3

The Unity of Reasoning?

JOHN BROOME

1. Introduction

In 'Intention, Belief, Practical, Theoretical' (chapter 2 of this volume), Michael Bratman raises an objection against an account of practical reasoning that I subscribe to. Bratman thinks this account makes practical reasoning, as he once said to me, 'just too cognitive'—too closely tied up with theoretical reasoning. This paper of mine responds to his objection. I shall argue that practical reasoning is genuinely more closely entangled with theoretical reasoning than one might expect. I doubt the two sorts of reasoning are as independent as Bratman would like them to be.

Before I can explain what is at issue, I need to outline my account of reasoning. Doing so will take me up to section 8 in this paper. Section 2 introduces my general approach to reasoning. In section 3 I describe a common model of reasoning, which I call the 'second-order' model, and then reject it. This model would have made reasoning highly unified, but I think it is mistaken. Sections 4 and 5 describe my own first-order model in the context of theoretical reasoning. Section 6 discusses how to extend the first-order model to attitudes other than beliefs; doing so requires us to extend our notion of the content of attitudes. Practical reasoning is reasoning that concludes in an intention, and section 7 therefore considers the content of intentions in particular. It argues that an intention of yours to do something is closely tied to a belief of yours that you will do it. This turns out to be one reason why practical reasoning has a strong cognitive element. Finally, section 8 introduces my first-order model of practical reasoning.

Section 9 describes Bratman's 'just too cognitive' objection to my model. Sections 10 and 11 consider two specific aspects of his objection. They argue that my model of reasoning is cognitive to just the right extent.

2. Reasoning and rationality

In textbooks of logic, the word 'reasoning' often refers to a linked sequence of sentences or propositions. But I use it differently: to refer to a process of a particular sort that goes on in a person's mind.

What particular sort? In this paper, I shall not try to characterize reasoning fully. But I shall mention some of its characteristics. One of them is that it is a process through which certain states of mind give rise to others. The states involved are of the sort that are commonly called *attitudes*. They are relations of a particular sort that hold between a person and a *content*. The content is generally taken to be a proposition, but in section 6 I shall adopt a broader notion of content. The content is the content of the state, and the state is an attitude towards the content. The attitudes I shall mostly be concerned with are beliefs and intentions. (I always use the word 'belief' to refer to an attitude, and never to a proposition that is the content of an attitude.)

When you reason, you acquire an attitude that you did not previously have. Alternatively, you may be confirmed in an attitude you did previously have. Either way, I shall call the attitude the *conclusion-attitude* of the reasoning. I classify sorts of reasoning by the nature of their conclusion-attitudes. I classify reasoning whose conclusion-attitude is a belief as *theoretical reasoning*, and reasoning whose conclusion-attitude is an intention as *practical reasoning*. There may also be reasoning whose conclusion-attitude is a desire or something else, but in this paper I shall concentrate on theoretical and practical reasoning.

Reasoning whose conclusion-attitude is a belief in a normative proposition, such as the proposition that you ought to do something, I classify as theoretical reasoning. It proceeds in exactly the same way as theoretical reasoning with other sorts of content, and answers to exactly the same standards. If you are a non-cognitivist and think beliefs with a normative content are not truly beliefs, this feature of reasoning is a problem you have to deal with; it is not something you can deny (Geach 1960). When we investigate the process of reasoning, it would only be confusing to lump in this type of reasoning with practical reasoning (Anscombe 1957: 60). I do not do so.

Another characteristic of reasoning is that it can bring you to satisfy some of the requirements of rationality. Below are some examples of these requirements. More exactly, these are examples of requirement-schemata; individual requirements are obtained from them by making suitable substitutions for the schematic letters. For '*N*' substitute a term that designates a person, and for '*p*', '*q*', '*e*' and '*m*' substitute sentences that designate propositions.

No contradictory beliefs. Rationality requires of *N* that *N* does not both believe *p* and believe not *p*

Modus ponens. Rationality requires of *N* that, if *N* believes that *p* and *N* believes that if *p* then *q*, and if it matters to *N* whether *q*, then *N* believes that *q*.

Instrumental requirement. Rationality requires of *N* that, if *N* intends that *e*, and if *N* believes that *e* will be so only if *m* is so, and if *N* believes that *m* will be so only if she herself intends that *m*, then *N* intends that *m*.

An obvious modification of *Modus ponens* yields its cousin *Modus tollens*. In each formula, all the attitudes mentioned are supposed to be contemporaneous. To lighten the wording, I have not attached 'at *t*' to each mention of an attitude, but I could have done.

No contradictory beliefs says that rationality requires you not to have contradictory beliefs. *Modus ponens* says that rationality requires you to believe what follows by *modus ponens* from the contents of your beliefs, if it matters to you. *Instrumental requirement* is a rendering of Kant's formula: 'Who wills the end, wills (so far as reason has a decisive influence on his actions) also the means which are indispensably necessary and in his power' (1948: 80–1).

My formula contains more detail than Kant's. For one thing, it contains the clause 'if *N* believes *m* will not be so unless she herself intends that *m*'. ('She herself' is a compound reflexive pronoun, representing in indirect speech the pronoun 'I' in direct speech.) This clause is obviously needed. Even if *N* believes she will not attain her end *e* unless a means *m* obtains, she needs to intend *m* only if she also believes *m* will not obtain unless she intends it. She need not intend it if, say, she thinks *m* will obtain anyway.

Given this clause, there turns out to be no need in the formula for *N* to believe *m* to be an actual means to *e*, rather than simply a necessary condition for *e*. This gives *Instrumental requirement* a little extra generality. Moreover, *N* has to believe *m* is a necessary condition for *e* only in the weak sense that *m* is so if *e* is so. She need not believe it is a necessary condition in the strong sense that, necessarily, *m* is so if *e* is so.

So, when *m* is a means to *e*, *N* has to believe she will not achieve the end unless she takes the means. She does not have to believe it is impossible for her to achieve the end unless she takes the means. *Instrumental requirement* would have few applications if it required this stronger belief (Wallace 2001). But actually it has very many applications. Throughout this paper, I shall use 'necessary condition' in the weak sense.

Often you are caused to satisfy requirements of rationality by unconscious processes that just happen. You intend to go to the cinema, you believe you

will not do that unless you buy a ticket, and you find yourself intending to buy a ticket. Here is another example that will come up again in section 5. The requirement *No contradictory beliefs* is generally enforced by unconscious processes. If you believe a proposition, you generally cannot come to believe its negation without at the same time losing your original belief. Unconscious processes will generally see to it that you do not have both beliefs together.

Perhaps if we were ideally rational, unconscious processes would protect us against any infringement of the requirements of rationality. But we mortals fail to satisfy many requirements. Reasoning is something we can do to improve our score of rationality. One characteristic of reasoning is that it is an activity—something we do. We might count as unconscious reasoning some of the unconscious processes that can bring us to satisfy requirements of rationality. But full-blooded, conscious reasoning is an activity, and that is the sort I shall be exclusively concerned with in this paper. It is an activity by means of which we can bring ourselves to satisfy some requirements of rationality. Section 4 explains how it is an activity.

3. The second-order model

My first task is to explain how the activity of conscious reasoning works. I shall do this in the context of reasoning that can bring you to satisfy *Modus ponens*—a paradigmatic sort of theoretical reasoning. Take this example. Suppose you go to sleep on a cruise ship and wake up hearing gulls. Because of what you hear, you believe there are gulls about. You have a long-established belief that, if there are gulls about, land is nearby. Moreover, it matters to you whether land is nearby. However, because you are sleepy and have not yet thought about where you are, you do not yet believe land is nearby. So you do not satisfy the requirement *Modus ponens* in this instance. You believe there are gulls about; you believe that, if there are gulls about, land is nearby; it matters to you whether land is nearby; but you do not believe land is nearby. By reasoning, you can surely bring yourself to satisfy the requirement. How will your reasoning proceed?

