

IX—AGAINST REQUIREMENTS OF RATIONALITY

A. W. PRICE

Are inferences, theoretical and practical, subject to requirements of rationality? If so, are these of the form ‘if ... ought ...’ or ‘ought ... if ...’? If the latter, how are we to understand the ‘if’? It seems that, in all cases, we get unintuitive implications (often involving bootstrapping) if ‘ought’ connotes having reason. It is difficult to formulate such requirements, and obscure what they explain. There might also be a requirement forbidding self-contradiction (not that one’s current beliefs can be consciously contradictory). It is a good question whether self-contradiction constitutes, or evidences, irrationality; but talk of a rational requirement causes trouble.

I

The following is a valid inference-schema:

- (1a) If p , then q .
- (1b) p
- So, (1c) q

So is the following, which is more complex:

- (2a) I will achieve end E .
- (2b) I can’t achieve end E without realizing means M .
- So, (2c) I will realize means M .

(1a–c) is a valid schema for pieces of theoretical reasoning that will be sound if they lead from truth to truth. (2a–c) is that, but also a schema for intelligible pieces of practical reasoning leading towards intentional action.

Saying this is as yet saying nothing—at least, nothing explicit—for or against actual or possible pieces of reasoning that are valid or intelligible in these ways. Hence it may seem insufficient either for

assessing or for explaining certain pieces of reasoning as *good* pieces of reasoning. A tradition going back at least to W. K. Clifford's essay 'The Ethics of Belief' (1877) suggests that we are subject to what may be called *doxastic duties*. Thus he writes, in a moralizing style, 'It is wrong, always, everywhere, and for anyone, to believe anything upon insufficient evidence.' This principle prescribes that one believe that *p* only if one has sufficient evidence to do so. Yet this may seem too restricted: what if a subject *refuses* to form a belief though sufficient evidence stares him in the face? And perhaps we also need prescriptions governing logical inferences. These might not only forbid us certain inferences as invalid, but require others of us, if only conditionally, as valid.¹ If such requirements are requirements of *rationality*, we should have reason—and perhaps conclusive reason—to observe them. It may seem that in some way, to the extent that we are rational, we are *bound* to try to reason so.

Yet how might this be captured and conceived? I wish to argue that it is misguided to apply such a strategy to deductive and practical inferences. (The formation of beliefs on the basis of evidence or induction is not my topic. Hence this paper is less general than its title.) There are difficulties both in formulating such requirements, and in finding them an explanatory role. So much I have argued elsewhere (Price 2008, pp. 72–8). What I shall especially argue here is that *if* the requirements constitute reasons for subjects (and why else should we attend to them?), they license derivative ascriptions of reasons to subjects in a manner that is bootstrapping.

II

Where might we start a search for requirements to help us? Perhaps it is inviting to turn to conditional 'ought's, and to try to define what else people *ought* to believe and intend *if* they believe or intend certain things. Here the range of options is wide, and in two respects. First, there are differences of scope. Is it conditionally upon his being in a certain mental state that *A* ought to believe or intend certain things? Or is it rather that this is how *A* ought to be: believ-

¹ Thus John Broome puts to me that we need requirements upon agents that 'make certain patterns of reasoning correct'.

ing or intending certain things conditionally upon his being in a certain mental state? Not that we should be quick to presume that we know what the second means. One can believe a conditional of the form of (1a); but what is it to believe one thing conditionally upon believing another?

Secondly, there are variations in the force of 'ought'. How *A* ought to act or be is a matter of what is *fitting* or *appropriate* for him—by which I mean either *uniquely* or *best* fitting. But fitting or appropriate *to what?* Practical 'ought's are commonly relative to an end, and a set of circumstances. They have less force if the set of circumstances leaves much that matters open. Their force is also a function of the standing of the end: some ends are imperative, others optional, others downright undesirable. It can be true that, in a sense, *A* ought to ϕ , because ϕ -ing fits a goal of his, even if there is nothing worthwhile in his achieving that goal. ('I intend to get drunk every night.' 'Then you ought to buy a distillery or work in a pub.')

Some 'ought's invite scepticism: 'Really?', we respond. Others invite indifference: 'So what?', we reply. We may distinguish a *minimal* from a *maximal* 'ought'. If *A* minimally ought to ϕ , his ϕ -ing is fitting to something or in some way, but nothing follows about whether he has any reason to ϕ . If *A* maximally ought to ϕ , he has decisive reason to ϕ . If we want to explain the force of rationality, it is surely 'ought's connoting reasons that may interest us.

Suppose that *A* believes (1a) and (1b): inasmuch as this is the case, does he then have reason to believe (1c)?² (By 'have reason' I mean 'have an undefeated reason'. This falls within a range that descends from 'have decisive reason' to 'have best reason', to 'have an undefeated reason', to 'have *a* reason'.³) For what end? Perhaps in order to meet a requirement of rationality. Practical reasoning is more complex, but this might hold: supposing that *A* intends (2a) and believes (2b), he has reason to intend (2c).

Strictly, of course, this is too simple: we need to specify times. Believing (1a) and (1b) at *some* time can't give *A* reason to believe (1c) at *all* times. More precisely, imagine that *A* is reflecting, at time t_1 , what to believe at time t_2 : he might well say to himself, 'Suppose

² Strictly, of course, I should speak of *A*'s believing things of the form of (1a) and the like; but that would become cumbersome.

³ Cf. Raz (2005a, p. 7).

that I believe ($\text{I}a$) and ($\text{I}b$) at t_2 : then I ought also to believe ($\text{I}c$) at t_2 .⁴ It may seem a problem that it takes time to conduct an inference. Yet the problem is superable: accepting certain premisses at t_1 , I come to draw a conclusion at t_2 —while still accepting the premisses. However, I shall generally disregard this complication.

