistency for partial intentions, or the proportionality requirement. In so far as there is need to preserve some degree of proportionality, it should rather come as a constraint on the intentions that one forms: given one's beliefs and preferences, one should all-out intend to put this many resources into this partial intention, and that many into that.

One reason that I have spent so long investigating these various coherence requirements is because they will be important to us in considering whether partial intention entails partial belief. But before we turn to that, let me briefly conclude this section by mentioning a further normative constraint on intentions that has been much discussed. This is what Bratman calls the constraint of *means-ends coherence*: roughly, we must intend the means to our intentions. This goes beyond a simple consistency requirement in that it can require us to add new intentions if these are needed to provide the means to our existing intentions; it is thus the intention analogue to dynamic requirements on belief formation, rather than to static requirements of consistency. The requirement as I have roughly formulated it is a very strong one; almost certainly too

²⁹ The overall approach has much in common with the standard decision theoretic way of inferring credences and preferences from actions. But we don't want to *define* credences in terms of the effort put into our actions, since it would again become trivial that one's credences correspond to one's intentions. And this is plausible anyway, since the concept of credence is also held in place by various other constraints, especially how they are formed, and what they lead one to think and say.

strong, though I shall not argue that here.³⁰ Here we need only observe that the introduction of partial intentions by itself forces no substantial revision on the requirement of means-end coherence. It is just that the intention to achieve the means to a partially intended end may itself be partial. Given that partial intentions from different subplans need not be consistent, you can partially intend the means to one partial intention, whilst at the same time partially intending something inconsistent with that means: one requirement for the success of your subplan that involves getting the tree company out is that you not phone to cancel. So that subplan had better not contain any such intention. Yet it is exactly part of the other subplans, those in which you do move it, that you do phone to cancel.

Intentions and Belief

We are finally in a position to return to the question with which we started, that of whether intention entails belief in success. That has now fractured into a set of four questions depending on whether the belief is all-out or partial, and similarly whether the intention is all-out or partial. I take them in turn.

Does partial intention entail all-out belief?

The answer to this question is obvious: a partial intention does not entail all-out belief in success. If one had all-out belief in success, one would have no need of a merely partial intention.

Does all-out intention entail all-out belief?

The answer to this second question, although it has been much discussed, is, I think almost as obvious. All-out intention does not entail all-out belief. An all-out intention is simply one that is not accompanied by an alternative intention, and that is quite compatible with a merely partial belief in success. One might, for instance, be simply resigned to a high chance of failure in achieving one's end, having no idea of any alternative intentions that one might pursue. Of course, I could have defined an all-out intention as one that is accompanied by an all-out belief in success; but then the claim that all-out intention entails all-out belief would have been trivial. Once it is understood as a substantial claim, we have no grounds for believing it.

To say this is to disagree with David Velleman, who has recently provided a set of arguments that purport to show that all-out intention does entail all-out belief (Velleman 2007). His first contention is that it is only if one believes that one will do something that one will be able to coordinate one's plans around it, and that others will

³⁰ It may, for instance, be more rational to abandon one's intention, than to go on to intend the means to realize it. This suggests the idea that 'oughts' in these contexts typically have wide scope: it is not that, if one has an intention, one ought to intend the means, but rather that one ought, if one has the intention, to intend the means; and this latter obligation can, and sometimes should, be discharged by abandoning the initial intention. For discussion of the plausibility of this move in application to intentions see Setiya, 2007.

be able to do likewise. He writes: 'If I am agnostic as to whether I will be in Chicago on Tuesday, why should anyone plan or act on the assumption of my being there?' Here he gives simple agnosticism as the alternative to belief, which is unsurprising since he is working in an all-out belief framework. But once we move to the kind of framework that I have proposed, the alternative to an all-out belief in success need not be agnosticism. Instead it can be a partial belief. Of course, an all-out belief in success will engender greater confidence in the success of my coordinations. But even a partial belief can provide enough structure to enable defeasible intra-personal coordination—after all, what else am I to do if I can see no other options? And defeasible inter-personal coordination is equally possible, though we would expect a warning of my uncertainty to my collaborators.