One answer is given by the 'the second-order model' of reasoning, as I shall call it. In this model your reasoning proceeds through your forming a normative belief about your attitudes. More precisely, you form the belief that you ought to have a particular attitude or combination of attitudes, or perhaps you form the weaker belief that you have reason to have a particular attitude or

combination of attitudes. In turn, your normative belief causes you to acquire the attitude or combination of attitudes.

So, in the example, you might form the belief that you ought to believe land is nearby, and that might cause you to believe land is nearby. That would be reasoning in accordance with the second-order model. It would bring you to satisfy *Modus ponens*.

Alternatively, your normative belief might be derived directly from the requirement *Modus ponens* itself. Applied to the example, *Modus ponens* says that rationality requires of you that, if you believe gulls are about and you believe that if gulls are about land is nearby, you believe land is nearby. You might have the corresponding normative belief that you ought, if you believe gulls are about and you believe that if gulls are about land is nearby, to believe land is nearby. You believe you ought to satisfy a particular conditional, and your belief might cause you to satisfy it. If it does, you satisfy the requirement *Modus ponens*. If this is what happens, your reasoning accords with the second-order model.

The second-order model requires you to have a belief about your attitudes. This is a second-order belief—hence the model's name. But one essential characteristic of the model is not registered in its name: the belief has a normative content. I do not count a process as second-order reasoning unless it involves a second-order *normative* belief.

The model easily extends to other sorts of reasoning besides theoretical reasoning. Second-order reasoning can follow the same pattern whatever its conclusion-attitude. It proceeds through your forming a normative belief about a mental state: that you ought to be in it or that you have reason to be in it. Then this normative belief causes you to enter that state. For instance, you might believe you ought to have a particular intention, and your belief might cause you to have it. So this is a very unified model of reasoning.

However, it is an unsuccessful model, at least for my purposes. I shall explain why only very briefly, since this is not the topic of this paper.[1] Both the cause of your second-order belief and its effect are problematic.

The cause first. You could not acquire your second-order belief except through reasoning. To see this, think of the two ways I suggested the second-order model might work in my example. In the first way, your second-order belief is the belief that you ought to believe land is nearby. That belief would have to be derived in some way from your premise-attitudes: your belief that there are gulls about and your belief that if there are gulls about land is

[1] There is a fuller argument in Broome 2006.

nearby. So you could acquire it only by reasoning. If the model is applied the other way, your second-order belief is the belief that you ought, if you believe gulls are about and you believe that if gulls are about land is nearby, to believe land is nearby. The content of this belief is an instance of *Modus ponens*. You could not acquire it except by deriving it through reasoning from a general belief in *Modus ponens*. Either way, reasoning is required to get you to your second-order belief. For similar reasons, in any application of the second-order model, reasoning will be required to get you to your second-order belief.

Therefore the second-order model must apply to your acquisition of your second-order belief, since it is supposed to apply to all reasoning. The model tells us you will have to acquire the second-order belief through a third-order belief. For example, to acquire the second-order belief that you ought to believe land is nearby, you will first need the third-order belief that you ought to believe you ought to believe land is nearby. Acquiring that belief will require a fourth-order belief, and so on. The process of reasoning therefore cannot begin. That is the problem over the cause of your second-order belief.

Now the problem over its effect. The model requires you to progress from believing you ought to have an attitude or combination of attitudes (or believing you have reason to have an attitude or combination of attitudes) to having that attitude or combination. How might that happen? I see two possible routes.

One route is through an unconscious process. You might be so constituted that, when you believe rationality requires you to be in such-and-such a mental state, you tend to enter that state without thinking about it. T. M. Scanlon (1998: 20) claims that some attitudes are 'judgement-sensitive', by which he means they are attitudes 'an ideally rational person would come to have whenever that person judged there to be sufficient reasons for them'. I assume Scanlon thinks that, when an ideally rational person believes she ought to have one of these attitudes, she will come to have it through an unconscious process.

If this is how the effect works, your conscious reasoning is finished once you have your second-order belief. That makes it theoretical reasoning by my definition; it is reasoning that concludes in a belief. If you only reason this way, you will never do any conscious practical reasoning. True, if your second-order belief is a belief that you ought to have a particular intention, it will cause you to have that intention. But that last step will not be conscious reasoning. If the second-order model works this way, it is not a model of conscious practical reasoning.

The second route to the effect goes through an intention. When you believe you ought to have such-and-such an attitude or combination of attitudes, this belief might first bring you to form the intention of having that attitude or combination. Then, second, the intention might cause you to have the attitude or combination.

The first stage of this process is forming the intention. This is the sort of thing that can happen through conscious reasoning. At least, I think so. I think that, when you believe you ought to do something, conscious reasoning can bring you to intend to do it. I cannot prove this is so, but I am willing to grant it here for the sake of argument. I shall also grant that, when you believe you ought to have a particular attitude or combination of attitudes, conscious reasoning could bring you to intend to have that attitude or combination.

But the second stage of the process can rarely succeed. Intending to have an attitude or combination of attitudes is rarely successful; it rarely causes you to have the attitude or combination. You can rarely alter your attitudes by intending to.

Sometimes you can. Sometimes you have a means available of coming to have a particular attitude. For example, going regularly to church may be a means of coming to believe there is a god. If so, an intention to believe there is a god might be effective. If you intend to believe there is a god, your intention might cause you to believe there is a god, through causing you to go regularly to church. But for many attitudes, no such means is available.

Without a means, you cannot alter your attitudes by intending to. You can do some things by intending to, without using a means. For example, you can raise your arm by intending to, without using a means. But we do not have that sort of control over our attitudes. Intending to believe there is a god will not normally cause you to believe there is a god, unless you use a means. I cannot support this claim here; I simply assert it.[2] It means the second-order model cannot work through this second route, because we do not have the sort of control over our attitudes that it would require.

So the second-order model is unsuccessful as an account of conscious reasoning. I turn to a different model.

[2] It is supported by Pamela Hieronymi 2006. My own arguments are different from hers. Bernard Williams (1973) argues that it is a necessary truth that we do not control our beliefs in this way. I have been convinced by Jonathan Bennett (1990) that it is only contingent.

4. The first-order model

I favour the first-order model. I shall use the same example of theoretical reasoning to explain it. You believe there are gulls about, and you believe that if there are gulls about, land is nearby, but you do not believe land is nearby. Moreover, it matters to you whether land is nearby. So you do not satisfy the requirement *Modus ponens* in this instance.

But you can bring yourself to satisfy it by saying to yourself that:

There are gulls about.

If there are gulls about, land is nearby.

So land is nearby.

Here, I have written down a sequence of sentences, which designate propositions. You do not necessarily say the sentences to yourself; you might reason in Swedish, say. But you do say to yourself the propositions that these sentences designate. You say to yourself that there are gulls about, and that, if there are gulls about, land is nearby, and then you say that land is nearby. I shall mention the word 'so' at the end of this section.

You initially believe the first two of these propositions; in saying them to yourself you are expressing your beliefs. You do not initially believe the third. But when you say it to yourself, you express a belief in it. By the time you come to say it, your reasoning has brought you to believe it. By this time, you satisfy *Modus ponens*. That is how the first-order model of reasoning works.

The propositions you say to yourself constitute the contents of your beliefs. You can reason with beliefs only because they are attitudes that have contents. Their content gives you something to reason about.

