Practical reasoning

John Broome

University of Oxford

Aristotle took practical reasoning to be reasoning that concludes in an action. But an action – at least a physical one – requires more than reasoning ability; it requires physical ability too. Intending to act is as close to acting as reasoning alone can get us, so we should take practical reasoning to be reasoning that concludes in an intention. Sections 1 and 2 of this paper argue that there is such a thing as genuine practical reasoning, concluding in an intention. It can be correct, valid reasoning, and section 2 explains how.

Section 3 deals with an incidental complication that is caused by a special feature of the concept of intention. Sections 4 and 5 then explore the normativity of practical reasoning. They argue that, although practical reasoning concludes in an intention, it gives the reasoner no reason to have that intention.

This paper considers only one type of practical reasoning, namely instrumental reasoning. Moreover, up to the end of section 5 it considers only instrumental reasoning that proceeds from an end to a means that the reasoner believes is necessary. This one special case is enough to demonstrate that genuine practical reasoning exists. But reasoning of this special type is rare, so the remainder of the paper investigates how successfully the particular conclusions of sections 1–5 can be extended. It considers instrumental reasoning when the means is not believed to be necessary. Sections 6 and 7 examine and reject the idea that this type of practical reasoning depends on normative reasoning. Section 8, 9 and 10 examine and reject the idea that decision theory provides a good account of it.

After those negative arguments, section 11 looks for a correct account of instrumental reasoning from an end to a means that is not believed to be necessary.

1. Intention reasoning

You might reason like this:

	I am going to buy a boat	(1a)
and	For me to buy a boat, a necessary means is to borrow money	(1b)
so	I shall borrow money.	(1c)

In this piece of reasoning, I mean (1a) to express an intention of yours to buy a boat, rather than a belief that you are going to buy a boat. I mean (1b) to express a belief of yours, and I mean (1c) to express a decision you make.

Think of this as a process of reasoning you might actually run through in your head. You might do so when you get your bank statement. You intend to buy a boat, you now form the belief that a necessary means of doing so is to borrow money, and then you go through the reasoning set out in (1). It takes you from two of your existing states of mind, an intention and a belief, to a new state of mind, a new intention. To form an intention by reasoning is to make a decision. So this piece of reasoning concludes in a decision.

I shall call reasoning that concludes with the forming of an intention 'intention reasoning'. Intention reasoning is *practical* reasoning; it gets as close to action as reasoning can.

Your reasoning process is a particular type of practical reasoning. It is *instrumental* reasoning, which means it is concerned with taking an appropriate means to an end. This paper considers instrumental reasoning only. There are other sorts of practical reasoning too, but if we want to understand practical reasoning, it is a good idea to start with instrumental reasoning, because it is less controversial than other sorts.

Your reasoning is a very special type of instrumental reasoning: it is reasoning from an end to a means that you believe to be necessary. Later in this paper I shall examine instrumental reasoning more generally, starting in section 6.

All this begs the question of whether intention reasoning like (1) is truly reasoning at all.

We could call almost any process of thought 'reasoning', but I shall use this term only for *correct* reasoning. (I shall reserve the term 'valid' for the content of reasoning, and use 'correct' as the corresponding term applied to the process of reasoning.) Intuitively, (1) expresses correct reasoning, but we need an explanation of why. Section 2 offers one.

2. The correctness of intention reasoning and belief reasoning
In explaining why (1) expresses correct reasoning, I shall make a few assumptions. I shall assume that intentions and beliefs are propositional attitudes. That is to say, they are states of mind that have contents, and their contents are propositions. I shall assume your name is 'Chris', and I shall assume that the proposition that Chris will buy a boat is the same as the proposition that you, Chris, would express with the sentence 'I am going to buy a boat'. So the content of your intention expressed in (1a) is the proposition that Chris will buy a boat.

Idiomatically, we generally described a person's intentions using an infinitive rather than a noun clause. For instance, we say, 'You intend tomorrow to be a restful day', rather than, 'You intend that tomorrow is a restful day'. But I take these two sentences to have the same meaning. In this context, an infinitive and a 'that' clause are alternative ways of expressing a proposition.

A sentence expresses a proposition. When you *say* a sentence to yourself or out loud, you often also express a particular attitude to the proposition that the sentence expresses. Sometimes the nature of this attitude is indicated by an inflexion or mood within the sentence, but sometimes not. In English, a first-person, future-tense, indicative sentence can be used to express either an intention or a belief. In saying 'I am going to buy a boat' you may express either an intention of buying a boat or a belief that you are going to buy a boat. Sometimes, but not always, the difference is idiomatically indicated by subtle inflections involving 'shall', 'will' and 'going to'.

I assumed you use (1a) and (1c) to express intentions. So, writing 'I' for 'you intend that' and 'B' for 'you believe that' – both operators on propositions – your reasoning in (1) can be described like this:

	I(Chris will buy a boat)	(2a)
and	B(For Chris to buy a boat, a necessary means is for Chris to borrow money)	(2b)
so	I(Chris will borrow money).	(2c)

This is a description of your reasoning, not an inference. From the fact that you intend to buy a boat, and the fact that you believe a necessary means of doing so is to borrow money, we cannot infer that you intend to borrow money. You may not have that intention, for instance if you are irrational.[1]

Compare this piece of intention reasoning with theoretical reasoning or, as I shall call it 'belief reasoning'. All reasoning, conceived as a process, starts from existing states of mind and concludes in a new state of mind. By 'belief reasoning', I mean reasoning that concludes in a belief. Here is an example of belief reasoning that you might go through:

	B(Chris will buy a boat)	(3a)
and	B(For Chris to buy a boat, a necessary means is for Chris to borrow money)	(3b)
so	B(Chris will borrow money).	(3c)

(To distinguish the beliefs in (3) from the intentions in (2), imagine the transactions will be made by your attorney, against your will.) Like (2), (3) describes a process of reasoning. It is not an inference; from (3a) and (3b), we cannot infer (3c). Even if you have the beliefs (3a) and (3b), you might not have the belief (3c), for instance if you are irrational.

The content of the reasoning process (3) is:

	Chris will buy a boat	(4a)
and	For Chris to buy a boat a necessary means is for Chris to borrow money	(4b)
so	Chris will borrow money.	(4c)

No doubt you will actually express these propositions to yourself in the first person:

	I shall buy a boat
and	For me to buy a boat, a necessary means is to borrow money
so	I shall borrow money.