Velleman's second claim about intention picks up on Anscombe's observation that the natural expression of an intention to do something is the simple assertion that one will do it. This may well be true in cases when we have an all-out intention and an all-out belief in success. But, as we have seen, where that belief is lacking, intention is more naturally reported by saying that one intends to act (or that one will try to act, if the act of trying can be separated out), often with a qualification that one is unsure of success.

Velleman's third contention concerns coherence requirements on intention. He writes: 'Why, for example should an agent be rationally obliged to arrange means of carrying out an intention, if he is agnostic about whether he will in fact carry it out?' A first response is again that the alternative to all-out belief is not agnosticism but partial belief. If the answer to our remaining question—whether intention entails partial belief—is positive, this may be enough to ground the normative constraints. So, before addressing Velleman's point fully, let us turn to consider that issue. We can do so by considering partial and all-out intentions together.

Does intention (whether partial or all-out) entail partial belief?

Here, at last, a positive response has some plausibility. If I intend to do something then it might seem that I must regard success as a real possibility. If I did not then I would be in no position to plan and coordinate around it; and it is only if I plan and coordinate around it that we can really see it as an intention. I raise two worries though.

The first concerns the strength of this claim. If it is only the claim that an intention entails a partial belief, then the claim is very weak. Indeed it is just the weak consistency requirement for partial intentions that was discussed in the last section, now elevated from the status of a normative requirement to that of a necessary truth. We might hope to make it stronger, but the only way to strengthen it that I can see is in the direction of the strong consistency or proportionality requirements of the last section, again transformed from normative requirements to necessary truths. But as we saw there, it is very hard to make them plausible even as defeasible normative requirements; I see little hope of establishing them as necessary truths.

We should stick then with the weaker claim that intention entails partial belief in success, that is, as seeing success as a live possibility, though perhaps one to which little confidence is given. My second worry concerns even this weak claim. It may be true that if one intends to perform an action, it follows that one will take the performance of

that action as a live possibility in one's deliberations about what to do *up until* that action. But is it so clear that it entails one to take it as a premise in what to do *subsequent* to the time at which the action will or will not have been performed? Take Anscombe's example of the man who is to be tortured, who believes he will break down, but who is determined not to. Might he not be sure that he will break down? (This does not mean that he has a credence of zero that he will fail; only that he does not treat it as a live possibility.) Or take Anscombe's other example of St Peter, determined not to betray Christ, but convinced, since he has it on Christ's own authority, that he will. Again, this seems possible.

It is not easy to get into the mind of either of these agents, but let us focus on the man who is to be tortured, since here at least we can avoid distracting issues of divine foreknowledge. His planning for what to do up until the torture takes it as a live possibility that he will not break down. We can imagine that he prepares himself in whatever way he can to resist (it is wise not to try to imagine this too far). But at the same time his thought about, and plans for, the period subsequent to the torture all involve the premise that he will have broken down. Such a state of mind, admirable though it is, is perhaps *rationally* criticizable; if he is convinced he will break down, should he not rationally give up his resolve to hold out? But we have yet to see a reason for denying that it is possible.

In this case perhaps it will be countered that he has not so much an intention as a hope that he will not break down; after all, it may seem that there is little that he can do to plan for it. But that is surely incidental to the form of the case. In general it seems that there may be cases in which agents have plans to achieve certain ends, but in which they are sure that they will fail. We might wonder why, if they are really not prepared to formulate partial intentions contingent on their success, they are even bothering to try. There are answers that one might give: perhaps they superstitiously fear that entertaining the possibility of success will doom them to failure; perhaps they think that they will be rewarded for having the intention³¹; perhaps, as in the torture case they think that it is right to have the intention, even if they are sure that they will fail. All of these do seem to involve grounds for having or not having an intention that are, in some sense, non-standard: the value is there whether or not one achieves the end. We might then try to embrace the restricted claim that agents' *standard* intentions entail the corresponding partial beliefs that they will succeed; but it is quite unclear that we will be able to define 'standard' in a way that makes this interesting. And even if we could, since intentions are typically arrived at on a variety of grounds, we would risk excluding too much.³²

³¹ Imagine a Kafkaesque case in which I am to be rewarded if I can somehow intend to do something that I believe impossible.