So the following is a possible principle:

(I) If A accepts ($\text{I}a$) and ($\text{I}b$), he ought to accept ($\text{I}c$).⁵

Here, 'if' has wider, and 'ought' narrower, scope. How plausible is (I)? There is no problem if its 'ought' is minimal: it *fits* my accepting ($\text{I}a$) and ($\text{I}b$) that I accept ($\text{I}c$). But this says nothing about whether I have any reason to accept the conclusion. If the 'ought' is maximal (or even if it just connotes having reason or a reason), (I) is less plausible. It actually distorts what can be expected of logical inference, which is something else. Given that ($\text{I}a$) and ($\text{I}b$) together entail ($\text{I}c$), a reason for taking ($\text{I}a$) and ($\text{I}b$) to be true will often be a reason for taking ($\text{I}c$) to be true, so that the grounding of the premisses is *transmitted* to the conclusion.⁶ The same holds, *mutatis mutandis*, of derivative intentions: a reason that A has to intend ($2a$) and believe ($2b$) may well be a reason for him to intend ($2c$).⁷ Yet it is questionable whether believing ($\text{I}a$) and ($\text{I}b$) is *itself* reason for believing ($\text{I}c$), or that intending ($2a$) and believing ($2b$) is *itself* reason for intending ($2c$). Surely this would amount to *bootstrapping*. Suppose that it is *crazy* of a subject to believe or intend the premisses: precisely to the extent that his acceptance of the conclusion rests upon his acceptance of the premisses, it will then also be crazy of him to believe or intend the conclusion. It isn't the task of inference to ground conclusions irrespectively of the grounding of

⁴ There is no call, and it would make no sense, to add a time-reference also to 'have a reason'.

⁵ Alternatively, (I) might state that A ought to *infer* ($\text{I}c$), which, in Niko Kolodny's terminology (2005, pp. 517–18), would make (I) rather a *process* than a *state* requirement.

⁶ Not that this holds universally. Reasons for beliefs, as for actions, are relative to sets of alternatives. If p and q are alternatives, and so are r and s , and p entails r whereas q entails s , a reason for believing p rather than q becomes by transmission a reason for believing r rather than s . But if, for instance, both p and q entail r , a reason for believing p rather than q will only *per accidens* yield a reason for believing r rather than s . (Dorothy Edgington alerted me to this. Elsewhere in this paper, when I talk of a reason for believing p I mean a reason for believing p rather than not believing p .)

⁷ However, the previous note applies here, *mutatis mutandis*. And the special nature of intention creates further exceptions; see §v, para. 4.

the premisses. In certain cases, (1a) is an axiom or a framework-proposition, or end *E* is self-recommending, so that they need no grounding. However, it is still not *A*'s *accepting* (1a) or (2a) that may ground his accepting (1c) or (2c).⁸

Concessive responses to this argument are not likely to help us. Beliefs and intentions may ground derivative beliefs and intentions *defeasibly*. Perhaps believing (1a) and (1b) is generally reason to believe (1c), and yet, if it is ungrounded but in need of grounding, no reason at all.⁹ Alternatively, it is the other reasons that are defeasible: in cases where the inference articulates the only way in which *A* can justifiably come to believe (1c), his reasons for believing (1a) and (1b) may only be reasons *for him* to believe (1c) if he believes (1a) and (1b). On either view, (I) could be amended as follows:

(1a) If *A* accepts (1a) and (1b) with reason, he ought to accept (1c).

Yet as a principle governing inferences (1a) would permit a surplusage of irregularity. *A* first believes (1a) and (1b) though he is aware that he lacks grounds (which is a candour uncommon but surely conceivable). He then, without discarding those beliefs, declines to infer (1c) from them since they are ungrounded. Which is not an intelligible option.

⁸ The bootstrapping objection may seem most evident (as John Broome, 1999, p. 405, has observed) in the simplest entailment of all, from *p* to *p*: it can't be that believing that *p* gives one reason to believe that *p*. This would be an Indian rope trick, with a belief sustaining itself. However, this may be ruled out not for the special reason that one belief is no reason for another just because its content entails the content of the second, but for the general reason that, while belief that *p* is indeed grounded in whatever grounds belief that *p*, nothing can ground itself. In the case of 'ought's, no belief can be fitting or appropriate to itself; rather, of course, belief that *p* fits whatever belief that *p* fits.

However, time-references make a difference. It is nonsense to say that believing that *p* at *t*₁ gives one reason to believe that that *p* at *t*₁. However, it might be that believing that *p* at *t*₁ gave one reason to continue to believe that that *p* at *t*₂—though this would depend upon contingencies. The reason could be truth-regarding: Christian Piller (unpublished) cites the case of a man who knows that his beliefs about a certain subject-matter, to which *p* belongs, have proved themselves to be reliable. Then his believing that *p* at *t*₁ can become part of his reason for continuing to believe that *p* at *t*₂, and the continuing belief will then be more richly grounded than the initial one. Or a contextual cost may attach to a change of mind, whether intrinsic (say, if it would count as one further piece of fickleness on the part of an inconstant subject), or consequential (say, if he has already invested resources in a certain belief). However, one can't form a belief on the ground of a non-truth-regarding reason; so these reasons are better viewed as reasons for, e.g., declining to reopen a question, even if retaining one's present beliefs is a foreseen consequence.

⁹ Such a view is advanced in Piller (unpublished).

III

We need to try again. It may seem that what we need is an ‘ought’ that, being maximal, is neither overridable nor defeasible, within a construction that doesn’t permit any inference from actually accepting a set of premisses to being right to accept the conclusion. This is precisely what we are offered by John Broome’s conception of ‘normative requirements’, which are requirements of rationality with a distinctive structure. What if we invert the scopes of ‘if’ and ‘ought’ within (I)? Using brackets for disambiguation, we might have this:

(II) *A* ought (to accept (1*c*) if he accepts (1*a*) and (1*b*)).

