

Replies

John Broome

I am immensely grateful to all the authors of these comments about my book *Rationality Through Reasoning*. I always feel honoured when someone bothers to read what I have written, and it is much more of an honour when people take such trouble to write penetrating comments. Readers of my responses will notice that I tend to be obstinate in sticking to my views. This does not mean I do not recognize the value of the comments. Each one has given me food for thought, and each deserves a much longer response than I have given it. There are so many good points among all these comments taken together that I have simply selected a few to respond to.

I especially value and appreciate the extremely generous remarks several of the authors have made about my book. I also very much want to thank Antonio Gaitán for his work on editing this symposium.

Olav Gjelsvik

I am very glad that Olav Gjelsvik has presented an approach to practical reasoning that is so radically different from mine. Readers will value having the two approaches set out next to each other, to help them make an informed comparison. Unfortunately, Gjelsvik has had only a few pages for setting out his approach, whereas I have taken up a whole book with mine, so the comparison can hardly be balanced. Without a more developed account from Gjelsvik, I am not in a position to make a fair and properly worked-out response to him. All I can do is mention a few potential difficulties I see for his approach. I imagine that a developed theory might well be able to overcome them.

The first is that, so far as I can tell, Gjelsvik has not adopted the broad aim I set for my own book, to answer a much-debated question about motivation. Gjelsvik describes this aim of mine at the beginning

of his paper. It is to explain how we can sometimes bring ourselves to intend to do what we believe we ought to do. We are mostly disposed, when we believe we ought to do something, to intend to do it. Often this disposition works itself out automatically, through subpersonal processes. For example, when you come to believe that you ought to take a break from your work, you find yourself intending to do so. But sometimes the disposition works through a process of active reasoning, which is something you do. You may start from a belief that you ought to take a break, and bring yourself to intend to take a break. You motivate yourself through reasoning. Your belief is a premise-attitude in the reasoning, and your intention the conclusion-attitude. The theory of enkratic reasoning set out in my book explains how this sort of reasoning works.

On page 2 of the book, I gave the name ‘sheep’ to people who are strongly disposed to intend to do what they believe they ought to do, and ‘goats’ to people who are not. Enkratic reasoning is a way of making yourself more sheep-like.

Gjelsvik rejects my account of enkratic reasoning, and replaces it with a very different one. It starts from a ‘practical relation’, as he calls it, that the reasoner bears to a proposition. This proposition serves as a premise in the reasoning. I am not entirely clear just what a practical relation is, but I assume that standing in a practical relation to a proposition involves intending that proposition. In the taking-a-break example, the proposition in question is the conditional proposition ‘If I ought to take a break now then I shall take a break now’. So in Gjelsvik’s reasoning you start with an intention to do what you believe you ought to do. You are already motivated, that is to say; there is no question of motivating yourself. You start as a sheep. That is why I say Gjelsvik has not adopted my aim of explaining how you motivate yourself.

Gjelsvik rejects my account of enkratic reasoning because it involves an incorrect inference from ‘I ought to take a break’ to ‘I shall take a break’. By an ‘inference’ he means a connection between propositions. Gjelsvik and I both think of reasoning as a process that goes from some relations the reasoner has to propositions, which I shall call ‘premise-relations’, to a new relation she has to a proposition, which I call the ‘conclusion-relation’. I call the corresponding propositions the ‘contents’ of the relations. Gjelsvik assumes that, for reasoning to be correct, the proposition that is the content of the conclusion-relation

must be validly inferrable from the propositions that are the contents of the premise-relations.

This is a plausible condition of correctness for reasoning with beliefs. When you reason to a belief in a conclusion, the conclusion should be validly inferrable from the premises. Gjelsvik thinks the same condition applies to practical reasoning. He ‘sees the correctness of the practical inference itself as reflected in standard logic’. He thinks the correctness of a process of reasoning is determined by inferential connections among the contents of the premise-relations and the conclusion-relation, independently of whether these relations are beliefs or something else.