Whether expressed in the first or the third person, the syllogism in (4) constitutes a valid inference. (4b) is stronger than it needs to be to make the inference valid; I have given it a special modality for a reason I shall explain in Section 3. A material conditional would have done. Still, (4) as it stands is certainly valid. The three propositions in (4) stand in a particular relation to each other, a relation such that, if the first two are true, so is the third. This makes the inference valid, which in turn makes the process of reasoning correct. It is correct to proceed from your first two beliefs to the third. The reasoning is correct because its content is a valid inference.

The validity of this content plays a part in other sorts of reasoning besides (3), because propositions do not need to be believed for them to play a part in reasoning. For example, the same content might feature in hypothetical reasoning, where you do not believe (4a) or (4b), but are working out what would be true if they were true. The same validity also plays a part in intention reasoning. Intention reasoning (2) has the same content as belief reasoning (3). This content, expressed in the third person, is the valid inference (4). It turns out that the validity of (4) makes (2) correct just as it makes (3) correct.

The difference between (2) and (3) is not in their content but in the attitude you take towards their content. For instance, in the belief reasoning (3) your attitude towards the proposition (4a) is to take it as true, whereas in the intention reasoning (2) your attitude is to be set to make this proposition true. In (3) you take (4a) and (4b) as true. Because (4) is a valid inference, if (4a) and (4b) are true, (4c) must also be true. So you cannot rationally take (4a) and (4b) as true without taking (4c) as true. This is why (3) is correct belief reasoning. In (2) you take (4b) as true, and are set to make (4a) true. Because (4) is a valid inference, if (4a) and (4b) are true, (4c) must also be true. So you cannot rationally be set to make (4a) true, and take (4b) as true, without being set to make (4c) true. That is why (2) is correct intention reasoning.

Both (2) and (3) appropriately follow the transmission of truth through the valid inference (4). (2) follows it in a truth-making way and (3) in a truth-taking way. Even if David Hume was right that reasoning is concerned only with truth,[2] he should still have recognized that reasoning can transmit the truth-making attitude as well as the truth-taking attitude. It can transmit intention as well as belief, so reasoning can be practical.

Intention reasoning and belief reasoning both follow truth. Their logic is the same, because belief and intention are both attitudes to propositions rather than features of propositions themselves, and the logic belongs to the propositions.[3] The attitudes are not part of the content of the reasoning.

To be sure, some propositions are about your attitudes. An example is the proposition that you intend to buy a boat. You may have attitudes of intention or belief towards these propositions too, and they can figure in reasoning. This is a correct piece of belief reasoning, for example:

| | B(Chris intends to buy a boat) |
| and | B(If Chris intends to buy a boat, Chris intends to borrow money) |

so   B(Chris intends to borrow money).

(In formal statements throughout this paper, 'if' stands for material implication.) No attitudes appear within the content of this reasoning, only propositions that you have particular attitudes. Peter Geach's well-known objection to noncognitivism in ethics is no objection to my account of practical reasoning.[4]

3. Restrictions on the notion of intention

If you believe the premises of a valid inference, then if you reason correctly you will believe the conclusion. It would be nice if I could say in parallel that, if you intend some of the premises of a valid inference, and believe the others, then if you reason correctly you will intend the conclusion. But in fact this is not so; the parallel between belief reasoning and intention reasoning is not complete.

To see why, notice first that you may believe the conclusion of the inference is true anyway, without your intending it. You may intend to buy a boat, and believe a necessary means is to borrow money, but not intend to borrow money because you believe you have already done so. Consequently, you see no need to intend to borrow money now.

Let us extend the notion of 'being set to make true' to cover propositions that you believe are true anyway, without your making them true. In this example, let us say that you do indeed set yourself to make it true that you borrow money. You happen to have an easy time of it, because you believe nothing is required of you to make it true.

Granted this extension, then whenever you are set to make true some of the premises of a valid inference, correct reasoning will bring you to set yourself to make the conclusion true. So with the extended notion of 'being set to make true' in place of intention, we have a good parallel with belief reasoning.

However, *intention* reasoning is not parallel, and we can now see why not. It is because our notion of intention does not coincide with the notion of being set to make true, especially now we have extended the latter notion.

The notion of intention is narrower than the notion of being set to make true. Several restrictions apply to the former but not the latter. The first is that intending does not cover propositions that you believe are true anyway, without your making them true. We have extended 'being set to make true' to cover these propositions, but 'intend' does not cover them.

The second restriction appears in this example. Consider the putative reasoning:

        I(Chris will buy a boat)
and   B(If Chris will buy a boat, Chris will find new friends)
so   I(Chris will find new friends).

The content of this reasoning is a valid inference, but the reasoning itself is not correct. You need not intend to find new friends, just because you intend to buy a boat and recognize that this will be the consequence. You may simply take it as a side effect. Certainly, you must adopt the truth-making attitude towards it; in fulfilling your intention of buying a boat you must make this proposition true. But you need not intend it. You need not intend whatever follows from an end that you intend.

That is why I formulated my original example with the strong modality 'necessary means' in the major premise, rather than simply a material conditional. Although you need not intend whatever follows from an end that you intend, you must intend whatever you believe is a necessary means to an end that you intend, unless (the first restriction) you believe it is true anyway, without your intending it.

The third restriction is more controversial. It arises from something Frances Kamm calls 'triple effect'.[5] Consider this putative reasoning:

> I(Chris will have more fun)

and > B(For Chris to have more fun, a necessary means is for Chris to find new friends)

so > I(Chris will find new friends).

This contains the 'necessary means' modality. Nevertheless, Kamm argues it may not be correct reasoning. Suppose you independently intend to buy a boat, and believe that finding new friends will be a side effect of doing so. In that case, Kamm thinks you need not separately intend to find new friends as a means of having more fun.

If Kamm is right, this third restriction is a variant of the first. You need not intend the proposition that you will find new friends, because you believe that proposition is true 'anyway', whether or not you intend it. It is true without your intending it, because it is a side effect of a different intention.

In summary, not all putative intention reasoning whose content is a valid inference is correct. It must satisfy some further constraints. I incorporated the second constraint in my example, by using the 'necessary means' modality. The first and third constraints require that you should not believe the conclusion of the inference is true anyway, independently of your intending it.

Whenever the constraints are satisfied, intention reasoning is genuinely correct reasoning, provided its content is a valid inference. It is a clear paradigm of practical reasoning.

4. Reasoning is not reason-giving

Because intention reasoning is reasoning, we may say it is normatively guided; precisely what this means will appear by the end of this section. But intention reasoning is normative in no other way. Its content is not normative; it is not about what you ought to do or have a reason to do. (I use term 'a reason' for a *pro tanto* reason. If you have a reason to do something, that means you ought to do it unless you also have a contrary reason not to.)