Broome (1999) recommends this shift in scope expressly in order to preclude detachment: (I) permits us to infer ‘*A* ought to accept (1*c*)’ given that *A* accepts (1*a*) and (1*b*), whereas (II) doesn’t.

This permits us to intensify the force of the ‘ought’ without licensing even less plausible inferences than those already rejected. The ‘ought’ of (I) seems unmerited even if it is *slack* (Broome 2001, p. 106), even, that is, if it is overridable or defeasible. Yet the ‘ought’ of (II) may be *strict* (Broome 2001, p. 106): if *A* accepts (1*a*) and (1*b*) without accepting (1*c*) he is definitely *out of order*, and not as he is *required* to be—though the right remedy may be to avoid or cease accepting (1*a*) or (1*b*), and not to accept (1*c*).

How does this connect with what *A* has *reason* to do? Here Broome is agnostic: he thinks it an open question whether we have reason to observe normative requirements. (For this reason, he now prefers not to express them as ‘ought’s: see Broome forthcoming.) This may puzzle us: How then are they requirements of *rationality*? And what is their point and authority? Of course, the phrase ‘requirements of rationality’ is ambiguous (as is talk of the requirements of morality, or of etiquette): are the requirements of rationality just necessary conditions for counting as ‘rational’, or are they requirements about how to think that rationality somehow places upon us?¹⁰ If they are the former, they should yield hypothetical imperatives that one has reason to regard to the extent that one aims to be rational. If they are the latter, but require things of us without our having any reason to comply, they surely merit a deaf ear. In either case, the concept of *rationality* is detached from that of *hav-*

¹⁰ Cf. Kolodny (2005, pp. 555–6), Broome (2007, pp. 361–2).

ing a reason. Yet little is explained if being rational is allowed to be optional. *A* may indeed accept (1*a*) and (1*b*), and yet irrationally fail to infer (1*c*); yet surely he can't evade an invitation to infer (1*c*) by *opting* to be irrational.

So I think that if (II) is to do any work, its 'ought' must rather be interpreted as maximal and also categorical: it ascribes to *A* a reason that is *peremptory* in that it can neither be disregarded at will, nor weighed in the balance against other reasons. And if (II) is *itself* to be explanatory as a *requirement*, it isn't enough that it *tracks* reasons (so that *A* always has reason to comply with it): it must *constitute* a reason (so that it is *its requiring* compliance of *A* that gives him reason to comply). Yet we all often fall short of what such a rule prescribes, and in ways that are venial. And its authority can't be conceived as free-standing—rather as if God's law had prescriptive force independently of his existence, like the smile without the Cheshire Cat. It must be a law of our own reason, and yet one that our reason observes not (as it were) spontaneously, but *out of respect* for it—as with Kant's Categorical Imperative. So it is a puzzling conception, about which much more would need to be said. (It isn't for me to anticipate what.)

I take it, then, that (II) aspires to constitute a peremptory reason for *A* to think in certain ways. How are we to interpret its content? There is a danger here. Anyone who introduces it precisely *in order to* block detachment isn't going to treat it as an open question, still to be investigated, whether its logical grammar invites something *like* detachment. (What is *strictly* detachment is precluded by the scope of the operators.) Here I can only state a position that I argue elsewhere.¹¹

Take the general form of a practical 'ought' with a conditional content: '*A* ought (to ϕ if *p*).' How might we infer '*A* ought to ϕ '? Certainly not within the scope of the supposition 'It ought to be that *p*.' For it may be that *A* has no call to ϕ unless it is *actually* the case that *p*. Nor, I concede, simply through supposing that *p*. What I take the sentence to predicate of *A* via the modality of 'ought' is ϕ -ing *conditionally upon its being the case that p*. Hence, once it is given that *p*, it follows that *A* ought to ϕ *because p*—if I may use the term in an unidiomatically weak sense that simply excludes mere coinci-

¹¹ See Price (2008, ch. 3, §VI). Here I define my proposal slightly differently, and add, in illustration, a brisk application to 'if's and 'can's.

dence. Now suppose that the closest possible world to this one in which A ϕ 's *when* p is a world in which he ϕ 's *because* p (in the same weak sense): then and only then, as it seems to me, we can infer 'A ought to ϕ ', given p . To distinguish this, I call it 'quasi-detachment'. Take a clear example, actually of an attitude rather than an action. Suppose that A ought to feel kindly towards B if B treats him as a friend, as B does: it will follow that A ought to feel kindly towards B , so long as, if he does, this will be at least in part *because* B treats him so. If A feels kindly simply out of habit or good nature, he won't be responding as he ought to B 's treating him so.¹²

A once-debated context with which this connects is the logic of 'if's and 'can's. Take 'I can ϕ if I choose.' This 'if' may be Austinian: 'I can ϕ —but whether I choose to do so is an open question.' Or it may connect two propositions: 'If I choose to ϕ , I am able to apply myself successfully to the task of ϕ -ing—as I am not if ϕ -ing is imposed upon me against my will.' But it can also, I suggest, mean this: one thing that I can do in ϕ -ing is act *upon* a choice to ϕ . It holds of certain values of ' ϕ ' that one can ϕ , though one can't do *that*: some things one can do easily enough (forget a familiar name, overlook a relevant consideration, make a remark on the spur of the moment), but not at will, i.e. not in execution of a choice or decision. These are not things that one can *do if one chooses*.

In the light of this, let us reconsider (II), 'A ought (to accept (IC) if he accepts (IA) and (IB)).' Suppose that the closest world in which A accepts (IC) *when* he accepts (IA) and (IB) is one where he accepts (IC) *because* he accepts (IA) and (IB). Then it seems to me to follow from (II), given that A does accept (IA) and (IB), that he ought to accept (IC). Yet this may well be unintuitive if 'ought' here connotes having reason, and he *oughtn't* to accept (IA) and (IB). The moral, I think, isn't that quasi-detachment fails, but that (II) is little more acceptable than (I) if their 'ought's connote having reason.