This cannot be right. The correctness of a piece of reasoning is not determined entirely by the inferential connections among its contents. For example, this is correct theoretical reasoning by *modus ponens*:

$$\begin{array}{l} \vdash_J \text{ (I shall visit Venice)} \\ \vdash_J \text{ (If I visit Venice, I shall be importuned by people selling} \\ \quad \text{handbags)} \\ \hline \vdash_J \text{ (I shall be importuned by people selling handbags)} \end{array}$$

(I use Gjelsvik’s notation to indicate the nature of the relations: ‘ \vdash_J ’ for a belief and ‘ \vdash_P ’ for a practical relation, which I assume to involve an intention.) But this is not correct reasoning:

$$\begin{array}{l} \vdash_P \text{ (I shall visit Venice)} \\ \vdash_J \text{ (If I visit Venice, I shall be importuned by people selling} \\ \quad \text{handbags)} \\ \hline \vdash_P \text{ (I shall be importuned by people selling handbags)} \end{array}$$

From a practical relation to visiting Venice, you plainly could not correctly derive a practical relation to being importuned by people selling handbags.

Whether or not reasoning is correct has to depend partly on the nature of the relations involved in it; it cannot depend only on a connection of valid inference among the contents. My own view, set out in my book, is that correctness depends on rational connections among the relations. Specifically it depends on ‘basing permissions of rationality’ as I call them, which hold among propositional attitudes.

I was myself deceived about this for some time, as the result of a confusion. Superficially, this can look like correct practical, instrumental reasoning:

$$\begin{array}{l} \vdash_P (\text{I shall visit Venice}) \\ \vdash_J (\text{If I visit Venice, I shall buy a ticket to Venice}) \\ \hline \vdash_P (\text{I shall buy a ticket to Venice}) \end{array}$$

It looks as if it brings you to have a practical attitude towards a means of visiting Venice – namely, buying a ticket to Venice – on the basis of a practical attitude towards visiting Venice.

The content of this putative piece of practical reasoning is the same as the content of this correct theoretical reasoning by modus ponens:

$$\begin{array}{l} \vdash_J (\text{I shall visit Venice}) \\ \vdash_J (\text{If I visit Venice, I shall buy a ticket to Venice}) \\ \hline \vdash_J (\text{I shall buy a ticket to Venice}) \end{array}$$

It is tempting to assume that both these pieces of reasoning are made correct by the same inferential connection among contents: modus ponens. That is what I once thought.¹

But it was a bad mistake. The putative practical reasoning is not correct at all, despite appearances. It has the same form as the obviously incorrect piece of reasoning I set out previously, which led incorrectly to a practical relation towards being importuned by people selling handbags. So this cannot be a form of correct practical reasoning. Correct instrumental reasoning takes a different form, which is described in my chapter 14.

Gjelsvik uses a different example, but I think he makes the same mistake. He is too impressed by a superficial similarity between modus ponens reasoning and instrumental reasoning.

The upshot is that correctness for practical reasoning is not determined by inferential connections among contents. So I think Gjelsvik's objection to my version of enkratic reasoning is unsuccessful.

María José Frápolli and Neftalí Villanueva

M. J. Frápolli and N. Villanueva address their comment to a core topic of my book, the requirement of rationality I call 'Enkrasia'. I re-

produce it here from page 170:

Enkrasia. Rationality requires of N that, if

- (i) N believes at t that she herself ought that p , and if
- (ii) N believes at t that, if she herself were then to intend that p , because of that, p would be so, and if
- (iii) N believes at t that, if she herself were not then to intend that p because of that, p would not be so, then
- (iv) N intends at t that p .

Condition (i) uses a deviant version of English grammar that I introduced on page 14. ‘Ought’ is treated as a lexical verb rather than auxiliary verb. I adopt this grammar because it makes it possible to say explicitly who owns an ought, whereas common English does not. The English sentence ‘Jane ought to have a better job’ does not tell us whether or not the ought is owned Jane. Is it up to Jane to find a better job, or is it perhaps up to her boss to give her one? In my deviant grammar, the subject of ‘ought’ is the owner of the ought. ‘Jane ought that Jane has a better job’ ascribes ownership to Jane. If we do not mean to ascribe ownership, we may say ‘It ought to be the case that Jane has a better job’.

In condition (i), the schematic letters ‘ N ’ and ‘ t ’ stand for a person and a time respectively, and ‘ p ’ stands for a proposition.

The pronoun ‘she herself’ in condition (i) corresponds in indirect speech to the pronoun ‘I’ in direct speech. If N were to express the belief described in (i) using my deviant grammar, she would say ‘I ought that p ’. If condition (i) did not include this reflexive pronoun, *Enkrasia* would not be true. Suppose, say, that Letizia believes that the owner of the dead donkey ought to bury it. Suppose Letizia is the owner of the dead donkey, but suppose she does not intend to bury it. Is she necessarily irrational? She is not. She might not realize she is the owner of the dead donkey. In that case, she may be rational even though she does not intend to bury it. She is irrational in not intending to bury the dead donkey only if she believes that she herself ought to bury it – that is to say, only if she has a belief that she could express by saying ‘I ought to bury the dead donkey’. *Enkrasia* requires ownership of the ought to be self-ascribed.