Furthermore, intention reasoning is not ought-giving nor even reason-giving; that is what I shall argue in this section and the next. In my example, intention reasoning takes you from your intention of buying a boat and your belief that borrowing money is the only the means of doing so, to an intention to borrow money. But it does not determine that you ought to borrow money, nor even that you have a reason to borrow money.

I could equally well say that the intention and belief on which the reasoning is premised are neither ought-giving nor reason-giving. Your intention of buying a boat, and your belief that to do so you must borrow money together give you no reason to borrow money. Nor do they make it the case that you ought to borrow money. That is what I shall argue.

Reasoning in general is neither ought-giving nor reason-giving. Once again, it is easiest to see this by looking at the more familiar example of belief reasoning. Belief reasoning is neither ought-giving nor reason-giving or, to put it another way, beliefs are neither ought-giving nor reason-giving. Suppose you believe some proposition p from which q can be inferred by a valid inference. It does not follow that you ought to believe q, nor that you have a reason to believe q. This section defends this claim. Section 5 returns to intention reasoning.

My claim has nothing to do with the complexity of the inference. Throughout, I shall assume that the inference from p to q is immediate and obvious.

From one point of view, a defence scarcely seems needed, because the conclusion is surely obvious. Suppose, for instance, that you ought not to believe p, though you do believe it. Then obviously it may not be the case that you ought to believe q, nor that you have a reason

to believe it.

Moreover, this conclusion can be supported by a simple argument. The proposition p itself follows from p by an immediate and obvious inference. But from the fact that you believe p it plainly cannot follow that you ought to believe p, or have a reason to believe p. Beliefs do not justify themselves. So it cannot be a general principle that if you believe p you ought to believe its immediate and obvious consequences, nor that you have a reason to do so.

Still, however obvious this may be, there is a plausible contrary thought. If you believe p, then surely in some sense or other that gives you a reason to believe its consequence q, when q is different from p itself. No doubt, we would not say that believing p gives you a reason to believe p itself, because you do not need a reason for that; you already believe it. But if you are going to believe a consequence of p that is distinct from p, you do need a reason for that, and surely believing p gives you one, in some sense or other.

So two plausible views conflict. The truth lies with the first, but is easy to explain why the second seems attractive. The truth is that a particular relation holds between your believing p and your believing q: one *normatively requires* the other, as I shall put it. It is not the case that believing p gives you a reason to believe q, but it is easy to mistake the relation of giving a reason to for the relation of normatively requiring.

In slightly formal notation, the truth is that

$$B p \text{ requires } B q, \tag{5}$$

where 'B' still stands for 'you believe that', and 'requires' stands for 'normatively requires you to see to it that'. I have examined the notion of normative requirement more thoroughly in 'Normative requirements', a companion paper to this one. Here I shall mention only one essential features.[6] (5) implies

$$O(B p \supset B q), \tag{6}$$

where 'O' stands for 'you ought to see to it that'. But (5) does not imply

$$B p \supset RB q, \tag{7}$$

where 'R' stands for 'you have a reason to see to it that'. (5) attaches normativity to the relation between believing p and believing q, not to believing q itself. It does not say literally that if you believe p you have a reason to believe q. On the other hand, (7) attaches normativity to the consequent rather than the relation. This means the consequent in (7) can be detached by modus ponens: from $B p$ and (7), we can infer $RB p$. If you believe p, then (7) says you have a reason to believe q. (5) does not allow detachment of that sort.

It is easy to confuse a reason and a normative requirement because they both involve a weakening of the notion of 'ought', and it is easy to muddle the two sorts of weakening. A reason is a weakened sort of ought; it is weakened by being made *pro tanto*. A normative requirement is also a weakened sort of ought; it is weakened by being made relative.

The difference between them can be described like this: a normative requirement is strict but relative; a reason is slack but absolute. A normative requirement is relative in the sense that it is a relation between two propositions. It is the truth of the first (such as your believing p) that requires you to see to the truth of the second (such as your believing q), and the requirement cannot be detached from its antecedent. But a normative requirement is strict because it is strictly a requirement: if you do not satisfy it, you fail in something that is required of you. (6) expresses this strictness. It says you ought to see to it that if you believe p you believe q. So if you believe p but not q, you fail to see to something you ought to see to. On the other hand, a reason is not relative, but it is slack in that it is only *pro tanto*. If you do not do what you have a reason to do, you may not have failed in any way; you may have performed exactly as you should have. You may have had a better reason not to do this thing,

and correctly followed the better reason.

The relation between believing p and believing q, when q follows from p by an immediate and obvious inference is plainly strict. That is to say, if you believe p and not q, you are definitely not entirely as you ought to be. If the relation was simply that believing p gave you a reason to believe q, it would be slack; you might believe p and not q, yet still be entirely as you ought to be. This could happen if, say, you had a better reason not to believe q. But actually this is not possible. To be sure, you might have a good reason not to believe q, and an appropriate response might be to stop believing p. That way, you can escape from the requirement that is imposed on you by your belief in p. But if you do not take this way out, and you believe p without believing q, you are not entirely as you ought to be. So the relation between believing p and believing q is strict, and it therefore cannot be the reason-giving relation. It is the relation of normative requirement.

To summarize, believing p does not give you a reason to believe its consequence q. That is to say, (7) may be false. On the other hand, believing p does normatively require you to believe q. That is to say, (5) is true.

When one proposition implies another (and the implication is immediate and obvious), believing the one requires you to believe the other. The relation of normative requirement merely reflects, at the level of beliefs, the relation of implication that holds between propositions. At the beginning of this section, I said that reasoning is normatively guided, and this is its only normative feature. The idea of normative requirement makes explicit what this normative guidance amounts to.

A final note. One sometimes comes across a third weakening of 'ought', besides a reason and a normative requirement. This is the *subjective* ought. People sometimes say that if you believe p then you subjectively ought to believe its consequence q. This idea is right in making the connection between believing p and believing q both strict and relative, but wrong in making it relative to the person rather than to her belief. To see its wrongness, imagine you have inconsistent beliefs: you believe p and also r, and q follows from p, but $\neg q$ follows from r. Then according to this story, you subjectively ought to believe q, and you subjectively ought to believe $\neg q$. That is hard to make sense of; it sounds contradictory. The right thing to say is that your belief in p requires you to believe q, and your belief in r requires you to believe $\neg q$. That is easy to make sense of, though it reveals a failing on your part. You ought not to have these inconsistent beliefs. They impose inconsistent requirements on you, as one would expect.

5. Intentions are not reason-giving

Everything I have said about belief reasoning applies to intention reasoning too. If you intend an end, and believe some act is a necessary means to it, your intention and belief normatively require you to intend the means. Reasoning will bring you to intend it. This is simply to restate the correctness of intention reasoning such at (2). It does not follow that you ought to intend the means, nor that you have a reason to intend it.