Saying something to yourself is an act. Sometimes, no doubt, you say things to yourself out loud, but more often you do it silently. In that case, I could alternatively have said you call the proposition to mind; 'saying to yourself' is just a more graphic way of describing what you do. One thing saying to yourself does is bring your beliefs together in your mind, if you have not previously done so.

Your acts of saying to yourself are part of your reasoning but not the whole of it. Your reasoning is the causal process whereby some of your attitudes cause you to acquire a new attitude. It includes a sequence of acts, and it is itself a complex activity. Some of your beliefs cause you to acquire a new belief, through some acts of saying to yourself. The process ends

when you acquire your new belief. That is the conclusion-attitude of the reasoning.

Your acquisition of this belief by reasoning is an intentional act. I said in section 3 that you cannot alter your attitudes by intending to, without using a means. I am not contradicting that claim. You use a means: you acquire this belief by means of reasoning. In any case, you rarely intend to acquire the particular belief you do acquire.[3] Generally, when you are doing some theoretical reasoning, you intend to acquire whatever belief emerges from the reasoning. If your reasoning brings you to believe land is nearby, your acquisition of this belief is intentional, but not because you intend to believe land is nearby. It is intentional because it is the same event as acquiring whatever belief emerges from your reasoning, and you intend to acquire whatever belief emerges from your reasoning.

Since reasoning is a process that takes place among your attitudes, acts of saying to yourself can only form a part of reasoning when they express attitudes of yours. In the example, in saying to yourself that there are gulls about, you must express a belief that there are gulls about. When you say to yourself that land is nearby, you must express a belief that land is nearby, and so on. In the context of belief, saying to yourself is asserting to yourself. True, even if you did not have the corresponding beliefs, you could say to yourself the sequence of sentences:

'There are gulls about.

If there are gulls about, land is nearby.

So land is nearby.'

But in doing that you would not be reasoning because you would not be going through a process that takes place among your attitudes.

In the course of your reasoning, you do not say to yourself any propositions about your attitudes; you say to yourself the propositions that constitute the contents of your attitudes. In the example, you do not say to yourself that you believe there are gulls about, nor that you ought to believe land is nearby. No second-order belief is involved. We may say you reason *with* your beliefs. You reason *about* the content of your beliefs.

The second-order model was supposed to set out from a normative belief about your beliefs. It was blocked because there is no route of reasoning from this second-order normative belief to actually modifying your first-order beliefs. On the other hand, the process I am now describing directly

[3] Simon Robertson pointed out to me that there are exceptions.

modifies your first-order beliefs, because it works on their contents. When you conclude that land is nearby, in doing that you are directly acquiring a new belief.

To summarize: reasoning is a process whereby some of your attitudes give rise to another attitude; the attitudes involved must be ones that have contents; in reasoning you say to yourself the propositions that constitute these contents, and you reason about these contents.

This cannot be a full characterization of reasoning. Not just any mental process that has these features is reasoning. For example, suppose you believe there are gulls about and that, if there are gulls about, land is nearby. Suppose you say to yourself that there are gulls about, and that, if there are gulls about, land is nearby, and suppose this causes you to believe sea-lions are red. That bizarre process is probably not reasoning.

What distinguishes true reasoning from bizarre processes like this? You might think it is the presence of a second-order belief. In my example of genuine reasoning, you moved from believing gulls are about and believing that, if gulls are about, land is nearby to believing land is nearby. You might think this process is reasoning only if you have the second-order belief that: rationality requires you to believe land is nearby if you believe gulls are about and you believe that, if gulls are about, land is nearby.

That is not so. A second-order belief of this sort is not necessary for the process to be reasoning. A reasoner does not have to have such a sophisticated belief as that. To do simple reasoning you do not need to have such a difficult concept as *rationality requires* or even *belief*. In any case, this is the wrong sort of belief for characterizing the process as reasoning. You might hold a weird theory of rationality according to which rationality requires particular relations to hold among your beliefs, quite independently of any inferential relations that may hold among the beliefs' contents. A process that brought you to satisfy one of these requirements you believe in would not count as reasoning. To be reasoning, a process needs to match your beliefs about inferential relations among contents, rather than your second-order beliefs about rational relations among attitudes.

It is much more plausible that a belief about inferential relations is needed to separate your reasoning process from other processes such as the bizarre one. For the process I described to be reasoning, you might need to believe that, from the proposition that gulls are about and the proposition that, if gulls are about, land is nearby, it follows that land is nearby. I do not deny that you may require a belief like this if you are to reason. We know from Lewis Carroll's 'What the Tortoise Said to Achilles' that we should not treat it as a further premise-attitude in the reasoning process. Your reasoning proceeds in the way

I originally described, but for it to be reasoning, you may need to have this belief in the background.

But my own view is that the way to characterize a process as reasoning is that it is computational. If I am right, it adds to the ways in which reasoning is an activity, since computation is something you do. You operate on the contents of your attitudes computationally. I think that, when you say to yourself the word 'so' or its equivalent in another language, it marks your computation. Computation is too big and difficult a topic to broach in this paper. I shall allow myself the assumption that reasoning is an operation on the contents of your attitudes.

5. Reverse reasoning

As deductive reasoning by *Modus ponens*, the gulls example is a paradigm of theoretical reasoning. But it represents only a small fraction of theoretical reasoning, and leaves a great deal to be explained. For example, theoretical reasoning often involves the attitude of hypothesizing, as well as the attitude of believing. You may reason on the basis of a hypothesis and conclude your reasoning by coming to believe a conditional proposition. How does that work? Fortunately, I do not need to answer that question, nor investigate all the complexities of theoretical reasoning here.

But I do need to mention one complexity. Reasoning often does not proceed in the linear fashion illustrated in the example, where your reasoning sets out from some initial beliefs and concludes with a new belief. Actual theoretical reasoning will often lead you to drop one of your initial beliefs, rather than acquire a new one. Dropping a premise-belief will bring you to satisfy the requirement of rationality *Modus ponens*, just as well as acquiring a conclusion-belief will. How does reasoning of that sort work?

Suppose you embark on the process of reasoning I described, but do not conclude it. You believe there are gulls about and that if there are gulls about land is nearby. You say to yourself that there are gulls about and that if there are gulls about land is nearby, but you do not end up believing land is nearby. You fail to come to believe a new proposition.

You remain in violation of the requirement *Modus ponens*. But you may yet be able to achieve rationality through reasoning. You may not be able to do so if it is some irrational obstruction that prevents you from believing the conclusion that land is nearby. But normally, when you fail to believe the

conclusion of reasoning you embark on, it is because you believe its negation. In the example, you believe land is not nearby. This means you have an alternative way to make yourself rational. You can put your reasoning into reverse. You can say to yourself in sequence that:

Land is not nearby.

If there are gulls about, land is nearby.

So, there are no gulls about.

At the end you believe there are no gulls about.

Rationality requires you not to believe both a proposition and its negation. I said in section 2 that this requirement is normally enforced by unconscious processes. You will not be able to come to believe there are no gulls about whilst still believing there are gulls about. So for you to complete your reverse reasoning, two things must happen: you must come to believe there are no gulls about and you must stop believing there are gulls about. Provided you do both, you will end up satisfying both *Modus tollens* and *Modus ponens*.

Since reasoning is not necessarily linear, it might go in either of two directions. In the example, it could have gone forward and brought you to believe land is nearby, but actually it went backward and brought you to believe there are no gulls about. This raises a new question. What controls the direction of your reasoning?

In the example, you start with competing beliefs: that there are gulls about, that if there are gulls about land is nearby, and that land is not nearby. In a sense, the direction of your reasoning must be determined by the relative robustness of these beliefs. How convinced are you that what you hear is gulls, or that land is not nearby? How willingly would you give up those beliefs? Robustness in this context is a complex notion; some of its complexity will appear in section 11. So to give a proper account of the direction of reasoning would be a substantial undertaking, which I cannot embark on here.