¹² This is an adaptation to a different grammatical form of P. F. Strawson's thought (1997) that, within 'If p , q ', 'if' expresses not just a truth-functional relation, but the relation of a *ground* to a *consequence*. He compares the role of 'so' within ' p ; so q .' (What I really need for my adaptation of this is a sense of 'because' that is not more specific than that of 'so'.)

Note two things. First, my 'because' excludes mere coincidence without requiring that the subject or agent consciously respond *for a reason*: A can count as ϕ -ing because p even if he isn't consciously aware that p , or doesn't take the fact that p as his reason. (It may not even be a cause: it might just be a corollary of a cause.) Secondly, 'if' doesn't mean *only if*. It may be that A ought to ϕ if p , and also if q . Suppose that p is true at t_1 , p and q at t_2 , and q at t_2 : A 's ϕ -ing can be a response at t_1 to p , at t_2 to q , and at t_2 to *both*. Thus he may be complying in ϕ -ing with 'A ought to ϕ if p ' even if he would also be ϕ -ing if it wasn't the case that p .

IV

Yet it is unclear how damaging this is to Broome, for two reasons. I have already noted that he is himself uncertain whether there is reason to comply with normative requirements (and hence prefers not to state them as 'ought's). In any case, he interprets (II) differently. I read it as instantiating the more general form 'A ought to ϕ if p ', which I take to relate A , by way of 'ought', to ϕ -ing conditionally upon its being the case that p . To such a schema, so understood, contraposition has no application. We can't have 'A ought not p if not ϕ ', for p isn't verbal and ϕ isn't propositional. We might try to circumvent this by a cumbersome paraphrase: 'A ought to make it the case that not p if he doesn't ϕ .' But this has two defects. First, it is free and ad hoc. Secondly, it differs from 'A ought to ϕ if p ' in its presuppositions: 'A ought to ϕ if p ' presupposes that, supposing that p , it is in A 's power to ϕ , whereas 'A ought to make it the case that not p if he doesn't ϕ ' would have to presuppose that, supposing that A doesn't ϕ , it is in his power whether p . Since neither presupposition entails the other, the two sentences can't be logically equivalent. Broome, however, intends (II) to be playable in two ways: as he means it, it requires A either to accept (Ic), or not to accept (Ia) and (Ib). Either would save him from accepting the premisses of a valid argument while failing to draw its conclusion.

Thus Broome reads (II) as roughly equivalent to this:

(IIa) A ought either to accept (Ic) or not to accept (Ia) and (Ib).¹³

Now suppose, as Broome leaves open, that A has reason to comply with (IIa). Then we must interpret it as offering A two alternative means to some implicit goal (presumably the avoidance of some kind of incoherence) that he can achieve either by accepting (Ic) or by not accepting (Ia) and (Ib). Though Broome contests this (and views what I shall infer from it as a *reductio*), I take it that, if A has reason (that is, an undefeated reason) either to ϕ or to χ , for the

¹³ In Broome's view (1999, p. 400), (II) entails (IIa), but further indicates that the two disjuncts stand in a relationship of relevance that is not purely truth-functional. Apparently this doesn't affect *what* is required of A . (Otherwise, there would be a logical fallacy: if p is stronger than q , p entails q , but Op doesn't entail Oq .) I shall take it that (IIa) is an adequate statement of *that*. Supposing otherwise would complicate the presentation, though not the substance, of my argument.

sake of some goal, then, for its sake, he has a reason to ϕ that is also a reason to χ , even if he has no reason to do *both*. I would add that the inference is defeasible: if it is to follow that *A* has a reason to ϕ , it must not already be given that he will χ if he either ϕ 's or χ 's. (This covers the case—which itself excludes his having a reason to ϕ —that he can't ϕ . It implies that, if he can and does ϕ , this can only, for the purpose of ϕ -ing or χ -ing, be superfluous.)¹⁴

This lands (IIa) with an unexpected corollary. It turns out to imply that, so long as it is at least open that *A* may accept the premisses of a possible inference, he has a reason to accept its conclusion, whatever this happens to be. Which is surprising.¹⁵

Further, given (IIa), it is hard to escape the following inferences,

¹⁴ Contrast Broome (2005, pp. 6–7) with Raz (2005a, pp. 12–13). Admittedly, Raz (2005b, pp. 1–2) appears to change his view, holding that, while 'A has reason to bring it about that *p* or *q*' entails 'A has a reason to bring it about that *p*', 'A has reason to ϕ or χ ' doesn't entail 'A has a reason to ϕ .' Which seems doubly mistaken. A may have reason to ensure that *p* or *q*, but no reason to ensure that *p*, or to ensure that *q* (either of which might intrusively spoil things—suppose that choice *between p* and *q* must be left to others). When Raz argues that 'A has reason to ϕ or χ ' is entailed by 'A has reason to ϕ ', he appears to confuse the former with 'A has reason to ϕ , or A has reason to χ .' He had previously and more wisely rejected 'a rule of inference allowing introducing a disjunction within the scope of a modal operator' (Raz 2005a, p. 12n.).

What Broome enjoyably calls my 'extravagant distribution of reasons over disjunction' is helpful in Buridan's ass cases. Suppose that *A* has reason to ϕ , which he can achieve by ψ -ing or ξ -ing, each of which is otherwise indifferent. It needn't worry us that he won't ψ if he is hyper-rational (a term I owe to Broome) to the extent that he will never intentionally ϕ by ψ -ing rather than by ξ -ing if he can give no reason for ψ -ing rather than ξ -ing. Will he also never ϕ by ψ -ing if he is rational to the degree that he won't intentionally ϕ by ψ -ing if he can give *no* reason for ψ -ing? I find this inference less palatable (though Broome responds by counting this agent too as hyper-rational rather than rational). So we should accept that, having reason to ϕ , and hence to ψ or ξ , he has a reason—indeed, sufficient reason—to ψ .