For example, if you intend to buy a boat, and believe borrowing money is a necessary means of doing so, it does not follow that you ought to intend to borrow money or have a reason to do so. For instance, if you ought not to intend to buy a boat in the first place, it might be that you ought not to intend to borrow money either. Still, instrumental reasoning will bring you to intend to borrow money. This intention is required by your intention of buying a boat and your belief that borrowing money is a necessary means to do so.

You might be tempted to think that your original intention and belief must give you some

sort of a reason to intend to borrow money. Similarly, you might be tempted to think that believing a proposition must give you some sort of reason to believe its obvious consequences. But I explained in section 4 that the latter view is a mistake. It is an easy mistake to make, because the truth is that believing a proposition normatively requires you to believe its obvious consequence, and it is easy to mistake a normative requirement for a reason. It is easy to make the same mistake with intention reasoning. In section 4, I applied the test of strictness to distinguish a reason from a normative requirement, and I can apply the same test here. The relation between intending an end and intending what you believe is a necessary means is plainly strict. If you intend an end and do not intend what you believe is a necessary means, you are definitely not entirely as you ought to be. Therefore, this relation must be normative requirement. If intending an end merely gave you a reason to intend what you believe is a necessary means, the relation would be slack.

Understanding the relation of normative requirement is essential to understanding how instrumental reasoning is even possible. If you intend an end, you must be able to reason correctly about how to bring it about, and you must be able to do this even if you have no reason to intend the end. Indeed, you must be able to do it even if you actually ought not to intend this end. The possibility of correct instrumental reasoning cannot depend on the rightness of the end. But if you ought not to intend the end, you may have no reason to intend the means. So if instrumental reasoning had to work by giving you a reason to intend the means, instrumental reasoning would not be possible when you ought not to intend the end. Fortunately, it does not work this way. It works through a normative requirement.

This point supports an argument of Michael Bratman's in *Intention, Plans and Practical Reason*. Bratman points out that a theory of practical reasoning must not imply 'bootstrapping', as he calls it.[7] Just because you form the intention of buying a boat, that intention cannot possibly create a reason for you to do so, if you do not already have one. An intention cannot justify itself, by providing itself with a reason. Still, your intention must play a role in your reasoning and lead you to other intentions, including the intention of borrowing money. The purpose of Bratman's book is to explain the vital role of intentions in the reasoned planning of our lives. This role creates a puzzle. If your intention leads to other intentions by reasoning, does it create a reason for those other intentions? If it did, that would still be a sort of bootstrapping; a reason would be being pulled into existence out of nothing. We have a solution to this puzzle. There is no bootstrapping. One intention gives rise to another by means of reasoning, but no reasons are involved. There is only the relation of normative requirement, which is given us by the correctness of the reasoning.[8]

Christine Korsgaard's 'The normativity of instrumental reason' is an important discussion of what Korsgaard calls the 'instrumental principle', the principle 'that practical reason requires us to take the means to our ends'.[9] Korsgaard is concerned with how instrumental reasoning is possible. One of the main conclusions she draws from her argument is that: 'Unless there are normative principles directing us to the adoption of certain ends, there can be no requirement to take the means to our ends'.[10] She says, in more detail:

> For the instrumental principle to provide you with a reason [to take the means to an end], you must think that the fact that you will an end *is a reason* for the end. It's not exactly that there has to be a *further* reason; it's just that you must take the act of your own will to be normative for you. And of course this cannot mean merely that you are *going* to pursue the end. It means that your willing the end gives it a normative status for you, that your willing the end in a sense makes it good. The instrumental principle can only be normative if we take ourselves to be capable of giving laws to ourselves – or, in Kant's own phrase, if we

take our own wills to be *legislative*.[11]

Korsgaard would be right that you must have a reason for your end *if* instrumental reasoning provided you with a reason to take a means. And if instrumental reasoning did this for you, I dare say her other conclusions would follow. But instrumental reasoning does not provide you with a reason to take a means. That is not how it works. Willing (or intending) an end normatively requires you to will whatever you believe is a necessary means to the end. It does not give you a reason to take the means, and it does not need to. So actually Korsgaard's conclusions do not follow. Willing an end need not give the end a normative status for you, for instance.

6. Normative reasoning and normative ascent

The intention reasoning illustrated in (2) is rare. We rarely have the opportunity to engage in reasoning just like that. It is available only when we believe some particular means is necessary to an end of ours, and we rarely encounter a means that we believe to be strictly necessary. Normally we recognize several alternative ways of achieving each of our ends. How does practical reasoning work then? For the sake of argument up to now, I have ignored the fact that you can buy a boat by other means than borrowing money. For example, you can sell your house or rob a bank. These are definitely inferior alternatives to borrowing money, but they are possible. How can you reason about means of buying a boat, bearing this in mind?

When you do not believe the means is necessary, it is tempting to turn to normative reasoning. Consider the inference that you would express to yourself as:

	I intend to buy a boat	
and	If I intend to buy a boat, I ought to borrow money,	
so	I ought to borrow money.	

In neutral terms:

	Chris intends to buy a boat	(8b)
and	If Chris intends to buy a boat, Chris ought to borrow money,	(8c)
so	Chris ought to borrow money.	(8d)

Let us assume a cognitivist account of normative statements. Given cognitivism, (8) expresses a valid inference concluding in a normative proposition (8d). So the following describes a correct piece of belief reasoning you might go through:

	B(Chris intends to buy a boat)	(9b)
and	B(If Chris intends to buy a boat, Chris ought to borrow money)	(9c)
so	B(Chris ought to borrow money).	(9d)

I shall use the term 'normative reasoning' for belief reasoning that concludes in a belief in a normative proposition. The reasoning from (9b) to (9d) is normative reasoning.

We might think that normative reasoning could form a component of a longer process of practical reasoning. Your whole reasoning process might be:

	I(Chris will buy a boat)	(9a)
so	B(Chris intends to buy a boat).	(9b)
Also	B(If Chris intends to buy a boat, Chris ought to borrow money)	(9c)
so	B(Chris ought to borrow money),	(9d)
so	I(Chris will borrow money).	(9e)

I shall call this reasoning by 'normative ascent'. From (9a) to (9b) is the step of ascent, from (9b) to (9d) is a process of normative reasoning, and (9d) to (9e) is the descent step from a normative belief to an intention.

However, this process of reasoning has a fatal flaw. The conditional normative proposition contained in (9c) is not in general true. Even if you intend to buy a boat, it may not be the case that you ought to borrow money. Perhaps you ought not to buy a boat in the first place, for example. I said in section 5 that your intention to buy a boat does not make it the case that you ought to intend to borrow money. Much less does it make it the case that you ought to borrow money.