Frápolti and Villanueva argue that my use of a reflexive pronoun in (i), corresponding to ‘I’ in direct speech, is not enough to establish the degree of self-ascription that Enkrasia requires. They think it does not rule out a particular sort of error of misidentification. I think their point is that you might have a belief that you are in a position to express by saying ‘I ought that p ’ but nevertheless not self-ascribe the ownership of the ought. If they are right, I need to change my formulation.

But I think they are not right. They explain their point using an example from François Recanati.² They report Recanati as describing a schizophrenic patient who believes he is conscious of some mental states that are not his but belong to ‘the Other’. For instance, he is conscious of a belief that could be expressed by ‘I am good and omnipotent’, but does not believe this is a belief of his. This is a common situation in one respect. You might be conscious of someone else’s belief that she herself is good and omnipotent without believing this is a belief of yours. For example, she might have told you about it. Recanati’s particular case is not common because the schizophrenic’s consciousness of the Other’s beliefs is like his consciousness of his own beliefs. But in any case, the schizophrenic does not believe that he himself is good and omnipotent. Nor is he in a position to say ‘I am good and omnipotent’. So Recanati’s case does not support Frápolti and Villanueva’s view that you might have a belief that you are in a position to express by saying ‘I ought that p ’ but nevertheless not self-ascribe the ownership of the ought.

That is one argument of Frápolti and Villanueva’s. Here is another. Condition (i) refers to self-ascription at one point, using a reflexive pronoun. Frápolti and Villanueva argue that it needs to contain self-ascription at a second place too, within the scope of ‘ought’. They think the condition should be formulated

N believes at t that she herself ought that she herself Fs .

They replace my propositional letter ‘ p ’ with ‘she herself Fs ’. Evidently, Frápolti and Villanueva assume that in a correct statement of Enkrasia it must be possible for N to express the proposition p by saying ‘ IF ’. N must self-ascribe the property of $Fing$.

I think this assumption is doubly incorrect. First, there may be genuine instances of Enkrasia in which the proposition p cannot be expressed using a subject-predicate sentence in which the subject de-

notes *N*. Second, even in cases where the proposition can be expressed this way, *N* need not self-ascribe the predicate, so the second pronoun need not be reflexive.

Suppose Letizia believes that she herself ought that the owner of the dead donkey buries it. Suppose conditions (ii) and (iii) are satisfied: roughly, this means that Letizia believes it is up to her whether or not the owner of the dead donkey buries it. Perhaps she believes she has perfect command over all donkey owners. If she does not intend that the owner of the dead donkey buries it, she is irrational.

In this application of *enkrasia*, the proposition *p* is the proposition that the owner of the dead donkey buries it. If we suppose for the moment that Letizia does not own the dead donkey, this proposition cannot be expressed by a sentence whose subject denotes Letizia. This illustrates the first way in which I think Frápolli and Villanueva go wrong.

Many philosophers think it is actually I who have gone wrong, because my example is impossible. If I understand them right, Frápolli and Villanueva are among those philosophers. They think it could not possibly be the case that Letizia ought that the owner of the dead donkey buries it, unless Letizia owns the dead donkey. I shall return to this view soon.

Before that, suppose now for a moment that Letizia owns the dead donkey. The proposition *p* in this case is that Letizia buries the dead donkey. So this is a case in which the proposition *p* can be expressed by a subject-predicate sentence in which the subject denotes *N* (Letizia). But suppose Letizia does not realize she owns the dead donkey. If she does not intend that the owner of the dead donkey buries it, is she necessarily irrational? She is; I have already said so. She is irrational because she self-ascribes the ought that the owner of the dead donkey buries it, and because she believes it is up to her whether or not the owner of the dead donkey buries it. It is irrelevant whether or not she self-ascribes burying the donkey. This illustrates the second way in which I think Frápolli and Villanueva go wrong.