But I need to say one thing. No doubt the direction of your reasoning may be influenced by your second-order beliefs about what you ought to believe or about what rationality requires you to believe. You may engage in second-order reasoning in order to form second-order beliefs of this sort. That is to say, you may engage in reasoning about your first-order beliefs. This shows that our first-order reasoning may be influenced by second-order beliefs and second-order reasoning. It has a 'superstructure', as I put it.

This is not inconsistent with my rejection of the second-order model of reasoning in section 3. I rejected that model on the grounds that second-order

beliefs cannot directly alter your first-order beliefs through your intentions, as the second-order model requires. You cannot control your beliefs that way. I am not now suggesting that second-order beliefs directly alter your first-order beliefs. I am saying they may influence the direction of your first-order reasoning. Since reasoning is something you do, it is to some extent under your control. Your second-order beliefs can therefore influence the direction of your first-order reasoning. If you do some second-order reasoning, it does not replace first-order reasoning; it directs it. It has to be first-order reasoning that ultimately alters your first-order beliefs.

6. The contents of attitudes

Now to reasoning other than theoretical reasoning. For most attitudes apart from beliefs, it is at first hard to see how we can reason with them, operating on their contents. Beliefs have a special feature that allows us to reason with them. When you say to yourself that gulls are about, you express your attitude of belief. You also, in a different sense, express the content of that belief. You say the proposition that is the content of your belief.

Reasoning requires this sort of double expression. You reason with your attitudes. When you say things to yourself in the course of reasoning, you must be expressing attitudes that you have; otherwise you would not be truly reasoning as I understand it. But, as well as that, reasoning is an operation on the contents of the attitudes. You need those contents before your mind, which means you have to express them to yourself. So your expression of your attitudes also has to express the contents of those attitudes.

At first sight, few attitudes share with beliefs the property that you can express them and their contents together. Take a desire, for example. Most philosophers take the content of a desire to be a proposition. Suppose you want to be in England. Then, according to this common view, the content of your desire is the proposition that you are in England. To express this content you would say 'I am in England'. But saying 'I am in England' is not a way to express a desire to be in England. So you do not express the desire and its content together.

This difficulty can be overcome by enlarging our notion of content. Philosophers commonly assume that attitudes of different sorts can have the same content, which is a proposition. So you might have a belief that you are in England, say, or a desire to be in England, and both those attitudes would have as their content the proposition that you are in England. The attitudes

have the same content, but in the two different cases you have a different attitude to the content—a believing attitude in one case and a desiring attitude in the other. That is the common view. The alternative is to take the content of an attitude to be a proposition together with a mark of some sort, which marks the type of attitude it is. In this way the differences in attitudes can be absorbed into the contents of the attitudes.[4]

For instance, if you desire to be in England, you might express the content of your desire by 'I am in England—nice'. 'I am in England' designates the propositional part of the content, and 'nice' designates the mark that marks the content as the content of a desire. I shall take the mark for belief to be the absence of any other sort of mark. So, if you believe you are in England, I shall designate the content of your belief by 'I am in England'. I shall give the name 'markers' to the linguistic items, such as the word 'nice', that designate marks.

If you say to yourself the sentence 'I am in England—nice', you are expressing your desire, and you are also expressing its content. In this way the mark gives a desire the special feature that a belief has: expressing the content of the desire is also expressing the desire itself. So, when you express the desire, you make its content available to be operated on, as reasoning requires. Desires become available for reasoning with. Other attitudes can also be made available for reasoning if we understand their contents to be propositions together with marks.

Marks solve another problem too. Different sorts of attitude are subject to different requirements of rationality. For example, what is required of your beliefs is different from what is required of your intentions. This should be obvious, but here is an example all the same. Rationality requires you to intend what you believe is the best means to an end that you intend. However, rationality does not require you to intend what you believe is the best means to something you believe you will do, when you do not intend to do it. Here is the difference spelt out. This is true:

> Rationality requires of *N* that, if *N* intends that *e*, and if *N* believes that the best way for *e* to be so is for *m* to be so, and if *N* believes that *m* will be so only if she intends that *m*, then *N* intends that *m*.

But this is false:

> Rationality requires of *N* that, if *N* believes that *e*, and if *N* believes that the best way for *e* to be so is for *m* to be so, and if *N* believes that *m* will be so only if she intends that *m*, then *N* intends that *m*.

[4] Important exponents of this idea are Richard Hare (1952) and Paul Grice (2001).

To help make sense of the comparison, use this example. For '*e*' substitute 'you reveal the secret' and for '*m*' 'you write to the newspapers'. Rationality requires of you that, if you intend to reveal the secret, and believe the best way to do so is to write to the newspapers, you intend to write to the newspapers. But suppose you do not intend to reveal the secret, but nevertheless believe you will reveal it because it will be wrung out of you by torture. Then you might not intend to write to the newspapers, and yet be perfectly rational.

Since reasoning can bring you to satisfy requirements of rationality, and since different types of attitude are subject to different requirements, reasoning with different types of attitude must be subject to different rules. Consequently, when reasoning, you must keep track of what type of attitude you are dealing with.[5] If reasoning is an operation on contents, that means the type of attitude needs to be registered within the content of the attitude. Marks do that. This is the second problem they solve.

This paper is concerned with conscious reasoning only. I take it for granted that when we reason consciously we express our attitudes to ourselves using language. So we can only reason about marked contents if we have markers in our language to designate them. This means that the reasoning we can do is limited by the contingencies of our language. That should not be a great surprise. I suggested in section 2 that conscious reasoning is an activity we undertake to bring ourselves to satisfy particular requirements of rationality, when we fail to satisfy them through automatic processes. It is a self-help activity that makes up for weaknesses in our contingent nature. Our language has presumably evolved to allow us to do the reasoning that, contingently, we need to do.

The study of reasoning cannot therefore be done on purely a priori grounds. We have to take note of the way we do, contingently, reason using our language. I shall take reasoning in English as my example.

In English we do have some markers. They are not much like the artificial construction I used as an illustration. In English, markers in a sentence are generally special moods. For instance, the marker for desire in English is the optative construction. Robert Browning said: 'Oh, to be in England now that April's there!' This sentence expresses the proposition that Browning is in England together with the mark for desire, designated by an optative construction. When Browning said to himself 'Oh, to be in England!', he expressed his desire to be in England, and he also expressed the content of his desire, understood as a proposition with a mark. He could alternatively have

[5] Nadeem Hussain impressed the importance of this point on me long ago.

adopted my mode of expression and said 'I am in England—nice'. That might not have served his poetic purpose so well.

The optative construction can make the contents of desires available to be operated on. So Browning could have reasoned with his desires. For instance, he could have said to himself these sentences:

'Oh, to be in England now that April's there!

England is rainy in April.

So: Oh, to be somewhere rainy in April!'

Actually, he did not reason like that, and he would not have been reasoning correctly if he had. I am only illustrating how in principle we could reason with attitudes other than beliefs. Which patterns would constitute correct reasoning is another matter.

7. Intention and belief

Now we have the materials for an account of non-theoretical reasoning. They make for a lot of disunity in reasoning. Since the different attitudes we might reason with—beliefs, intentions and so on—have different contents, we can expect no homogeneity in the way that reasoning with them goes.