Since section 5, I have relaxed the assumption that borrowing money is a necessary means of buying a boat, but that makes no difference to this point. Your intention of buying a boat may normatively require you to intend to borrow money. But we cannot detach the conclusion that you ought to intend to borrow money. Much less can we conclude that you ought to borrow money. No detachable normative conclusion is available, and hence no material conditional proposition such as the content of (9c).

The content of (9c) might be true by accident, but it is not generally true, so it could not itself be supported by any general process of correct reasoning. Consequently, (9) cannot serve as a model for practical reasoning.

We can invent variants of (9), in which the content of (9c) is replaced with some other conditional. For example:

> B(If Chris will buy a boat, Chris ought to borrow money)   (9c′)

or   B(If Chris intends to buy a boat, Chris ought to intend to borrow money).   (9c″)

The rest of the reasoning in (9) would need to be adjusted accordingly. But the same objection applies to all these variants. These conditionals with a normative consequent are not in general true, and could not be supported by a general process of correct reasoning.

7. Reasoning and metareasoning

What can be done about that? We might try to accommodate the lesson of section 5 by setting up this different sort of reasoning process, which incorporates normative ascent in a different way:

> I(Chris will buy a boat)   (10a)

so   B(Chris intends to buy a boat).   (10b)

Also   B(Chris's intending to buy a boat normatively requires Chris to intend to borrow money)   (10c)

so   I(Chris will borrow money).   (10d)

In appropriate circumstances, the content of (10c) might be true and supported by a general process of correct reasoning. For example, that might be so if you believe that borrowing money is the best means of buying a boat. So in this respect (10) can serve better than (9) as a model for practical reasoning.

Nevertheless, (10) is misbegotten. To see why, let us take as an example a set of circumstances in which we know already that (10c) is a true belief. It is a true belief in the case I investigated earlier, where you believe that borrowing money is a necessary means of buying a boat. So let us return to that case for a moment.

In that case, we already have correct reasoning that can bring you to an intention of borrowing money. It is described in my original example of practical reasoning, (2). Compared with (2), (10) has at least two dubious features. One arises from the ascent to the normative and the other from the descent from the normative.

First, your intention of buying a boat figures in (10) only by supplying you with the belief that you have this intention. This is the ascent step. Your final intention of borrowing money is supposed to derive from beliefs only, the beliefs (10b) and (10c). This is not plausibly the

way one of your intentions rationally induces another. It does not do so by means of creating in you the belief that you have it, but more directly. In any case, the step from (10a) to (10b) is not plausibly reasoning.

Second, the descent step to the intention (10d) is very hard to fathom. In our earlier formulation (9), the corresponding descent from (9d) to (9e) can easily be understood. It sets out from your belief (9d) in the detached normative proposition that you ought to borrow money. It is plausible that, if you believe you ought to do something, this belief normatively requires you to intend to do it.[12] If so, the step from (9d) to (9e) is normatively required. It is much harder to understand how this intention can be directly derived from (10b) and (10c), without the aid of detachment. And we know that detachment is not available.

By now it should be obvious that (10) is malformed. We already have in (2) an accurate description of correct intention reasoning in the special case we are considering. To say that your intention of buying a boat normatively requires you to intend to borrow money (in the circumstances that you believe this is a necessary means of buying a boat), is merely to say that the reasoning in (2) is correct. It is a remark that belongs to metareasoning, not to reasoning. But in (10) it is injected into the reasoning itself through the belief (10c). That is why (10) is a mess; it is a muddle of reasoning and metareasoning. The correct reasoning (2) does not call on either the belief (10b) or the belief (10c), and it concludes directly in the intention (2c) or (10d). So it does not make the worrying steps of ascent and descent. Ascent and descent are only called on when metareasoning becomes incorrectly entangled with the reasoning. Normative ascent should be no part of instrumental reasoning. Normative reasoning belongs to metareasoning, not to instrumental reasoning itself.

We need metareasoning. We have to determine when reasoning is correct. In Section 2, I argued that one sort of instrumental reasoning, represented by (2), is correct. Its correctness depended on deductive logic. But now we are dealing with means that are not necessary, deductive logic is not going to be enough. You are not entitled to the belief (2b). We need to discover other correct patterns of instrumental reasoning.

There surely are some. When there is a necessary means to your end, the reasoning is clear cut. When there is a means that is not necessary, but still very much the best, reasoning cannot simply leave you in the lurch. By reasoning, you must still be able to arrive at an intention to take this means.

Once again, we can look at belief reasoning for an analogy. Correct belief reasoning can rest on deductive logic, like the reasoning described in (3). But there are other bases for correct belief reasoning, too: induction, perhaps, or inference to the best explanation. For the sake of argument, let us suppose this is correct reasoning.

> B(The road is wet)

and B(The best explanation of the road's being wet is that it has been raining)

so B(It has been raining).

Then this metastatement will be true: believing the road is wet and believing that the best explanation of the road's being wet is that it has been raining normatively requires you to believe it has been raining. The metastatement does not belong to the reasoning; it supports it.

Inevitably, principles of inference other than logic will be more controversial than logic, and I do not insist that inference to the best explanation is a valid principle of inference. It is only an example. There must surely be some valid principles of belief reasoning other than deductive logic.

Whatever these nonlogical principles are, we must expect them to impose weaker normative requirements than logic does. There are degrees of normative requirement. To

infringe logic is a worse offence than to infringe some other principle of inference.

We need to find analogous nonlogical principles for intention reasoning. We need to identify correct reasoning of the form

	I(Chris will buy a boat)	(11a)
and	B(. . .)	(11b)
so	I(Chris will borrow money).	(11c)

Some belief or other must substitute for the dots in (11b). This belief must be such that, together with your intention of buying a boat, it normatively requires you to intend to borrow money. Evidently it must in some way link your buying a boat with your borrowing money. We know already that (11) is correct reasoning when we substitute (4b) – 'For Chris to buy a boat, a necessary means is for Chris to borrow money' – for the dots. In that case, (11) comes down to (2). Now you are no longer entitled to believe (4b), we must find a replacement.

8. Decision theory and conditional goodness
Decision theory is generally supposed to supply an account of instrumental reasoning, so I shall start with decision theory. It will fail in the end, but it will take us a long way, and we can learn from its failure.

We want to know how you can reason about how to buy a boat when borrowing money is not the only way of buying one. Let us make the example more testing. Let us suppose that, although you intend to buy a boat, it would be better if you did not, because you cannot afford to. In these circumstances, it may be that you ought not to borrow money. Nevertheless, you must still be able to derive an intention of borrowing money from your intention of buying a boat. Instrumental reasoning must be able to do that for you. As I said in Section 5, instrumental reasoning cannot break down just because your end is one you would be better off without.