However, Broome now permits such a distribution of reasons when the alternatives are equal or indifferent. I am doubtful of this further restriction, and would rather say that one can have a reason to ψ (*viz.*, as a way of ϕ -ing) although one has *better* reason to ϕ by ξ -ing. (In fact, I agree with 'satisficers' that one can have sufficient—and so undefeated—reason to ψ even though one has better reason to ξ instead; for the *better* need not be the enemy of the *good enough*. We surely don't want to preclude that here.) If Broome's restriction holds, the argument that follows in my text loses some but not all of its force: (IIa) will imply that *A* has a reason to accept the conclusion of any inference whose premisses he may accept, and has no undefeated reason not to accept. Which remains surprising, I think.

Note, by the way, that while *A* may well have no reason to ϕ and χ if he has reason to ϕ or χ , he may have a reason to ψ , which hangs together as a way of ϕ -ing and χ -ing, and hence also of ϕ -ing or χ -ing. Relevant here is Anscombe (1995, p. 13).

¹⁵ This isn't an objection to Broome's current view, which leaves the relation of normative requirements to reasons open. The logic of 'be required' is different from that of 'have reason'. If *A* has reason either to ϕ or to χ , then (I think) he defeasibly has a reason to ϕ , and a reason to χ . But, if he is required either to ϕ or to χ , he is neither required to ϕ , nor required to χ . Even if it is true, say, that he can't help not ϕ -ing but can avoid not χ -ing, it doesn't follow that he is *required* to χ . We could then say that he *has* to χ —which Broome need not resist if this doesn't entail that he *has any reason* to χ .

whatever the force of the 'ought'. Suppose that, in fact, *A* accepts (*Ia*) and (*Ib*). This may be fixed or mutable. If it is fixed, then, as it seems to me, it clearly follows that he ought to accept (*Ic*). After all, he ought to comply with (*IIa*); he *had* two ways of doing so (not accepting (*Ia*) and (*Ib*), or accepting (*Ic*)), but now only one way is open to him (accepting (*Ic*)). If he accepts (*Ia*) and (*Ib*) but might conceivably be persuaded not to, then it isn't unqualifiedly that we can say that he ought to accept (*Ic*); but we can say it until and unless, improbably, he changes his mind. (It doesn't follow that the 'ought' is really conditional in content, so that there is nothing to retract.)¹⁶ In either case, if we relativize the 'ought' to a set of circumstances, he ought to accept (*Ic*) relatively to his accepting (*Ia*) and (*Ib*). We can equally infer, supposing that he doesn't accept (*Ic*), that he ought, relatively to that, not to accept (*Ia*) and (*Ib*). Yet these are inconvenient claims, since they amount to blatant bootstrapping if the 'ought' is maximal; for believing (*Ia*) and (*Ib*) then gives him decisive reason—relatively to as inclusive a set of circumstances as one cares to specify—to believe (*Ic*).

Note that '*A* ought to ϕ ' can't mean simply that *A* will be out of order unless he ϕ 's. (It doesn't follow, if I assure you that I didn't do it, that I oughtn't to have done it, though only then will I be in order.) It rather alludes to some set of ends and circumstances relatively to which it is fitting that he ϕ . If my argument is right, (*IIa*) implies that, if *A* accepts (*Ia*) and (*Ib*), then he ought, relatively to that, to accept (*Ic*), and that, if he doesn't accept (*Ic*), then he ought, relatively to that, not to accept (*Ia*) and (*Ib*). If the 'ought' is maximal, then, if *A* is subject to (*IIa*), he has decisive reason (even taking all relevant circumstances together) to accept (*Ic*), or not to accept (*Ia*) and (*Ib*).

This fits a feature of my original schema, 'If *p*, then *q*; *p*; so *q*.' Though this inference reasons towards *q*, the same logical relations that it exploits also lie behind the converse schema, 'If *p*, then *q*; $\neg q$; so $\neg p$ ', which reasons towards $\neg p$. So if we are looking for a corollary of those logical relations among normative requirements, we may prefer a requirement that can also be played by way of *modus ponens* or *modus tollens*. If *A* respects (*IIa*), he will take accepting (*Ia*) and (*Ib*) as reason for accepting (*Ic*), rejecting (*Ic*) as reason

¹⁶ Compare this: we make perceptual claims qualifiedly, since we might conceivably be hallucinating; it doesn't follow that *what we claim* is really conditional. And the range of things that might upset our expectations cannot be specified in advance.

for rejecting (1a) or (1b), and not accepting (1c) as reason for not accepting (1a) and (1b). Though this last is not an instance of *modus tollens*, it connects with a possible instance: if A leaves it open whether (1c) is false, he leaves it open whether one could infer from its falsity the falsity of (1a) or (1b); so he should not accept (1a) and (1b). In these ways (IIa) still permits A to bootstrap, which is bad, but in both directions, which seems apt.

However, it is problematic whether this carries over to practical reasoning. Consider the analogues of (II) and (IIa) for practical reasoning from intention to intention. Analogous to (II) is this:

(III) A ought (to intend (2c) if he intends (2a) and accepts (2b)).

Like (II), this is subject to quasi-detachment. Analogous to (IIa) is this:

(IIIa) A ought either to intend (2c), or else either not to intend (2a) or not to accept (2b).

(IIIa) identifies a variety of ways in which an agent may count as respecting Kant's principle, 'Whoever wills the end also wills (in so far as reason has decisive effect upon his actions) the indispensably necessary means to it that are within his power' (Kant 1997, 4:417). However, the parallelism of (IIa) and (IIIa) may be suspect. Accepting (1a) and (1b), A may come to accept (1c) (if he doesn't already). Not accepting (1c), he may come not to accept (1a) and (1b). Equally, A may come not to accept (1a) if he accepts (1b) but not (1c), or not to accept (1b) if he accepts (1a) but not (1c). But (2a-c), when it serves for practical reasoning and not just for theoretical reasoning, is more complex. Can one reason intelligibly from not intending (2c)?