I am particularly concerned with practical reasoning, which involves intentions. Do we have a marker for intention? Some philosophers treat the imperative mood as the marker for intention; they take intentions to be expressed in imperative sentences. This is certainly incorrect in English, where imperative sentences are used for the speech-act of commanding. You use an imperative when you think a person does not have a particular intention, and hope to induce her to have it. You say 'Turn left!' when you think the driver does not intend to turn left, and you hope to induce her to intend to turn left. When you say an imperative sentence to yourself, it serves the same purpose. I say to myself 'John, stop staring out of the window and get on with your work!' in the hope of inducing myself to intend to do what I tell myself to do. If I already had that intention I would not issue the command. I form most of my intentions without this palaver, so I rarely need to use the imperative mood in speaking to myself. This may explain why English does not contain a genuine first-person imperative form.

I do not think the philosophical idea that an intention is expressed by the imperative mood is really meant seriously. It is a fantasy. In their writing, philosophers often need a way to distinguish the content of an intention from

the content of a belief, and they find the imperative mood a handy artificial device for doing so.

These philosophers are up against the problem that the ordinary way to express an intention in English is the same as the ordinary way to express a belief: by a sentence in the indicative mood. To express your intention of being at the conference on Thursday, you say 'I shall be at the conference on Thursday'. The difference between a belief and an intention may sometimes be signalled by subtle inflexions. According to the Fowlers (1906: ch. 2), the first-person singular 'I will' is more coloured than 'I shall', and may be used to indicate the presence of will or intention. But a coloured grammatical form is supposed to be reserved for special, exceptional cases, whereas expressing intentions is the commonest use we have for first-person singular future-tense sentences. I therefore doubt that any special coloured form was used as the standard expression of an intention even in the Fowlers' day. In any case, in modern usage the distinction is almost non-existent, and the very same sentence, with either 'shall' or 'will', may express either an intention or a belief.

This is puzzling. It is not only philosophers who need to distinguish the content of a belief from the content of an intention. Members of the public do too. When you are reasoning, you must keep track of when you are expressing a belief and when an intention; yet your means of expression apparently does not make the distinction. How do you do it?

The answer is that your expressions of intention contain inaudible markers, which distinguish them from your expressions of belief. In principle, we could use an audible marker and say such sentences as: 'I shall reveal the secret—will', where 'will' designates the mark for intention. But actually the marker is inaudible. To express your intention, you say 'I shall reveal the secret', and the context shows you are expressing an intention.

Inaudible markers certainly exist in English. We sometimes put questions or commands in the indicative mood. You might say as a command: 'No member of the team will drink alcohol tonight'. This sentence has an inaudible imperative marker. However, the inaudible marker in an expression of intention functions differently from the one in questions and commands, and because of that a serious puzzle remains.

When you say 'No member of the team will drink alcohol tonight', meaning it as a command, you are not saying *that* no member of the team will drink alcohol tonight. You are not asserting that proposition. You may sincerely issue the command without believing the proposition. The inaudible marker in the sentence makes this possible: it ensures that you are not asserting the proposition that the sentence would designate if it did not have the

marker. The marker cancels out the ordinary assertoric effect of saying the sentence.

But suppose you say 'I shall be at the conference on Thursday', expressing an intention of yours to be at the conference on Thursday. Though you are expressing an intention, you are in fact saying—asserting—that you will be at the conference on Thursday. You could not sincerely say this if you did not believe it. A burglar overhearing you will know that Thursday would be a good day for burgling your house. So far as the burglar is concerned, it makes no difference whether your sentence expresses an intention to be at the conference or merely a belief that you will be at it. The inaudible marker for intention does not cancel the assertoric effect of the sentence. In this it differs from the inaudible markers for commands and questions.

When you assert something sincerely, you express a belief. You therefore cannot sincerely express an intention to do something unless you believe you will do it. At the same time as you express your intention, you express your belief. When you say 'I shall reveal the secret', expressing an intention, you are in fact expressing two attitudes of yours, an intention and a belief. With marks, the contents of these attitudes are, respectively, 'You will reveal the secret—will' and 'You will reveal the secret'.

This is puzzling; how can the expression of an intention also express a belief? It means that, when you express an intention sincerely, you must have the corresponding belief. But does anything guarantee that is so?

It is sometimes thought that, if you intend to do something, you must believe you will do it. If that were true, it would provide a simple solution to the puzzle, but actually it is not true. There is a complication, which I need to stress because it will become important in section 10. You can have an intention without believing you have it, and conversely you can believe you have an intention without actually having it. Michael Bratman gives a good example. On Sunday, you believe you intend to shop on Thursday and you do not believe you intend to shop on Friday, whereas actually you intend to shop on Friday and you do not intend to shop on Thursday.

How could that happen? Here is one possibility. You regularly shop on Thursdays. However, a couple of weeks ago you discovered the shops will be closed this Thursday, so you decided to shop on Friday this week. You have now temporarily forgotten your change of plan. You will remember it when you start making more detailed plans for the week, so your intention to go on Friday is still in place. But at the moment you do not believe it is in place.

Because you do not believe you intend to shop on Friday, you do not believe you will shop on Friday. Yet you do intend to shop on Friday. So

you can intend to do something without believing you will do it. But if you *believe* you intend to do something, in that case you must believe you will do it. Since you believe you intend to shop on Thursday, you believe you will shop on Thursday. The true connection between an intention and a belief is that, if you believe you intend to do something, you believe you will do it. More accurately:

> *Belief–intention link.* If *N* believes that she herself intends that *p*, then, because of that, *N* believes that *p*.

This is what allows an intention and a belief to be expressed by the same sentence. You cannot sincerely express an intention to do something without believing you have that intention. Consequently, when you sincerely express an intention, you must believe you will do what you intend. That is why you are in a position to express this belief simultaneously.

The evidence I offer for the existence of this belief-intention link is just that expressing an intention is also expressing a belief; both take the form of saying the same indicative sentence. So you cannot express an intention without expressing a belief that you will do what you intend. And you cannot do that sincerely without having the belief.

I have to admit that we call some attitudes 'intentions' even though they do not satisfy the belief-intention link (Bratman 1987: 37–8). For example, you might intend to go round by the library on the way home, and believe you intend it, whilst not fully believing you will do it. You might know about your own forgetfulness. You could not express this intention in what I have called the 'ordinary' way; you could not say 'I shall go round by the library on the way home'. To say that would also be to express a belief that you will go round by the library on the way home, which you cannot sincerely do. Instead, you might express your attitude by saying 'I intend to go round by the library on the way home'. In that way you avoid expressing the corresponding belief.

My account of practical reasoning applies only to intentions that do satisfy the belief-intention link. These ones can be expressed in the ordinary way, and it will quickly appear in the next section that this feature is essential to my account. Other attitudes, even if they are correctly called intentions, are beyond the scope of my account. I do not know whether we can do reasoning with these other attitudes, and if we can, I do not know how.

In this paper, I apply the term 'intention' only to attitudes that satisfy the belief-intention link. If you like, you may think of this formula as partly specifying what I mean by 'intention'.

8. Practical reasoning

The connection between intention and belief makes the problem of understanding practical reasoning more complex. Practical reasoning is reasoning whose conclusion-attitude is an intention. But if you form an intention by explicit reasoning, in which you express the conclusion to yourself in English, you will also express a belief. Your reasoning must therefore lead to a belief as well as an intention. The result is that theoretical and practical reasoning are entangled together.

With one sort of practical reasoning, the entanglement is very tight. This is a sort of instrumental reasoning: specifically reasoning from intending an end to intending what you believe is a necessary condition for that end. I shall now present an account of how this sort of practical reasoning works. The entanglement will become plain as the account develops.

Take this example. You intend to visit Venice. You believe you will not do so unless you buy a ticket to Venice. You also believe you will not buy a ticket to Venice unless you intend to. (So, for example, you have not authorized your agent to do your ticket-buying.) However, suppose at present that you do not intend to buy a ticket. You do not satisfy *Instrumental requirement*, therefore. We can expect you to be able to acquire this intention through practical reasoning, and thereby come to satisfy the requirement. How might that work?