How can decision theory help you? Most often, decision theory is formulated as a theory about the structure of preferences; it describes what structure your preferences ought to have. But I am not going to use it in this form. The reason is that I cannot see how in this form it can contribute to practical reasoning. I cannot envisage any process of reasoning that could take you from a preference to an intention. A preference can cause an intention, but I do not see how this causal process could be one of reasoning. But that is a long story, which I shall not develop here.

Instead I shall simply use decision theory in another form that is more promising. The theory can easily be divorced from preferences, to serve as a theory about the structure of goodness instead.[13] That is how I shall use it.

One purpose of decision theory is to take account of uncertainty. Uncertainty always intervenes between an act and its results; you never know exactly what the result of any act will be. For each act you might do, there is a range of possible results that may emerge if you do it. If you do the act, each of its possible results will have some probability, and each will also be good to some degree.

For the purposes of my argument, it does not matter what is the source or nature of this goodness. It might be impartial universal goodness, or it might be just your own good. We might be valuing the results of your acts from the point of view of the universe, or from your own self-interested point of view, or from some other point of view. We can apply decision theory in any case.

Related to a result's goodness is a quantity that decision theorists call its 'utility'. The utility of a result need not be exactly its goodness, but it *represents* its goodness, which

means that one result is better than another if and only if it has more utility. The distinction between utility and goodness makes no difference to my argument here, and it is safe to ignore it.

Each act has an expected utility, which is defined as the mathematical expectation of the utilities of its possible results. Utility is defined in such a way that, of all the acts that are available to you, the best is the one that has the highest expected utility.

The probabilities of the results of your acts will depend on circumstances. We are supposing that, given the general circumstances of the world, including your financial situation, it would be better for you not to buy a boat. This means that the expected utility of your buying a boat is less that the expected utility of your not doing so. The reason results from the probabilities: if you buy a boat you will probably get into financial difficulties; a high probability is attached to this bad result.

Moreover, let us assume that the expected utility of your borrowing money is less than the expected utility of your not doing so. Again, this results from the probabilities: if you borrow money, that too will probably plunge you into financial trouble. So when we compare borrowing money with not doing so, given the general circumstances of the world, borrowing money is the worse alternative.

However, we are assuming that actually you are going to buy a boat, foolish though that is. This is a further fact to be added to the general facts we have taken into account so far. It makes a difference to the probabilities. We now need to recalculate on the basis of conditional probabilities, conditional on your buying a boat. We must recalculate the expected utilities of your acts, using the new conditional probabilities.

Given the new fact that you are going to buy a boat, then if you do not borrow money you will probably end up selling your house or robbing a bank. Consequently, the expected utility of borrowing money is now greater than the expected utility of not doing so. When we compare borrowing money with not doing so, given the new probabilities that arise from the circumstance of your buying a boat, borrowing money is better. *Conditional on your buying a boat*, it is best to borrow money. That is the conclusion of decision theory.

The suggestion that emerges from all this is that, in place of the dots in (11b), we should put:

Conditional on Chris's buying a boat, it is best for Chris to borrow money. (12)

We get:

I(Chris will buy a boat) (13a)

and B(Conditional on Chris's buying a boat, it is best for Chris to borrow money) (13b)

so I(Chris will borrow money). (13c)

Expressed to yourself, the content is:

I am going to buy a boat

and Conditional on my buying a boat, it is best for me to borrow money,

so I shall borrow money.

This is a decision-theoretic account of instrumental reasoning. Is it a good account?

9. Merits of decision theory

How does the belief that it is conditionally best to borrow money fit into this reasoning – how does it help to justify an intention to borrow money? Some assumption of teleology is evidently implicit. Teleology is the theory that, when faced with a choice, you should choose whichever alternative is the best. To support (13), we shall need a version of teleology appropriately tailored to whatever notion of goodness is incorporated in (13b). If it is

goodness from the point of view of the universe, we must assume impartial teleology; if it is goodness for you only, self-interested teleology.

The appropriate version of teleology allows us to infer from (12) that

$$\text{Conditional on Chris's buying a boat, Chris ought to borrow money.} \qquad (14)$$

We can plausibly understand this to be saying that your intention of buying a boat normatively requires you to borrow money. So it provides metareasoning to justify the reasoning process (13).

Proposition (14) is not a material conditional, so it can provide no support for reasoning like (9), which I have already rejected. The 'ought' in (14) cannot be detached. (14) does not imply the material conditional:

$$\text{If Chris is going to buy a boat, Chris ought to borrow money.} \qquad (15)$$

If it did, it would risk contradiction. From (15) and the fact that you are going to buy a boat, we would be able to deduce that you ought to borrow money. But, given teleology, this is false. I assumed earlier that, unconditionally, borrowing money is worse than not doing so. Consequently, teleology implies you ought not to borrow money.

In decision theory, conditionalized propositions like (14) do not imply material conditional propositions like (15). Conditionalized propositions stem from conditional probabilities, and propositions about conditional probability do not imply material conditional propositions about probability. Conditional probability does not permit detachment.

Here is the demonstration. For some propositions e, f and g, let x be the probability of e conditional on f:

$$\text{Prob}(e|f) = x. \qquad (16)$$

Let y be the probability of e conditional on $f \& g$:

$$\text{Prob}(e|f\&g) = y. \qquad (17)$$

Assume x and y are different, as is certainly possible. Now suppose these conditional probabilities did imply the corresponding material conditionals. From (16) we would have:

$$f \supset \text{Prob}(e) = x.$$

Hence, by strengthening the antecedent:

$$f\&g \supset \text{Prob}(e) = x.$$

But from (17) we would have:

$$f\&g \supset \text{Prob}(e) = y.$$

Since x and y are different, this is a contradiction. Therefore, conditional probabilities do not imply the corresponding material conditionals.

So (14) does not imply (15), and is not subject to the risk of contradiction I mentioned. The upshot is that (13) has some satisfactory features as an representation of instrumental reasoning. It seems to fit what Michael Bratman has in mind as the role of intentions in practical reasoning. Bratman says:

> Prior intentions and plans . . . provide a *background framework* against which the weighing of . . . reasons for and against various options is to take place.[14]

The idea of providing a background framework is well captured by the process of conditioning on prior intentions. In a note, Bratman mentions this decision-theoretic interpretation as a possibility, though he does not endorse it:

> It may be possible to see expected utility theory as an account of how one is to bring to bear, in decisionmaking, utility and probability assignments concerning those options that are relevant and admissible, *given one's background framework of prior plans and flat-out beliefs*.[15]

10. Decision theory is not instrumental reasoning

Despite a good start, decision theory's account of instrumental reasoning is unsatisfactory in the end. One unsatisfactory feature is that it depends on teleology, and genuinely correct intention reasoning should not depend on such a specific normative theory as teleology. This might not be disastrous in itself, but it points us to the real problem with the reasoning described in (13).