³² Such counterexamples are lent some support by the clinical literature which seems to show cases of subjects with frontal damage maintaining intentions in the knowledge that they are acting against them. For various such cases see Luria 1966, Ch. 5; for discussion of these and others see Monsell 1996. I say that these only lend support, since the findings are somewhat unclear. In particular, most of Luria's cases involve agents who have been given commands but fail to act on them, whilst nevertheless acknowledging what they should be doing so. An alternative but equally

So I am sceptical even of the weak hypothesis that intention entails partial belief. The thesis is better kept as the normative one that intentions rationally require partial belief. If this is right though, it brings us back to the third of Velleman's arguments, discussion of which was postponed above. Velleman asks how we are to understand the source of the normative constraints on intention—most obviously the various coherence requirements discussed above. He argues that, if we understand intention to entail belief, then the normative constraints on belief will be inherited by intention. As we might put it, practical rationality is premised on theoretical rationality.

One response to this is to insist that the normative requirements on practical rationality are *sui generis*.³³ But the considerations raised here suggest another. Velleman's approach has much plausibility when we consider just all-out belief and all-out intention, for there we clearly have a better take on the coherence constraints on theoretical rationality. In particular, we have a very good take on (static) consistency for belief, which we might hope will provide a basis for consistency of intention. But matters are rather different when we turn to the partial case. We have no obvious analogue of consistency when it comes to credences or partial belief. The best that we have is the idea of conformity to the axioms of the probability calculus. And the standard arguments for the rationality of that commitment rely on arguments that themselves appeal to practical rationality: Dutch Book arguments, that aim to show that agents who fail to meet the commitment will accept bets that they will be bound to lose; and Representation Theorem arguments, that aim to show that agents will meet the commitment on the assumption that they have certain preferences, and that they will seek to maximize utility. Some have claimed that here theoretical rationality is grounded in practical rationality.³⁴ But one does not have to go that far. We might say instead that the distinction between theoretical and practical rationality is less clear. Rather than grounding one on the other, we should see them as forming a package that will stand or fall together. And once we think this about partial intentions and partial beliefs, maybe we should explore the idea that this holds for all-out intention and all-out belief too.

Conclusion

I have argued for a novel interpretation of partial belief, and for the existence of partial intentions. Largely independently of either of those controversial claims, I have argued that we should be sceptical of the intention-entails-belief thesis in any of its forms. Lying behind this is a more general thought: that beliefs and intentions are distinct existences, characterized by distinct features. We should not expect necessary connections between them. What we can expect are normative connections. I hope I

plausible explanation is that they have failed to form the right intentions, or that, whilst they have formed them, they are failing to realize that they are not acting on them.

³³ This is the approach taken in Bratman forthcoming.

³⁴ For an effective response to that, see Christensen 1996.

have made some start on identifying what these are, once we move away from all-out beliefs to the more complex structures given by partial beliefs.³⁵