Intuitively, like this. You might say to yourself the sentences:

'I shall visit Venice.

I shall not visit Venice if I do not buy a ticket to Venice.

So I shall buy a ticket to Venice.'

With the first sentence, you express your initial intention to visit Venice. With the second, you expresses your initial belief that your buying a ticket is a necessary condition for your doing so. With the third, you express an intention to buy a ticket. You did not have this intention initially, but you acquire it by means of your reasoning. This way, you bring yourself to satisfy *Instrumental requirement*. This is an intuitively satisfying example of practical reasoning. But we need a proper account of how it works. Here is mine.

When you say to yourself 'I shall visit Venice', whatever else you do, you assert to yourself that you will visit Venice. Indeed, in your reasoning

as a whole, you assert three propositions to yourself in sequence. You assert that:

> You will visit Venice.
>
> You will not visit Venice if you do not buy a ticket to Venice.
>
> You will buy a ticket to Venice.

When you say the first of these propositions, you express a belief that you will visit Venice. You next express a belief that you will not visit Venice unless you buy a ticket there. From the two propositions that you believe, it *follows* that you will buy a ticket. Naturally it matters to you whether you buy a ticket. So, if you do not believe you will buy a ticket, you violate, not just the practical *Instrumental requirement* but the theoretical requirement *Modus tollens* as well.

Indeed, the sequence of your beliefs would constitute a piece of theoretical reasoning of the sort I described in section 4. Straightforward theoretical reasoning would bring you to believe you will buy a ticket, and thereby put you in conformity with *Modus tollens*. However, you are limited by a special constraint, which prevents you from doing straightforward theoretical reasoning. You believe you will not buy a ticket unless you intend to do so. Therefore, if you are rational, you cannot come to believe you will buy a ticket, as *Modus tollens* requires, unless you also come to believe you intend to buy a ticket.

Compare the reverse theoretical reasoning I described in section 5. Given that you believe land is not nearby and that if there are gulls about land is nearby, reverse theoretical reasoning will take you to a belief that there are no gulls about. But you cannot acquire this belief unless you drop your belief that there are gulls about. Unconscious process will see to that. Here, unconscious process will see to it that you cannot acquire the belief that you will buy a ticket unless you also acquire the belief that you intend to buy a ticket.

In normal circumstances—I shall mention abnormal ones in section 10—you cannot acquire the belief that you intend to buy a ticket except by actually acquiring the intention of buying one. So, if you are to complete your reasoning and come to believe you will buy a ticket, you must also come to intend to buy one. Both the intention and the belief must click into place. Provided the reasoning proceeds smoothly, you will end up with both. You may then say to yourself 'I shall buy a ticket', thereby expressing both a newly acquired intention and a newly acquired belief.

Your reasoning is practical because it concludes in an intention, and it is theoretical because it concludes in a belief. It brings you to satisfy both

the practical *Instrumental requirement* and the theoretical requirement *Modus tollens*. It is a single piece of reasoning that has both theoretical and practical aspects.

The practical and theoretical aspects cannot be separated. You might think that the theoretical reasoning leads the practical reasoning: that you first come to believe you will buy a ticket through some theoretical reasoning, and this belief then causes you to intend to buy a ticket. But that is not so. Merely believing you will buy a ticket will not cause you to intend to buy one. Believing you will do something does not cause you to intend to do it. The causal connection is the other way round: intending to buy a ticket causes you to believe you will buy one. You cannot complete the theoretical reasoning till you form the intention. Given that you believe you will not buy a ticket unless you intend to, no purely theoretical reasoning can take you from your premise-beliefs—that you will visit Venice and you will not visit Venice if you do not buy a ticket—to a conclusion-belief that you will buy a ticket. You can acquire that belief only by acquiring the intention of buying a ticket.

Like pure theoretical reasoning, the reasoning I have described may not go through smoothly in a linear fashion. Your acquisition of a new belief might be blocked. You might find you cannot believe you will buy a ticket. For example, you might believe you will not buy a ticket—perhaps you believe you have too little money. This belief may prevent you from forming the opposite belief that you will buy a ticket. If it does, the theoretical part of the reasoning is blocked.

The practical part will be blocked as a consequence. The belief-intention link ensures that, because you cannot believe you will buy a ticket, you cannot believe you intend to buy one. But you cannot form an intention by explicit reasoning without at the same time coming to believe you have the intention. Since you cannot have that belief, you cannot come to have the intention by explicit reasoning. Your reasoning cannot bring you to intend to buy a ticket.

Since your reasoning is blocked, you will remain in violation of the two requirements of rationality *Instrumental requirement* and *Modus tollens*. But you have an alternative way to satisfy these requirements through reasoning: you can throw your reasoning into reverse. You say to yourself that:

> You will not buy a ticket to Venice.
>
> You will not visit Venice if you do not buy a ticket to Venice.
>
> So, you will not visit Venice.

You end up with a belief that you will not visit Venice.

On the face of it, this reasoning is purely theoretical, not practical. The two premise-attitudes are beliefs that are not linked to intentions, and the conclusion-attitude is a belief. Your reasoning here is comparable to the reverse theoretical reasoning described in section 5. In the example of that section, you could not acquire the conclusion-belief that there are no gulls about without dropping your previous belief that there are gulls about. Here, you cannot acquire the conclusion-belief that you will not visit Venice without dropping your previous belief that you will visit Venice. The reasoning concludes in a disbelief as well as a belief.

But the present example is more complex than the one in section 5. Because of the belief-intention link, you cannot drop your previous belief that you will visit Venice without also dropping your belief that you intend to visit Venice. And you cannot do that by explicit reasoning except by dropping the actual intention to visit Venice. To complete your reasoning, you will have to give up this intention. So the reasoning concludes in a non-intention, as well as in a belief and a disbelief. I originally defined practical reasoning as reasoning that concludes in an intention. Let me now include reasoning that concludes in a non-intention. Then this piece of reverse reasoning has a practical aspect as well as a theoretical one.

It brings you to lose both your belief that you will visit Venice and your intention of visiting Venice. As a result, you will satisfy both *Instrumental requirement* and *Modus tollens*.

9. Just too cognitive

That is my account of reasoning from intending an end to intending what you believe is a necessary condition for your end. It is a development of an earlier account of mine (Broome 2002), which attracted some accurate criticism from Jay Wallace (2001, particularly 18–19).[6] My present account owes a lot to Wallace's paper.

In my account, theoretical and practical reasoning are inextricably entangled. We could say that this account exhibits a close unity of reasoning. In one way that is its attraction; it makes it a good antidote to scepticism about practical

[6] Wallace's arguments (and some of mine in my earlier paper, I now realize) are aimed at supporting a requirement of rationality, namely a version of *Instrumental requirement*. I am here looking for something different: a process of reasoning that can bring you to satisfy the requirement. One effect of the difference is that Wallace calls on a different connection between intention and belief from mine. Rather than the belief–intention link I specified, he assumes that, if you intend to do something, you believe it is possible for you to do it.

reasoning. Scepticism has been rife since David Hume (1978, book 3, part 1, section 1) announced that 'reason is the discovery of truth and falsehood'. Practical reasoning is plainly not the discovery of truth and falsehood, so if Hume was right, there is no practical reasoning. But I think no one should doubt that you can reason from intending an end to intending what you believe is a necessary condition for your end. The Venice example is an intuitively attractive example of this sort of reasoning, even if my account of it is mistaken. Furthermore, if my account is not mistaken, the example shows that this sort of reasoning is made correct by the same valid syllogism as makes the corresponding theoretical reasoning correct. This strengthens the example as an antidote to scepticism.