But is my example impossible after all? Now I return to this question. It is true that it cannot be described using grammatical English, in which ‘ought’ is an auxiliary verb. The nearest we could get in grammatical English would be ‘Letizia believes that she herself ought to bury the dead donkey’. This sentence contains within the scope of ‘believes’ just the one compound verb ‘ought to bury’, which can have only one subject. The sentence does not distinguish the subject of

‘ought’ from the subject of ‘bury’. The subject of both is the pronoun ‘she herself’, which denotes Letizia, and represents self-ascription on Letizia’s part. My example is very different from this.

Frápolti and Villanueva’s view is that the metaphysics of ought follows the English grammar of ‘ought’ in this respect. They think ought is a relation between a person and a first-level predicable, rather than a relation between a person and a proposition as I assume. An infinitival construction such as ‘to bury the dead donkey’ expresses a first-level predicable. It does not allow the subject of ‘ought’ to be distinguished from the subject of the infinitive.

What is the evidence for this view? The grammar itself is only weak evidence. The grammar of ‘ought’ is plainly inadequate for metaphysical purposes. ‘Ought’ does not even have the grammatical inflexions we need. It has no tenses, but we would not deduce that oughts cannot exist in the past or the future. It has no nominalization such as a gerund, and for that reason philosophers often find they have to use ‘ought’ ungrammatically as a noun. I attach little metaphysical weight to the fact that it is grammatically an auxiliary verb.

Better metaphysical evidence for Frápolti and Villanueva’s view comes from the idea of ought implies can. It can be the case that you ought something only when you have some control over that something. But this metaphysical point gives little support to the idea that ought is a relation between a person and a first-level predicable, since first-level predicables do not only represent things you have control over. For example, suppose Florence ought to get a medal for her heroic acts. She may have no control over whether or not she gets a medal for her heroic acts, even though there is a first-level predicable denoted by ‘to get a medal for her heroic acts’. Conversely, if you control the donkey owners, you may control whether or not the owner of the dead donkey buries it, even though there is no corresponding predicable.

So I think the evidence for Frápolti and Villanueva’s view is weak. Nevertheless, I recognize that the metaphysical argument can continue.³ In my book I therefore tried to circumvent it by pointing out the greater generality of the propositional structure compared with the first-level predicable structure. I said that anything we can say using English grammar can be said using my grammar, whereas the converse is not so.

Frápolti and Villanueva deny this. They think that the degree of

self-ascription required to make Enkrasia correct cannot be properly expressed using the propositional structure. I think they are wrong. For condition (i) they themselves provide a formulation using the propositional structure but nevertheless referring to self-ascription in the two places where they think it is needed:

N believes at *t* that she herself ought that she herself *F*s.

True, they do not think the degree of self-ascription expressed by the reflexive pronoun ‘she herself’ is sufficient. But I have explained that I am not convinced by that either.

Moreover, an important point can be made using the propositional structure that cannot be made without it. It is the point I have been stressing: the truth of Enkrasia does not require self-ascription within the scope of ‘ought’. Ownership of the ought must be self-ascribed, but the property of *F*ing need not be self-ascribed. The second reflexive pronoun in Frápolli and Villanueva’s formulation is not needed. If we use the propositional letter ‘*p*’, there need be no self-ascription within *p*. This is a truth that is obscured by the formulation with a first-level predicable:

N believes at *t* that she herself ought to *F*.

Here we cannot separate the subject of ‘ought’ from the subject of ‘to *F*’, so we cannot say that one needs to be a reflexive pronoun and the other does not.

Conor McHugh and Jonathan Way

I shall start my response to Connor McHugh and Jonathan Way with the second part of their paper. There they develop an account of correctness for reasoning, which is based on the idea that reasoning has a point. They take the point of reasoning to be to preserve correctness: if you start with correct premise-attitudes, reasoning takes you to a correct conclusion-attitude.

They eventually formulate their account of correctness as:

It is correct for *S* to reason from attitudes $P_1 \dots P_n$ to attitude *C* if and only if (i) other things equal, if $P_1 \dots P_n$ are all correct, *C* is correct and (ii) *S* is in a position to be sensitive to (i).

Clause (ii) is included to deal with a particular problem, and so is the 'other things equal' clause. In this response, I shall leave aside those particular problems, and concentrate on the rest of clause (i). This first needs some explaining and refining.

What does it mean for an attitude to be 'correct'? McHugh and Way evidently mean that it is permissible rather than that it is required. For example, they offer this as a possible description of correctness for intentions:

We might think, for instance, that it is correct to intend to *F* just when you are permitted to *F* and it is up to you... whether to *F*.