Practical inferences differ from theoretical in not being reversible from *modus ponens* into *modus tollens*. Take (2a) and (2b) in turn. Given that one accepts (2b), 'I will achieve end E; so, I will realize means M' is intelligible as the derivation of one intention from another, but 'I won't realize means M; so, I won't achieve end E' is not. In some cases, one may form an intention not to realize a necessary means (say because it is unacceptable), and so *decide* not to achieve an end—but the second abstention is a corollary of the first, not a way or means of achieving it. Given that A intends (2a), 'I can't achieve end E without realizing means M; so, I will realize means M' is intelligible as the derivation of an intention from a belief; but 'I won't realize means M; so, I *can* achieve end E without realizing

means *M*' is intelligible only as a piece of theoretical reasoning. So (III*a*) differs from (II*a*) in not being defensible by appeal to the relation of *modus ponens* to *modus tollens*.

What sense does (III*a*) make? In my view, so long as its 'ought' connotes having reason, it licenses *A* to infer, supposing that he doesn't intend (2*c*) but accepts (2*b*), that he has reason not to intend (2*a*). It licenses him to infer, supposing that he doesn't intend (2*c*) but intends (2*a*), that he has reason not to accept (2*b*). But the first is blatant bootstrapping. (What would entail that *A* has reason not to intend (2*a*) is rather that (2*b*) is true, and he has an undefeated reason not to intend (2*c*) which derives from the undesirability of realizing that means.) And the second inference is bootstrapping—and worse. Failing to intend some means to an intended end can indeed be accommodated psychologically, without discarding the end, by a doubt whether the means is necessary for the end; but there is no intelligible inference (except given an unusual setting) from declining to adopt a means to doubting its necessity.¹⁷ This is respected by (III), precisely because it doesn't admit anything like contraposition. It is neglected by (III*a*). This may distinguish various ways in which *A* is safely precluded from certain types of incoherence; but it is not satisfactory—if I am right about its implications—as the formulation of a rational requirement.

I have postponed what should be a fundamental question: how may requirements of rationality *explain* why we reason as we do? We do not, and could not, apply them as rules of inference in order, for example, to infer (I*c*) from (I*a*) and (I*b*).¹⁸ (II) and (II*a*) are only formulable by reference to (I*a*–*c*) for the reason that it is *given* that (I*a*) and (I*b*) together entail (I*c*). To derive (I*c*) from (I*a*) and (I*b*) *via* a mechanical application of (II*a*), as if one was manipulating uninterpreted formulae and not because one *saw* that (I*c*) follows validly, would not be to understand (II*a*) as a rule of *rationality*—nor to *infer* (I*c*) *logically*. So the claim would have to be that *A*'s willingness—if not his ability—to infer (I*c*) from (I*a*) and (I*b*) is somehow underpinned by his *implicit* acceptance of (II*a*), though he neither does, *nor could*, apply (II*a*) as he makes the inference. Again, it

¹⁷ Just conceivably, someone might find that, by some providential mechanism, he always fails to resolve upon any means that he is mistakenly inclined to accept as necessary to his ends.

¹⁸ If we did, we would need to detach or quasi-detach or the like (all of which Broome excludes).

isn't for me to make this intelligible.¹⁹

My conclusions so far are various. If we are to express what commits us to thinking rationally by postulating the existence of rational requirements, we have reason to prefer conditional 'ought's whose 'ought' is maximal, and has wider scope than 'if'. However, even here, where detachment is excluded, what I call quasi-detachment may generate ascriptions of reasons that are counter-intuitive. One may make use instead of wide-scope 'ought's that are disjunctive. But these too generate unwelcome reason-ascriptions; and they suit theoretical reasoning better than practical. Two fundamental questions I have simply put on the table. One was about the *authority* of rational requirements: why can they be infringed, but never flouted? The other was about their *role*: how do they explain successful reasoning? Answers are awaited.

V

Suppose that *A* instantiates one or both of the following schemata:

- (A) *Failure of Closure of Beliefs under Implication*: *A* believes (1a) and (1b), but doesn't believe (1c).
- (B) *Failure of Consequent Intentions*: *A* intends (2a) and believes (2b), but doesn't intend (2c).

Are these automatically failures of rationality? If so, do we need something like *norms* of rationality that proscribe such failures?

By '*A* believes (1a)' I mean that he believes something of the form 'If *p*, then *q*'; and so for the rest. I am permitting myself that inexact formulation for concision—but there is an explanation of inferential failure to which this must not blind us. *A* can't count as believing that *p* unless he has sufficient grasp of the proposition that *p*; but *sufficient* need not be *total*—or few if any of us could count as believing, for example, that anyone *knows* or *perceives* anything (given the uncertainty of how best to analyse those concepts). A subject who fails to make an inference from a proposition because of an imperfect grasp of its logical form is conceptually defective, but not *ipso facto* irrational. On certain analyses of 'know' and

¹⁹ See, slightly more fully, Price (2008, pp. 74–6).

'perceive', applications of them entail conditionals and counterfactuals. If someone is uncertain of the corresponding inferences, this need not show either that the analyses are incorrect, or that he lacks the beliefs he thinks he has, or that he is irrational.

Start with (A). I have just conceded that *A* may fail to infer something of the form of (1c) from premisses of the form of (1a) and (1b) because his grasp of the logical form of something within the premiss-pair is imperfect. Consider further Lewis Carroll's premiss-sets (1939, pp. 1119–40), of many of which the following is true: someone who sufficiently grasps the logical form of each of their members may still lack the logical intelligence to see what their conjunction entails. A failure by *A* either fully to grasp the logical form of the premisses, or to see what they entail, should not itself count as irrational. However, a general ability to infer by *modus ponens* is a condition of counting as holding beliefs with conditional contents. An omission to apply that ability in a particular case may involve a failure of rationality, though there can be other explanations, some of which may impute no rational failing of any kind. *A* may have no reason to bring (1a) and (1b) together in a way that would cause him to come to believe (1c); for example, he may have no interest in the truth of (1c).