However, Michael Bratman objects to my account just because it tangles theoretical and practical reasoning so tightly together. Intuitively, practical reasoning should be more independent of theoretical reasoning, he thinks. As he put it to me once, the account makes practical reasoning *just too cognitive*.

I agree it makes theoretical and practical reasoning unintuitively close. However, I cannot see how to disentangle them in this particular case. Moreover, I do not think this 'just too cognitive' objection amounts to a real argument against the account. I can see two more specific objections that could be developed from it, and when I look at those two, I think they fail. If anything, they lead me to think better of my account, because I think they reveal genuine features of reasoning. In the next two sections, I shall set out these two more specific objections and respond to them.

10. First objection: intending vs believing you intend

There is first an objection Bratman himself mentions in 'Intention, Belief, Practical, Theoretical' (chapter 2 of this volume). In my account of practical reasoning, what really matters are not your intentions themselves, but your beliefs about what intentions you have. This is so for both the premise-intention and the conclusion-intention. Bratman concentrates on the latter, so I shall start with that one.

Look again at the forward—not reverse—piece of reasoning I described. To satisfy the theoretical requirement *Modus tollens*, you need to come to believe you will buy a ticket. This cannot happen unless you also come to believe you intend to buy a ticket. I said that in normal circumstances you cannot come to believe you intend to buy a ticket except by actually coming

to intend to buy one. But if you did happen to believe you intend to buy a ticket, without actually intending to, then the reasoning could come to an end with your believing you will buy a ticket, without your actually intending to buy one. At the end, you would satisfy *Modus tollens*, but not *Instrumental requirement*. In that case, my account could not explain how reasoning might bring you to satisfy *Instrumental requirement*.

It could indeed happen that you believe you intend to buy a ticket without actually intending to buy one. I do not see how it could be caused to happen by your reasoning—at least if you are rational—but you might already have the belief before you start the reasoning. You might previously have intended to buy a ticket, but subsequently changed your mind and dropped the intention. Then, by the time you come to do your reasoning, you might have forgotten you changed your mind. The shopping example in section 7 shows how this sort of thing can happen.

Suppose you are in this position before you start your reasoning. As before, you say to yourself 'I shall visit Venice', expressing an intention and a belief. You say to yourself 'I shall not visit Venice if I do not buy a ticket', expressing a belief. Then you say to yourself 'So, I shall buy a ticket'. You there express a belief, and in a sense you express an intention. But since you already believe you have this intention, you are in a position to say this to yourself sincerely, without actually having the intention. No intention need click into place. The process of reasoning I described will therefore not get you to intend to buy a ticket.

Bratman thinks that, nevertheless, true practical reasoning should be able to arrive at an actual intention. He concludes that my account of your reasoning cannot be right.

But I think my account properly represents the facts in this case. There is simply no way you can reason your way to an actual intention, past your belief that you have an intention. Let me fill out the example with some more detail. With the help of your travel agent, you have planned a world tour. You intend to go tomorrow to her office to buy the various tickets she has assembled for you. You believe a ticket to Venice is among them, so you believe you intend to buy a ticket to Venice.

However, a while ago you asked your agent to cancel the detour to Venice you had previously arranged, so actually she has no ticket to Venice for you. You have temporarily forgotten that. You will remember when you check her invoice this evening. And you actually do not intend to buy a ticket to Venice. You are not disposed to take any steps towards buying a ticket to Venice, beyond what you are going to do anyway, which is to go to your agent's office and buy the tickets that are waiting there. When you remember

that you cancelled your detour to Venice, you will do nothing towards buying a ticket to Venice.

Nevertheless, at present you do genuinely intend to visit Venice, though without much enthusiasm. The evidence for this is that there are things you are disposed to do to get yourself there. For instance, on the Internet, you are just about to book a water-taxi from Venice airport to the city.

Your position is this, then. You intend to go to Venice, you believe you will not go to Venice unless you buy a ticket to Venice, and you believe you will not buy a ticket to Venice unless you intend to buy one. Yet you do not intend to buy a ticket. You therefore fail to satisfy *Instrumental requirement*. If *Instrumental requirement* is a genuine requirement of rationality, you are irrational.

I think no reasoning you could do would get you out of this position. You believe you intend to buy a ticket to Venice, and you are unable to reason your way round this belief to arrive at an actual intention to buy one. Nor, alternatively, could reasoning lead you to drop your intention to visit Venice. Once you discover the error in your beliefs—you remember your agent has no ticket to Venice for you—at that point you can reason your way to dropping your intention. But, as things stand, I think your false belief blocks any reasoning that can bring you to satisfy *Instrumental requirement*.

Now a different example of a similar difficulty. This one concerns a premise-intention rather than a conclusion-intention. Suppose you intend to visit Venice, but do not believe you have this intention. (You have arranged to go to a conference there, but have temporarily forgotten your arrangement.) Suppose, as before, that you believe you will not go to Venice unless you buy a ticket to Venice, and that you will not buy a ticket to Venice unless you intend to buy one. But suppose you do not intend to buy a ticket to Venice. Then you do not satisfy *Instrumental requirement*. According to my account, no process of reasoning is available to you that will bring you to satisfy this requirement. This is because, on my account, a premise-intention can only enter into your reasoning through your believing you have it.

Yet another example is the converse of that one. Suppose you believe you intend to visit Venice but actually do not intend to. Then according to my account, you can reason your way to an intention to buy a ticket, even though you do not actually intend to visit Venice.

All these consequences of my account seem correct to me. I do not think they show that my account has implausible consequences. According to the account, an intention only enters your reasoning through your belief that you have the intention; I think the examples suggest this is intuitively correct. Indeed, it seems inevitable for conscious reasoning. The attitudes that can

figure in your conscious reasoning must be ones you are conscious of, which is to say ones you believe you have. Conversely, if you believe you have an attitude that actually you do not have, the fact that you do not have it cannot impinge on your conscious reasoning.

The consequences are disturbing from a theoretical point of view. They leave us with an unpleasant dilemma. In two of the examples I described, reasoning cannot bring you to satisfy *Instrumental requirement*. We shall have to conclude either that *Instrumental requirement* is not a genuine requirement of rationality, or alternatively that it is a genuine requirement but not one that reasoning can always bring you to satisfy. I am not comfortable with either conclusion. In this paper I shall not venture to decide between them.

11. Second objection: intending vs believing you will

Now the second 'just too cognitive' objection. In my account of instrumental reasoning, what does the work in the premise position is not your intention to go to Venice but your belief that you will go there. The belief is caused by the intention, but the reasoning would be no different if you had the belief without the intention. That may seem unsatisfactory.

To make it easier to examine this problem, I shall set aside the previous one. So I shall assume that your beliefs about your intentions match your actual intentions: you believe you intend to do something if and only if you actually intend to do it.

Now, go back to the example as I originally described it, and change it in just one respect. Now suppose you do not intend to go to Venice (nor believe you intend to go to Venice) but merely believe you will go to Venice. (Perhaps you have noticed that everyone in your set goes to Venice at some time.) Suppose everything else is the same as in the original example. You believe you will not visit Venice unless you buy a ticket to Venice. You believe you will not buy a ticket to Venice unless you intend to buy one. This means you believe you will not buy a ticket unless you *now* intend to buy one. (Perhaps the last tickets to Venice are now being sold, because Venice is about to sink.) Finally, in this version of the example, you do not intend to buy a ticket (nor believe you intend to). Nor do you believe you will buy one.

Can reasoning bring you to intend to buy a ticket? It cannot. You could not possibly reason your way to intending to buy a ticket to Venice when you

do not even intend (or believe you intend) to go there. If you intended to go to Venice, you could reason your way to intending to buy a ticket, but not when you merely believe you will go to Venice.