This is a weak sort of correctness. It might be correct to intend to get lunch at Mario's and also correct to intend to get lunch at Carla's, even though you could not do both. 'Permissible' could be substituted for 'correct' in what follows.

Even together with clause (ii), clause (i) is not a sufficient condition for a piece of reasoning to be correct. Correctness must also depend on there being some appropriate connection between the premise-attitudes and the conclusion-attitude. For example, suppose it is correct for you to believe *p*, to believe *q*, and to believe that if *p* then *q* – suppose each of these propositions is true, say. Suppose you do believe *q* and also believe that if *p* then *q*. Suppose you reason from these premise beliefs to a conclusion belief that *p*. Since your premise-attitudes and your conclusion-attitude are all correct, your reasoning satisfies McHugh and Way's condition. Yet it is not correct reasoning; it is fallacious.

When you reason, your conclusion-attitude is based on your premise-attitudes. Your reasoning can be correct only if it is correct for it to have this basis. I think any account of the correctness of reasoning will have to refer to basing of some sort. So McHugh and Way's condition needs to be strengthened to something like: if your premise-attitudes are correct, it is correct for you to have the conclusion-attitude on the basis of your premise attitudes.

How should we interpret this condition when your premise-attitudes are not correct? It is possible to reason either correctly or incorrectly on the basis of incorrect premises. We must interpret the condition in any way that recognizes both these possibilities. I think the best interpretation is to read it as a subjunctive conditional: if your

premise-attitudes were correct, it would be correct for you have the conclusion-attitude on the basis of your premise-attitudes. I shall call this 'condition (i')'; it is my refinement of condition (i).

It hits a problem if there are attitudes that are necessarily incorrect. It seems plausible there are, though this depends on what makes an attitude incorrect. McHugh and Way apparently provide an example, since they apparently think that a false belief is incorrect. If that is so, a necessarily false belief is necessarily incorrect. If one of your premise attitudes is necessarily incorrect, the antecedent of condition (i') is necessarily false. Yet on the face of it, it is possible to reason either correctly or incorrectly on the basis of a necessarily incorrect attitude. Suppose, for example, that the Goldbach conjecture is true. Since it is an arithmetic formula, this means it is necessarily true. If you believe it is false, your belief is necessarily incorrect. Nevertheless, on the basis of this necessarily incorrect belief, you might correctly reason to a conclusion belief-that the Goldbach conjecture is not true, or you might incorrectly reason to a conclusion-belief that it is true. Or so it seems at least. This cannot be explained by condition (i').

My own condition for reasoning to be correct is not much different from (i'). The main difference is that I give the conditional a wide scope. I think that, when you reason from some premise-attitudes to a conclusion-attitude, your reasoning is correct if and only if it is correct for you, if you have the premise-attitudes, to have the conclusion-attitude on the basis of the premise-attitudes. This is what I called a 'basing permission'. One advantage of my formula is that it avoids the problem condition (i') encounters over attitudes that are necessarily incorrect.

McHugh and Way might complain that my formula misses the point of reasoning. The point is to preserve correctness: to take you from correct attitudes to another correct attitude. But if this is indeed the point of reasoning, my formula does not necessarily miss it. One could argue that my formula actually entails condition (i'). This would not be a straightforward argument to make, and I shall not attempt it here. If it goes through, it means that my formula embodies what McHugh and Way see as the point of reasoning.

But actually, preserving correctness cannot be the point of reasoning. At least, not if the property of correctness is as weak as McHugh and Way suggest it is in the example of theirs that I quoted above. It might be correct for you to intend to get lunch at Mario's and

also correct for you to intend to get lunch at Carla's. It might also be correct for you to intend to turn left, and also correct for you to intend to turn right. Suppose left is the way to Carla's and right is the way to Mario's. Suppose you intend to get lunch at Carla's. On the basis of this correct intention, reasoning should be able to get you to intend to turn left. But the aim of preserving correctness would equally well be served by intending to turn right, since intending to turn left and intending to turn right are both correct. This example shows that the point of reasoning is to acquire attitudes that stand in an appropriate relationship to each other, not merely attitudes that are correct.

You might think that intending to turn right is not correct once you intend to get lunch at Carla's. This is to adopt an idea of correctness for attitudes that is relative to other attitudes. I agree the point of reasoning is to preserve this sort of relative correctness. Preserving relative correctness is just making sure your attitudes stand in an appropriate relationship to each other. That is exactly what I think correctness is.

The upshot is that, once McHugh and Way's account of correct reasoning is refined as it has to be, I think it is not far different from my own account. Where it differs, I think mine is preferable.