Equally, (B) may involve no failure at all. It is characteristic of fully practical thinking starting from an end that it is *itself* conducted *for the sake of* that end. A successful piece of practical thinking is non-defective in content. (It doesn't, for example, identify a satisfactory means to an end accidentally through a false middle term.²⁰) It also serves the achieving of the end in context. (So it will fail if the thinker is forgetful of, or counter-suggestible to, his own deliberations.) Now, if *A* intends the truth of (2a), he can't be indifferent to the truth of (2c) so long as he believes (2b); so a lack of interest can't explain a failure to intend (2c). However, there remain several possibilities. Even if he will need to realize means *M* if he is going to achieve end *E*, he may not need to realize *M* *intentionally*; or, if he presently lacks an intention to realize *M*, he may have *no need yet* to form one (since he can do that later); or he may be *in no position yet* to do so (since he knows that, if he formed the intention now, he would forget it before the time came to act on it).

Any requirement of rationality, unlike a valid inference-schema, would have to be sensitive to such pragmatic considerations. So the

²⁰ See Aristotle, *Nicomachean Ethics*, VI 9, 1142b22–6.

schemata ($1a-c$) and ($2a-c$) are not mechanically translatable into requirements. However, intentions look less problematic than beliefs, just because they typically have a practical goal. So we can refine (B) as follows:

(B*) A intends ($2a$) and believes that he can't achieve E without both realizing M and now intending to realize M , but doesn't now intend ($2c$).

(A) is another matter. We might attempt the following:

(A*) A believes ($1a$) and ($1b$), but doesn't believe ($1c$), though the truth of ($1c$) is of interest to him.

Yet (A*) fails where (B*) succeeds: A can consciously hold back from forming an intention that ($2c$) just because he lacks the extended belief; yet he can't consciously hold back from inferring ($1c$) from ($1a$) and ($1b$) just because of a lack of interest. It marks a contrast between belief and intention that (A*) is an unhappy, while (B*) is a happy, mixture of logic and pragmatism. So we might retain (B*), but retreat from (A) to this:

(C) *Inconsistency in Belief*: A believes ($1a$) and ($1b$), but disbelieves ($1c$).

A requirement forbidding (C) is more plausible than one forbidding (A). To forbid (A), and (much more generally) any failure to think through the implications of one's beliefs, is to prescribe that they be closed under implication—which suits God better than man. To forbid (C) is to forbid a form of inconsistency. Yet, if (A) demands too much, (C) on its own demands too little: a subject who remains consciously agnostic about ($1c$), though he believes ($1a$) and ($1b$) and connects them, isn't in breach of (C).

I doubt whether there is any satisfactory formulation of a requirement governing logical inferences.²¹ What, anyway, would its function be? It can't reasonably tell us to make inferences we don't see to be valid; and to *see* an inference to be valid is to *make* it.²² If it told us always to look out for valid inferences, that would be an endless task. In the case of practical reasoning, we no more need rules to propel us from initial to derivative intentions than from intending to

²¹ See Price (2008, pp. 72–3).

²² See Anscombe (1995, pp. 3–4).

acting. It is a criterial of intending an end that one attempts any deliberation or action that one takes to be necessary for its fulfilment.²³

I shall end by discussing a different requirement, one that proscribes self-contradiction:

(IV) *A* ought not both to believe that *p* and to believe that $\neg p$.²⁴

It is possible for a subject to breach this requirement by keeping contradictory beliefs in different mental compartments. There are hypocrites whose hypocrisy isn't a matter of dissembling (saying that *p* when one consciously believes that $\neg p$), but of spontaneously and unconsciously shifting between different mental registers as the occasion invites. What isn't possible is consciously (and seriously) to hold contradictory beliefs.²⁵ To believe that $\neg p$ while consciously believing that *p* would be to believe that *p* and $\neg p$ —which denies what it asserts and asserts what it denies, and hence cancels its own content. (Elizabeth Anscombe (1959, p. 76) once likened this to taking a globe and painting its whole surface black.) Hence it is impossible to *be* in conscious breach of (IV), though one may be conscious of *having been* in breach of it.

What should we say of *A* if, unawares, he both believes that *p* and believes that $\neg p$? In what way or ways does he then count as irrational? It is obscure how it explains anything to suppose that he is in breach of a *requirement*. To take it that such a requirement would be a *rational* requirement is to appreciate already that this *is* a case of irrationality; identifying it as such doesn't wait upon the postulation of a requirement. It is true that a person may take more or less trouble to avoid unconscious self-contradiction; but it isn't clear that we make his goal more intelligible by adding that he wishes to avoid that out of respect for a requirement.

We have still to ask: what is bad about unconscious self-contradiction?

²³ Criteria, of course, are looser than necessary conditions. One can be neglectful of an end that one hasn't discarded. However, it is *failure* to try to think or act in accordance with an intention that calls for explanation; and blatant or brazen failure shows that one doesn't *really* intend the end.

²⁴ Simultaneously, of course; but I continue to omit time-references for simplicity. It is again for concision that I make use of (IV) rather than a more exact principle such as the following: 'A ought not to hold a pair of beliefs that relate to one another as *p* relates to $\neg p$.' As earlier, I set aside cases where *A* gets into trouble through an imperfect grasp of the logical form of his own beliefs.