So it makes a decided difference to what you can achieve by reasoning whether you set out with an intention to visit Venice or merely a belief that you will. The rule in this case is: intention out—intention in.[7] An intention cannot come out of your reasoning if no intention goes in (or, more exactly, if no belief that you have an intention goes in). Do not misunderstand me. This rule applies *only* to reasoning of the form I am describing. I am not saying in general that an intention can only arise from a previous intention; I think that is false. An intention can be derived from a belief through other sorts of reasoning. For example, if you believe you ought to buy a ticket, I think you can derive by reasoning the intention of buying a ticket. Moreover, you can acquire an intention in other ways than by reasoning. An intention can be caused in all sorts of ways, and its causal antecedents do not have to include a previous intention. Still, for instrumental reasoning on the pattern I am describing, intention out—intention in does appear to be a rule.

If it is a rule, we should expect it to register in the content of your reasoning. But it does not register according to my account. When you intend to visit Venice, your intention figures in the reasoning only through the belief it gives you that you will visit Venice. So it makes no difference to the content of the reasoning whether you have the intention or merely the belief.

This does not mean the example is inconsistent with my account. Here is what my account says about it. You believe you will visit Venice and you believe you will not do so unless you buy a ticket. However, initially you do not believe you will buy a ticket. This means you are violating the theoretical requirement of rationality *Modus tollens*. Theoretical reasoning is available to bring you into conformity with that requirement. You express your two beliefs to yourself by saying 'I shall go to Venice' and 'I shall not go to Venice if I do not buy a ticket'. But you will find yourself blocked from concluding that you will buy a ticket. You cannot acquire this belief without acquiring the intention of buying a ticket, because you believe you will not buy a ticket unless you intend to, and nothing will cause you to acquire the intention.

So the only way you will be able to bring yourself to satisfy *Modus tollens* is by throwing your reasoning into reverse. You will be able to do this, provided you believe you will not buy a ticket. If you have that belief, by backwards reasoning you will be able to arrive at the belief that you will not visit Venice.

[7] I am adapting a remark of Jay Wallace's (1990).

At the same time, you will lose your belief that you will visit Venice. Once you do that, you will satisfy *Modus tollens*.

Your reasoning may go either forward or backward, and my account is consistent with either. I said in section 5 that the direction of reasoning is a separate matter from the reasoning itself. The direction of reasoning can be controlled by what I called the superstructure, which includes second-order beliefs and second-order reasoning. One feature of the superstructure is what I called the 'robustness' of beliefs. The example shows that an intention can give a sort of robustness to its corresponding belief. If you believe you will visit Venice, and if this belief is supported by an intention to visit Venice, it may be robust enough to push your first-order reasoning through to an intention to buy a ticket. But in a case like the one I am describing, your belief is not robust at all if it is not supported by an intention. In the absence of an intention, your first-order reasoning will have to go into reverse.

So my account is consistent with the example. Still, it may not appear to be the right account of the rule of intention out—intention in. You might think this rule should not be enforced only by the superstructure of reasoning. You might think it ought to play a role in the first-order reasoning itself. That is the objection. My account does not sufficiently distinguish between the role of an intention and the role of a mere belief in first-order reasoning.

However, if anything is wrong here, it is with my account of theoretical reasoning, not with my account of practical reasoning. In the present version of the example, which starts from a pure belief that you will visit Venice, your reasoning is purely theoretical. Its premise-attitudes are a belief that you will visit Venice, and a belief that you will not visit Venice unless you buy a ticket. Neither has anything to do with an intention. If you could reach a conclusion-belief that you will buy a ticket, that would be purely theoretical reasoning. But you cannot get that far. The theoretical reasoning that would take you to this conclusion-belief is blocked.

Various things may block theoretical reasoning, and prevent it from taking you directly through to a conclusion-belief. The example in section 5 illustrates one sort of block: you may believe the negation of the conclusion. In our present case, the cause of the blockage is your belief that you will not buy a ticket unless you intend to buy one, together with your lack of an intention to visit Venice (or lack of a belief that you have such an intention). It is a feature of theoretical reasoning that it can be blocked in this particular way, among others.

The presence or absence of the block does not figure in your first-order theoretical reasoning, according to my account. Does this mean that my account is inadequate? If it did, my account of *theoretical* reasoning would be

just too cognitive, because it does not register at the first order whether or not you intend to go to Venice. But I think this is an inappropriate complaint. I think it perfectly appropriate to locate the source of the blockage in the superstructure. The source is itself partly a second-order belief. In the example, it is partly your belief that you will not buy a ticket unless you intend to. This is a belief about an attitude—an intention—of yours. There is no place for this belief anywhere within your first-order reasoning. On my account, it influences the direction of your first-order reasoning, but it is not part of the first-order reasoning. I think that is appropriate.

12. Summary

This paper investigated the activity of reasoning. I started with theoretical reasoning, principally so as to distinguish the second-order model of reasoning from the first-order model, and to reject the former. Then I went on to a special sort of practical reasoning: reasoning from intending an end to intending what you believe is a necessary condition for the end.

I pointed out that, when you express an intention of yours, you also express a belief that you will do what you intend. A consequence is that practical reasoning is closely entangled with theoretical reasoning, according to my account of it. I went on to consider the objection that the entanglement is implausibly close, and rejected it. I concluded that practical reasoning does indeed have a strong cognitive element, as my account of it implies.[8]

References

Anscombe, Elizabeth (1957). *Intention*. Oxford: Blackwell.

Bennett, Jonathan (1990). 'Why is Belief Involuntary?', *Analysis* 50: 87–107.

Bratman, Michael E. (1987). *Intention, Plans and Practical Reason*. Cambridge, MA: Harvard University Press.

Broome, John (2002). 'Practical Reasoning', in *Reason and Nature: Essays in the Theory of Rationality*, ed. J. Bermùdez and A. Millar. Oxford: Oxford University Press (2002), 85–111.

——(2006). 'Reasoning with Preferences?', in *Preference Formation and Well-Being*, ed. Serena Olsaretti. Cambridge: Cambridge University Press (2006), 183–208.

[8] I did much of the work for this paper while I held a Leverhulme Major Research Fellowship. I am very grateful to the Leverhulme Trust for its support. I am also very grateful to Olav Gjelsvik, John Skorupski and Simon Robertson for their excellent advice.

Carroll, Lewis (1895). 'What the Tortoise Said to Achilles', *Mind* 4: 278–80.
Fowler, H. W., and Fowler, F. G. (1906). *The King's English*. Oxford: Oxford University Press.
Geach, Peter (1960). 'Ascriptivism', *Philosophical Review* 69: 221–5.
Grice, Paul (2001). *Aspects of Reason*. Oxford: Oxford University Press.
Hare, R. M. (1952). *The Language of Morals*, Oxford: Oxford University Press.
Hieronymi, Pamela (2006). 'Controlling Attitudes', *Pacific Philosophical Quarterly* 87: 45–74.
Hume, David (1978). *A Treatise of Human Nature*, ed. L. A. Selby-Bigge and P. H. Nidditch. Oxford: Oxford University Press.
Kant, Immanuel (1948). *Groundwork of the Metaphysic of Morals*, translated by H. J. Paton under the title *The Moral Law*. London: Hutchinson (1948).
Scanlon, T. M. (1998). *What We Owe To Each Other*. Ithaca, NY: Belknap Press of Harvard University Press.
Wallace, R. Jay (1990). 'How to Argue about Practical Reason', *Mind* 99: 355–85.
——(2001). 'Normativity, Commitment, and Instrumental Reason', *Philosophers' Imprint* 1.
Williams, Bernard (1973). 'Deciding to Believe', in his *Problems of the Self*. Cambridge: Cambridge University Press (1973),136–51.