My account left what I called on page 258 'a serious gap in this book's argument'. In the book I was unable to say exactly what the relationship is between basing permissions of rationality and synchronic requirements of rationality, but I did insist there has to be one. The first part of McHugh and Way's paper emphasizes that this indeed a serious gap. Some of the things I said were evidently misleading. For one, I talked of deriving permissions from synchronic requirements, but I did not mean to suggest permissions are derived from synchronic requirements only. As McHugh and Way point out, they cannot be, because synchronic requirements have a particular sort of symmetry that permissions do not have. Something else has to be included, which determines a direction for permissible basings. I take this something else to be a rational prohibition, and as McHugh and Way point out, it cannot be derived from a synchronic requirement.

McHugh and Way's approach was intended to help close this serious gap in the argument. I am grateful to them for the attempt, but I do not think their approach is sufficiently different from mine to succeed in this aim.

Miranda del Corral

Miranda del Corral criticizes the requirement of rationality I call 'Persistence of Intention'. This is a good target to choose. I do not feel confident in this requirement, and I would not be surprised to find that some of the details of its formulation are mistaken. It tries to tread a delicate path between being too strong and being too weak. On the one hand, when you intend to do something, clearly you may at least sometimes change your mind and drop your intention, without any irrationality. On the other hand, if you change your mind too often or too readily, that is irrational. There has to be a persistence requirement of rationality that fits somewhere between these extremes. To speak roughly, the requirement I propose is that, when you have an intention, you should not simply drop it without reconsidering it, unless some 'cancelling event', such as discovering new information, occurs in the meantime.

I recognize that my formulation may be wrong, but I do not think del Corral has correctly identified faults with it. So far as I can tell, her arguments depend on some misunderstandings. For example, she takes me to deny that amnesia caused by a blow to the head is a cancelling event. But I do not deny that. On pages 180-1 I described three types of cancelling events, and said there are other cancelling events that do not fall within these types. I mentioned death as an example. Amnesia caused by a blow is another. If you lose an intention that way, I think it is not a failure of your rationality.

For another example, del Corral takes a time t_2 to lie between another time t_1 and itself t_2 , and assumes that I use 'between' this way. But I do not. As I understand 'between', t_2 does not lie between t_1 and t_2 .

I tried to formulate Persistence of Intention as precisely as I could, navigating the delicate path between the extremes I mentioned above. Details of the formulation make a big difference. A criticism needs to take them into account accurately.

Del Corral is motivated by very natural thought that rationality requires us not to be weak-willed. I agree, and Persistence of Intention is my attempt to formulate this requirement. Del Corral thinks that sometimes the requirement not to be weak-willed requires us actually to do what we intend. I explained in an earlier chapter [on pages 151-2] that, if you intend to do something, normally you do it. This is because, if you do not do it and nothing prevents you, that implies you do not intend to do it. So Persistence of Intention normally does the job del

Corral wants done. But I explained in the same place that rationality cannot require you actually to act as you intend, or even to try to act as you intend.

To explain why not, I gave the example of a slow watch. Suppose you intend to get on the bus by 10.00, but your watch is slow and 10.00 goes by without your knowing it. Then it may happen that you do not get on the bus before 10.00 or even try to get on it before 10.00. Yet you may be rational; the fault lies with your watch.

Fernando Broncano and Jesús Vega

Fernando Broncano and Jesús Vega offer an account of how rationality can be non-derivatively normative. Rationality has its own characteristic achievement, which is appropriately adjusting our attitudes in a way that is not just lucky. Adjusting attitudes in this way has value, and is meritorious. They make an illuminating comparison with the value of knowledge, which goes beyond the value of merely true belief.

I like their paper, and I think it may well be right, but at present I find myself with little to say about it. An adequate response would demand more thought than I have been able to give it. It is easy to be sceptical about sources of normativity. In this case, I might ask for an explanation of why the achievement of rationality or of knowledge is more meritorious than other achievements such as accurately counting the blades of grass in a lawn. Broncano and Vega say ‘a rational subject sees her rational activity as something valuable and as something she has reason to promote’. True, but a grass-counter may see her activity in the same light, and she would be wrong. We need to know whether rational activity is truly valuable and truly something a person has reason to promote. How people see it is not enough.

There really is normativity in the world, which means that scepticism must sometimes be answerable. There may be a good answer on behalf of rationality to the scepticism I have described, and I definitely do not want to press a sceptical response to Broncano and Vega.