The real problem lies in the notion of goodness incorporated in (12) and (13b). We started with some general notion of good. I call it a general notion, but remember is it is not necessarily good from the point of view of the universe; it may be your own good only. From this general good, we determined the comparative goodness of your acts by means of expected utility. We imposed the condition that you are going to buy a boat, and calculated the expected utilities of acts on the basis of conditional probabilities. In the calculation, the probabilities were conditional on your buying a boat, but the notion of goodness was not. It remained as it started: goodness in general, not goodness in any way relative to your buying a boat. We assessed your alternatives of borrowing money or not, according to how good they are, given that you are going to buy a boat. We did not ask how good they are for buying a boat. (12) says that borrowing money is the best thing to do, given that you are going to buy a boat. It does not say that borrowing money is the best way for you to buy a boat.

These are quite different matters. The best thing to do, given that you are going to buy a boat, might not even be a means of buying a boat. For example, it could turn out that the best thing to do, given that you are going to buy a boat, is to join a course in seamanship. But this is not a means of buying a boat.

That is why we had to call on teleology to support the reasoning in (13). We had to presume you are acting for the sake of goodness, not particularly to buy a boat. Your pursuit of good is merely *constrained* by the condition that you are going to buy a boat. Consequently, the reasoning had to call on goodness in general to provide an external objective for you to aim at. The objective of buying a boat was not itself enough to support the reasoning..

Furthermore, once we have identified good as your aim, we can see that something else has gone wrong. If good is your aim, you ought not to condition on buying a boat, and pursue what is best conditional on that. You intend to buy a boat, but you might not succeed. The bank manager might not cooperate, or you might change your mind before you make the purchase. You ought to allow for the possible failure of your intentions. Consequently, if you are pursuing what it best, you should condition on your intention of buying a boat rather than on your actually buying one.

But conditioning on intention takes the decision-theoretic reasoning process even further from choosing an appropriate means to your end. Suppose you intend to kill yourself. Appropriate means might be to slit your wrists in the bath or jump off a cliff. But the best thing to do, conditional on your intention of killing yourself, may be to get help with your problems, so that in due course you will no longer have this intention.

The truth is that, despite some merits, (13) does not represent instrumental reasoning at all. It is not about how best to buy a boat, but about how to do what is best. I conclude that, whatever else it might be, decision theory does not provide an account of instrumental reasoning. It tells you how to achieve good, constrained by your intentions, not how to fulfil your intentions.

11. Correct intention reasoning again

I think we may take it for granted that a satisfactory replacement for the dots in (11b) will be something like

> For Chris to buy a boat, the best means is for Chris to borrow money

or The best way for Chris to buy a boat is to borrow money.

From section 10, we have learnt that the notion of 'best' in these propositions must be in some way relative to your end of buying a boat. Your intention reasoning will go like this, say:

> I(Chris will buy a boat)   (18a)

and B(The best way for Chris to buy a boat is to borrow money)   (18b)

so I(Chris will borrow money).   (18c)

Its content, expressed to yourself, will be:

> I am going to buy a boat

and The best way for me to buy a boat is to borrow money,

so I shall borrow money.

This seems intuitively correct, and a natural extension of (2).

Another example of intuitively correct intention reasoning is:

> I(Chris will commit suicide)

and B(The best way for Chris to commit suicide is for Chris to jump off the cliff)

so I(Chris will jump off the cliff).

This reasoning concludes in an intention to jump off the cliff. If the reasoning is correct, this intention is normatively required by your intention of committing suicide and your belief that jumping off the cliff is the best way of doing so. I am not suggesting you ought to jump off the cliff or that you have a reason to do so. I made it clear in Section 4 that no conclusion about what you ought to do or have a reason to do follows from the relation of normative requirement.

I have given two examples of intention reasoning that seem intuitively correct. I am sorry to say that is now the best I can do. I can point out what seems intuitively correct, but I have run out of arguments. I have left two tasks undone. One is to explicate the notions of 'the best means' or 'the best way', which figure in the reasoning. The other is to explain why reasoning like this is indeed correct.

I recognize these unaccomplished tasks are both major ones. So this paper is really only the beginning of a theory of instrumental reasoning. I hope I have accounted fully for reasoning to a means that you believe is necessary. But about reasoning to a means that you do not believe is necessary, I have not gone beyond negative conclusions. I have argued that normative reasoning is not a correct account of this type of reasoning, and nor is decision theory. A positive account is awaited.

I can mention some of the agenda for the two tasks. So far as the first task is concerned, we know already that the idea of goodness contained in 'the best means to the end' must be relative to the end itself. The goodness of a means to an end is not goodness in general. The best means to an end is not the same as what is best, given the end. For this reason, the correctness of (18) does not depend on teleology; it does not call on the external objective of goodness. The objective is internal to the reasoning. Taking the best means is normatively required by the end itself, not by the pursuit of good.

Undoubtedly, one component of goodness for an end must be reliability in achieving the end. Jumping off a cliff is a good way to kill yourself, because it will probably work. Taking an overdose of aspirins is not a good way, because it often fails. Also, it will often be true that external goodness is a component of goodness for an end, although it cannot be the whole of

it. What is good, given the end, in many cases plays a significant part in determining what is the best means to the end. The best means of buying a boat is partly determined by the effort and cost involved. Saving effort and money is good in a general sense; its goodness does not issue from your intention of buying a boat. So we must properly integrate external goodness with goodness specific to the end. That is part of the first task I mentioned.

A problem within the second task is this. I started this paper by arguing that intention reasoning and belief reasoning are parallel, and both follow truth. I made that argument in Section 2, for the special case of a necessary means. In this special case, both types of reasoning are supported by deductive logic. What about the more general case?

As it happens, belief reasoning parallel to (18) is plausibly correct:

	B(Chris will buy a boat)	(19a)
and	B(The best way for Chris to buy a boat is to borrow money)	(19b)
so	B(Chris will borrow money).	(19c)

This reasoning is defensible. It is supported by an inference to the best explanation. If you are going to buy a boat, the best explanation of how this will happen is that you will borrow money. So you might infer that you will borrow money. Any instance of intention reasoning to the best means like (18) will be parallelled by belief reasoning to the best explanation like (19). However, the parallel is not robust. (19) constitutes reasoning to the best explanation only *because* (18) is reasoning to the best means, and because you may be expected to reason to the best means. If you were irrational, (19) would not be defensible reasoning.