References

- Anscombe, G.E.M. 1963: *Intention*, Second Edition. Oxford: Basil Blackwell.
- Bratman, Michael 1987: *Intention, Plans, and Practical Reason*. Cambridge, MA: Harvard University Press.
- 1992: ‘Practical Reason and Acceptance in a Context’. *Mind* 101 pp. 1–15.
- Forthcoming: ‘Intention, Belief, Practical, Theoretical’. Forthcoming in Jens Timmerman, John Skorupski, and Simon Robertson, (eds) *Spheres of Reason*.
- Budescu, David and Thomas Wallsten 1995: ‘Processing Linguistic Probabilities: General Principles and Empirical Evidence’. In J. Busmeyer, D. Meldin and R. Hastie (eds) *Decision Making from a Cognitive Perspective*. San Diego: Academic Press, pp. 275–318.
- Charniak, Eugene 1991: ‘Bayesian Networks without Tears’, *AI Magazine* 12, pp. 50–63,
- Chater, Nick, Joshua Tenenbaum and Alan Yuille 2006: ‘Probabilistic Models of Cognition: Conceptual Foundations’. *Trends in Cognitive Sciences* 10, pp. 287–91.
- Christensen, David 1996: ‘Dutch Books Depragmatized: Epistemic Consistency for Partial Believers’. *Journal of Philosophy* 93, pp. 450–79.
- Frankish, Keith, Forthcoming: ‘Partial Belief and Flat-out Belief’. Forthcoming in F. Huber and C. Schmidt-Petri (eds) *Degrees of Belief*. Oxford: Oxford University Press.
- Grice, Paul 1971: ‘Intention and Uncertainty’. *Proceedings of the British Academy* 57, pp. 263–79.

³⁵ This paper was presented at seminars at MIT, Stanford and Monash from which it benefited greatly. I am thus very grateful to the audiences there, and to Michael Bratman, Agnieszka Jaworska, Nico Kolodny, Rae Langton, Agustin Rayo, Susanna Rinard, Kieran Setiya, Michael Smith, Brian Weatherston, Timothy Williamson and the referees for *Mind*.

- Hampshire, Stuart, and H. L. A. Hart 1958: 'Decision, Intention and Certainty'. *Mind* 67, pp. 1–12.
- Harman, Gilbert 1976: 'Practical Reasoning'. *Review of Metaphysics*, 29, pp. 431–63.
- 1988: *Change in View*. Cambridge, MA: MIT Press.
- Holtan, Richard 2004: 'Rational Resolve'. *Philosophical Review*, 113, pp. 507–35.
- Kahneman, Daniel and Amos Tversky 1973: 'On the Psychology of Prediction'. *Psychological Review*, 180, 237–257.
- 2000: *Choices, Values and Frames*. Cambridge: Cambridge University Press.
- Kaplan, Mark 1996: *Decision Theory as Philosophy*. Cambridge: Cambridge University Press.
- Lewis, David, 1996: 'Elusive Knowledge'. *Australasian Journal of Philosophy*, 74, pp. 549–67.
- Luria, Aleksandr 1966: *Higher Cortical Functions in Man*. New York: Basic Books.
- Monsell, Stephen 1996: 'Control of Mental Processes'. In Vicki Bruce (ed) *Unsolved Mysteries of the Mind*. Hove: Psychology Press, pp. 93–148.
- Owens, David forthcoming: 'Freedom and Practical Judgement'. Forthcoming in Lucy O'Brien and Matthew Soteriou (eds) *Mental Action*. Oxford: Oxford University Press.
- Setiya, Kieran 2007: 'Cognitivism about Instrumental Reasons'. *Ethics*, 117, pp. 649–67.
- Stalnaker, Robert 1984: *Inquiry*. Cambridge MA: MIT Press.
- Stewart, Neil, Nick Chater, and Gordon D. A. Brown 2005: 'Decision by Sampling'. *Cognitive Psychology* 53, pp. 1–26.
- Vallacher, Robin and Daniel Wegner 1987: 'What Do People Think They Are Doing? Action Identification and Human Behavior'. *Psychological Review* 94, pp. 3–15.
- Velleman, David, 1985: 'Practical Reflection'. *Philosophical Review* 94, pp. 33–61.
- 2007: 'What Good is a Will?'. In Anton Leist and Holger Baumann, (eds) *Action in Context* (Berlin/New York: de Gruyter/Mouton).

Weatherson, Brian, 2005: 'Can We Do Without Pragmatic Encroachment?' *Philosophical Perspectives*, 19, pp. 417-43.

Williamson, Timothy, 2000: *Knowledge and Its Limits* (Oxford: Clarendon Press).