Nicholas Shackel

Nicholas Shackel examines the arguments in my chapter on the normativity of rationality. He gives them useful detailed scrutiny. Not surprisingly, he is particularly interested in whether they might cause any difficulties to his own favoured theory of rationality. His theory is

very different from mine, but some of my formal arguments might carry over to rationality as he sees it, and that is why he investigates them.

If I understand him correctly, Shackle believes that, necessarily, if you are locally as you ought to be, then rationality requires you to F if and only if you ought to F . ‘You are locally as you ought to be’ means that your relevant attitudes are ones you ought to have. For example, your relevant beliefs are true and the relevant things you care about are ones you ought to care about. You ‘care rightly’ as he puts it.

Shackle’s explanation of why rationality requires you to F if and only if you ought to F is that ‘the internally available correctness normativity of rationality will line up with and thereby represent the external directive normativity in those localities where we care rightly’. I think this is a mistaken theory of rationality, and since theories like this are common,⁴ it is perhaps worth my explaining my objections to them.

One is that often you ought to F when F ing is a physical act. It might be an act of kindness, say. But rationality cannot require you to do a physical act, because rationality supervenes on the mind. If the minds of two people were in the same state, then they would be equally rational. Take two people whose minds are in the same state; in particular they both intend to do a particular act of kindness. They are equally rational. Yet one might do the act and the other not, perhaps because she is unexpectedly prevented from doing it by an external intervention. So failure to do the act cannot be a failure of rationality.

I have another objection. Suppose you ought to do some act F of kindness, and this is on grounds of morality. Suppose among your relevant attitudes is an intention of doing acts of kindness. Suppose you ought to have this attitude, also on grounds of morality. Suppose your other relevant attitudes are as they ought to be. Then Shackle would conclude that rationality requires you to F . But this cannot be right. We may suppose it is necessarily true that you will F if you are rational *in addition to* having the attitudes you ought to have. But it is *morality* that makes it the case that you ought to have the intention of doing acts of kindness. So it cannot be rationality alone that requires you to F . At best it is rationality and morality together.

To generalize the point, suppose you ought to F and your relevant attitudes are A_1, A_2 and so on. Suppose you ought to have attitudes A_1, A_2 and so on, and actually you do have them. It might be that rationality requires of you that, if you have A_1, A_2 and so on, you F . Does it follow that rationality requires of you that you F ? It does not.

To reach that conclusion we would have to detach a narrow-scope rational requirement from a broad-scope one, and our detachment rule would have to mix the operators ‘ought’ and ‘rationality requires’. I do not think there is a valid rule of that sort. If anything requires you to *F*, it would have to include whatever makes it the case that you ought to have attitudes A_1, A_2 and so on. It could not just be rationality.

Nevertheless, for the sake of pursuing the argument, I am happy to adopt Shackle’s theory of rationality temporarily. Let us ask whether either of my formulae

Normativity of Rationality. Necessarily, if rationality requires you to *F*, that fact is a reason for you to *F*

or

Weak Normativity. Necessarily, if rationality requires you to *F*, there is a reason for you to *F*,

is true within Shackle’s theory. Shackle would like at least Weak Normativity to be true.

Normativity of Rationality is false within his theory. According to this theory of rationality, in cases where you care rightly, what is rationally required of you lines up with and represents what you ought to do. The direction of explanation goes from what you ought to do to what rationality requires you to do. What you ought to do is determined by ‘external directive normativity’. The fact that rationally requires you to *F* does not in any way contribute to explaining why you ought to *F*. As I explain in chapter 4, a reason is something that explains or contributes to explaining an ought fact. So the fact that rationality requires you to *F* is not a reason to *F* in cases when you ought to *F*. It therefore cannot be a reason to *F* in cases where it is not the case that you ought to *F*.

Shackle correctly criticizes a parallel argument against Normativity of Rationality that appears in chapter 11 of my book. I would like to take this chance to correct the mistake in my argument there.

I used the example of a Bayesian requirement of rationality: specifically, the requirement to have degrees of belief that add up to one. I assumed that an argument based on a Dutch Book shows you have a prudential reason to have degrees of belief that add up to one – the de-

tails of the argument do not matter. I then argued that the fact that rationality requires you to have degrees of belief that add to one is not itself a reason to have degrees of belief that add up to one. I said 'Even if rationality did not require you to have degrees of belief that add up to one, you would still have just the same prudential reason to do so'.