²⁵ I add 'and seriously' parenthetically in order to set aside any conscious and eccentric self-contradictions inspired by philosophy.

diction? No doubt many things, some of which have nothing directly to do with rationality. It may, for instance, be an unstable state, and therefore a potentially unsettled and uncomfortable one. But what, precisely, is *irrational* about it? There seem, broadly, to be two alternative views. One finds it *in itself* irrational to hold one belief in one context, and its contradictory in another. The other holds instead that self-contradiction is the upshot of prior irrationality: if *A* has sufficient reason to believe that *p*, he can't also have sufficient reason to believe that $\neg p$. Hence, if, unawares, he believes that *p* (in one context) and believes that $\neg p$ (in another), either his belief that *p*, or his belief that $\neg p$, or both, must be contrary to reason.²⁶

These alternatives are interestingly different.²⁷ While the first finds irrationality within a certain pattern of beliefs, the second finds it in a failure to respond appropriately to reasons for belief.²⁸ Yet we can't say simply that the first views rationality as mind-centred, and the second as world-centred, for facts about states of mind can well be reasons (though not by way of bootstrapping). And nothing follows without argument about other values and dis-values. There might, for instance, be a vice of fickleness that was displayed even in being too quick to discard one ungrounded belief for another—though it might vary whether, in this context, constancy retained more than a vestigial value.

In my view, (IV) makes things worse. First, it entails that, if *A* does believe that *p*, however irrationally, then, relatively to that, he ought not to believe that $\neg p$. But this is only a plausible claim if the 'ought' is minimal. Secondly, if (IV) captures a peremptory reason, then *A* has, for any *p*, reason not to believe that *p* and that not $\neg p$, and so reason, equivalently, either not to believe that *p* or not to believe that $\neg p$.²⁹ It follows, so long as it is not already given how he will respect this reason if he does respect it, that he has a reason not to believe that *p*, and also a reason (the same one) not to believe that

²⁶ Cf. Peter Geach (1977, p. 166): 'A man who falls into inconsistency does not incur the further evil of a special sort of wrongness; it is only that logic suffices to show that somewhere or other (logic does not say where) he is wrong in a non-logical way.'

²⁷ There is much to be learnt here from Kolodny (2007).

²⁸ T. M. Scanlon (2007, p. 90) remarks, 'It does seem clearly irrational to have an attitude that one explicitly judges oneself to have conclusive reason not to have.' Certainly *something* is then irrational. Is it the attitude or the judgement, or their combination?

²⁹ To avoid this, we would need to rewrite (IV) as 'A ought not to believe that *p* if he believes that $\neg p$ '—understanding this in my way, not Broome's. This would then be subject to quasi-detachment.

$\neg p$. But does this really identify a reason against belief, or disbelief, that applies, though defeasibly, to *any* proposition?

I conclude that, *if* my logical intuitions are sound, those who appeal to requirements of rationality face a dilemma. If these constitute reasons for belief or intention, they generate implausible reason-ascriptions (often involving bootstrapping). If they don't, we may disregard them. Let me end with a flourish of liberation philosophy: rationality is a capacity; it is not directed by requirements.³⁰

Department of Philosophy
Birkbeck College London
Malet St
London WC1E 7HX
UK
a.price@bbk.ac.uk

REFERENCES

- Anscombe, G. E. M. 1959: *An Introduction to Wittgenstein's Tractatus*. London: Hutchinson.
- 1995: 'Practical Inference'. In Rosalind Hursthouse, Gavin Lawrence and Warren Quinn (eds.), *Virtues and Reasons: Philippa Foot and Moral Theory*, pp. 1–34. Oxford: Clarendon Press.
- Broome, John 1999: 'Normative Requirements'. *Ratio*, 12, pp. 398–419.
- 2001: 'Are Intentions Reasons? And How Should We Cope with Incommensurable Values?'. In Christopher W. Morris and Arthur Ripstein (eds.), *Practical Rationality and Preference: Essays for David Gauthier*, pp. 98–120. Cambridge: Cambridge University Press.
- 2005: 'Have We Reason to Do as Rationality Requires? A Comment on Raz'. *Journal of Ethics and Social Philosophy*, Symposium 1, pp. 1–8. Online at <www.jesp.org>.
- 2007: 'Wide or Narrow Scope?'. *Mind*, 116, pp. 359–70.
- forthcoming: 'Is Rationality Normative?'. In the conference volume of the European Congress of Philosophy, 2005.
- Carroll, Lewis 1939: *Complete Works*. London: Nonesuch Press.
- Clifford, W. K. 1877: 'The Ethics of Belief'. Available online at <www.infi-dels.org/library/historical/w_k_clifford/ethics_of_belief.html>.

³⁰ I am indebted to John Broome for critical comments, and grateful to the Aristotelian Society for a constructive discussion. Revisions after the event were prompted by Dorothy Edgington and David Wiggins.

- Geach, Peter 1977: *The Virtues*. Cambridge: Cambridge University Press.
- Kant, Immanuel 1997: *Groundwork of the Metaphysics of Morals*, ed. Mary Gregor, introd. Christine M. Korsgaard. Cambridge: Cambridge University Press.
- Kolodny, Niko 2005: 'Why be Rational?'. *Mind*, 114, pp. 509–63.
- 2007: 'How Does Coherence Matter?'. *Proceedings of the Aristotelian Society*, 107, pp. 229–63.
- Piller, Christian unpublished: 'The Normativity of Rationality'.
- Price, A. W. 2008: *Contextuality in Practical Reason*. Oxford: Clarendon Press.
- Raz, Joseph 2005a: 'The Myth of Instrumental Rationality'. *Journal of Ethics and Social Philosophy*, 1.1, pp. 1–28. Online at <www.jesp.org>.
- 2005b: 'Instrumental Rationality: A Reprise'. *Journal of Ethics and Social Philosophy*, Symposium 1, pp. 1–19. Online at <www.jesp.org>.
- Scanlon, T. M. 2007: 'Structural Irrationality'. In Geoffrey Brennan, Robert Goodin, Frank Jackson and Michael Smith (eds.), *Common Minds: Themes from the Philosophy of Philip Pettit*, pp. 84–103. Oxford: Clarendon Press.
- Strawson, P. F. 1997: "If" and "⊃". In his *Entity and Identity and Other Essays*, pp. 162–78. Oxford: Clarendon Press.