I think the truth is that, when we move away from necessary means, intention reasoning and belief reasoning diverge. Both still follow truth, one in a truth-making way and the other in a truth-taking way. But the former is concerned with the *best* way of making the end true, and the latter with the *most likely* way the end will be true. When there is only one means, and so only one way the end will be true, these concerns collapse into one, but otherwise they do not. The case when there is just one means is also the case where both types of reasoning are supported by deductive logic.

I do not think this should cast doubt on the conclusion of Section 2. I still think both types of reasoning follow truth. Because of this, they are supported by the same deductive logic, when they are supported by deductive logic at all. That is what Section 2 claimed. But part of the second task – giving an account of correctness for intention reasoning in general – will be to check my claim that intention reasoning, like belief reasoning, always follows truth.

12. *Summary*

Like all reasoning, practical reasoning is a process that takes you from some of your existing mental states to a new one. Theoretical reasoning concludes in a belief; practical reasoning in an intention. If a piece of reasoning is correct, it concludes in a state that is normatively required by the states it is derived from. But it does not follow that you ought to be in the concluding state, or have a reason to be in it, even if you are in the states it is derived from.

This paper considered only one sort of practical reasoning: instrumental reasoning. If you intend an end, from this intention together with an appropriate belief, instrumental reasoning leads you to intend a means to the end.

The simplest cases of instrumental reasoning are those in which you believe a particular means is necessary to the end. In those cases, the correctness of the reasoning is ensured by the logical validity of its content.

More commonly, you will not believe that the means is necessary to the end, but instead that it is the best means to the end. In those cases, your reasoning does not rest on logical

validity; it requires some other principle of reasoning. This principle remains to be worked out. But instrumental reasoning is definitely not to be interpreted as normative reasoning. Nor does decision theory supply an appropriate principle of instrumental reasoning.

Notes

1. In his paper 'On so-called practical inference' (p. 53 in the reprinted version), G. H. von Wright claims that statements of the form (2) do actually constitute a valid inference. I understand the temptation to think so: if you believe borrowing money is necessary for your buying a boat, but you do not intend to borrow money, it is tempting to think you cannot *really* intend to buy a boat. But if that were so, there would be no room for practical reasoning; there would be no opportunity for forming an intention by reasoning like (1). If you had the intention expressed in (1a) and the belief expressed in (1b), you would necessarily, even without reasoning, have the intention expressed in (1c). Because people sometimes fail to reason properly, (2) cannot be an inference.

2. Hume, *Treatise*, Book 2, Part 3, Section 3.

3. My views about intention and practical reasoning are similar to Bruce Aune's, set out in his *Reason and Action*. See particularly pp. 156–8. Aune recognizes that intention reasoning and belief reasoning can share the same logic. But he seems not to have recognized this is because they share the same content. As a result, he does not recognize that belief reasoning and intention reasoning are both concerned with truth, and both track truth in their own way. Instead, he assigns a special sort of valuation to intentions, different from truth and falsity.

My views also resemble Richard Hare's in some respects. In several of the papers reprinted in his *Practical Inferences*, Hare argues that imperatives share the same logic as indicatives, even though they are expressed in different moods. However, Hare's views and mine differ greatly. I am not concerned with imperatives. Moreover, Hare embeds imperatives within sentences, and that exposes him to the problem raised by Peter Geach's in 'Assertion'. I do nothing similar.

4. Geach, 'Assertion'.

5. Frances Kamm, 'The doctrine of triple effect'.

6. To avert confusion: my notion of normative requirement is not what Roderick Chisholm calls 'requirement' in his 'Practical reasoning and the logic of requirement'.

7. pp. 23–7.

8. This argument is developed more thoroughly in my 'Are intentions reasons?'.

9. This definition of the instrumental principle is implicit on p. 215. Although I have eventually concluded that Korsgaard's argument is mistaken, her paper was a major stimulus for this paper of mine.

10. p. 220.

11. pp. 245–6.

12. This idea is investigated in my 'Normative practical reasoning'.

13. See my *Weighing Goods* and the Introduction to my *Ethics Out of Economics*. In his *Logic of Decision*, Richard Jeffrey treats decision theory as a theory of 'desirability', which I take to be synonymous with goodness.

14. *Intentions, Plans and Practical Reason,* pp. 33–4. Italics original.

15. p. 180, n 12. Italics original.

References

Aune, Bruce, *Reason and Action*, Reidel, 1977.

Bratman, Michael E., *Intention, Plans and Practical Reason*, Harvard University Press, 1987.

Broome, John, 'Are intentions reason? And how should we cope with incommensurable values?', in *Practical Rationality and Preference: Essays for David Gauthier*, edited by Christopher Morris and Arthur Ripstein, Cambridge University Press, 2001, pp. 98–120.

Broome, John, *Ethics Out of Economics*, Cambridge University Press, 1999.

Broome, John, 'Normative practical reasoning', *Proceedings of the Aristotelian Society*, Supplementary Volume 75 (2001), pp. 175–93.

Broome, John, 'Normative requirements', *Ratio*, 12 (1999), pp. 398–419. Reprinted in *Normativity*, edited by Jonathan Dancy, Blackwell, 2000, pp. 78–99.

Broome, John, *Weighing Goods*, Blackwell, 1991.

Chisholm, Roderick, 'Practical reasoning and the logic of requirement', in *Practical Reason*, edited by S. Körner, Blackwell, 1974, pp. 2–13, reprinted in *Practical Reasoning*, edited by Joseph Raz, Oxford University Press, 1978, pp. 118–127.

Geach, Peter, 'Assertion', *Philosophical Review*, 74 (1965), pp. 449–65.

Hare, R. M., *Practical Inferences*, Macmillan, 1971.

Hume, David, *A Treatise of Human Nature*, edited by L. A. Selby-Bigge and P. H. Nidditch, Oxford University Press, 1978.

Jeffrey, Richard C., *The Logic of Decision*, Second Edition, University of Chicago Press, 1983.

Kamm, Frances, 'The doctrine of triple effect and why a rational agent need not intend the means to his end', *Proceedings of the Aristotelian Society*, Supplementary Volume 74 (2000), pp. 21–39.

Korsgaard, Christine, 'The normativity of instrumental reason', in *Ethics and Practical Reason*, edited by Garrett Cullity and Berys Gaut, Oxford University Press, 1997, pp. 215–54.

von Wright, G. H., 'On so-called practical inference', *Acta Sociologica*, Vol. 15, No. 1, pp. 39–53, reprinted in *Practical Reasoning*, edited by Joseph Raz, Oxford University Press, 1978, pp. 46–62.