That was my mistake. As Shackel points out, my counterfactual claim is dodgy, because the antecedent may be necessarily false. I should have referred to an explanatory connection rather than a counterfactual one. 'A reason' is an explanatory notion, and it is the lack of an explanatory connection that leads correctly to my conclusion. I should have said 'The fact that rationality requires you to have degrees of belief that add up to one does not contribute to explaining why you ought to have degrees of belief that add up to one, when you ought'. From this it follows that the fact that rationality requires you to have degrees of belief that add up to one is not a reason for having degrees of belief that add up to one, since a reason must either explain or contribute to explaining an ought fact.

So Normativity of Rationality is false given Shackel's theory of rationality. What about Weak Normativity? Is it true that, necessarily, if rationality requires you to *F*, you have a reason to *F*? It is not. First, I have just argued that the fact that rationality requires you to *F* is not a reason to *F*. Second, there may be no other reason. According to Shackel's theory, any other reason would have to come from 'external directive normativity'. According to the theory, if you are locally as you ought to be, external directive normativity makes it the case that you ought to *F* on occasions when rationality requires you to *F*. But if you are not locally as you ought to be, external directive normativity might give you no reason to *F* even on occasions when rationality requires you to *F*. So even when rationality requires you to *F* you might have no reason to *F*.

It is easy to make a mistake about this. It is no doubt likely that you are locally as you ought to be. Consequently, the fact that rationality requires you to *F* may be evidence that you ought to *F*. In that case, the fact that rationality requires you to *F* is a reason to believe you ought to *F*. This can give some attraction to the idea that this fact is a reason to *F*.⁵ But that idea is mistaken.

At least two mistaken lines of thought could bring you to this conclusion. One stems from the assumption that rationality requires of you that, if you believe you ought to *F*, you *F*. This is a bastard ver-

sion of the requirement I call 'Enkrasia'; let us temporarily accept it for the sake of argument. It can be mistaken for the view that believing you ought to *F* is a reason for you to *F*, which in turn might lead you to think that a reason for believing you ought to *F* is a reason to *F*. But this is a mistake.⁶

The second line of thought is that, if you *F*, the fact that rationality requires you to *F* might explain your *F*ing in the following way. Because this fact is a reason for you to believe you ought to *F*, it might through some epistemic process lead you to believe you ought to *F*, which in turn might lead you to *F*. This explanation of why you *F* involves your rationality in a way that may make it appropriate to say that you *F* for the reason that rationality requires you to *F*, and not just *because* rationality requires you to *F*. We might naturally draw the conclusion that the fact that rationality requires you to *F* is a reason for you to *F*. But that would be a mistake. To say that you *F* for the reason that rationality requires you to *F* is to say that the fact that rationality requires to *F* is a 'motivating' reason to *F*. It is not to say it is a normative reason to *F*, and we are concerned only with normative reasons.

I have argued that, if we accept Shackel's theory of rationality, Weak Normativity of Rationality is false. Shackel criticizes some of the more complex arguments that appear in my chapter on the normativity of rationality, but my argument above is different and simpler. It applies to Shackel's theory of rationality specifically.

*Faculty of Philosophy
Oxford University
Corpus Christi College
Oxford OX1 4JF, UK
and
Australian National University
Camberra, Australia
E-mail: john.broome@philosophy.ox.ac.uk*

ACKNOWLEDGMENTS

Research for this paper was supported by ARC Discovery Grant DP 140102468.

NOTES

¹ See my 'Normative Practical Reasoning', *Proceedings of the Aristotelian Society*, Supplementary Volume 75 (2001), pp. 175-93.

² François Recanati, *Perspectival Thought*, Oxford, Oxford University Press, 2007.

³ It is pursued further in my paper 'Williams on *Ought*', in *Luck, Value and Commitment: Themes from the Ethics of Bernard Williams*, edited by Ulrike Heuer and Gerald Lang, Oxford, Oxford University Press, 2012, pp. 247-65.

⁴ For example, it resembles Derek Parfit's theory in *On What Matters*, Oxford, Oxford University Press, 2011, p. 111.

⁵ See Stephen Kearns and Daniel Star, 'Reasons: Explanation or Evidence', *Ethics*, 119 (2008).

⁶ See my 'Normative Requirements', *Ratio*, 12 (1999), pp. 398-419, reprinted in *Normativity*, edited by Jonathan Dancy, Blackwell, 2000, pp. 78-99.