

Responses*

Répliques

John BROOME

SUMMARY: I. *Acknowledgments*. II. *Fernando Rudy*. III. *Daniel Fogal*.
IV. *Alex Worsnip*. V. *Carlos Nuñez*. VI. *References*.

I. ACKNOWLEDGMENTS

I am honoured that my four commentators have taken a great deal of trouble to understand my book *Rationality Through Reasoning*, and find their way through its complexities. I am immensely grateful to them for the care they have taken, and I very much appreciate their unreasonably generous comments. Naturally, in replying, I have to pick out points where I disagree with them, and this is likely to conceal how far I have been convinced by their arguments. So let me say now that I have learnt a great deal from them in studying their papers.

I also found our discussions at UNAM very helpful. I especially want to thank Juan Vega for so generously inviting us all to the meeting, and arranging it so well. Very sadly, it was suddenly brought to an end by the tragic earthquake of 19 September 2017. Two of the authors, who had come so far to Mexico City, were not able to present their papers. Nevertheless, we did manage to have an excellent philosophical conversation.

* Artículo recibido el 17 de agosto de 2017 y aprobado para su publicación el 30 de noviembre de 2017.

JOHN BROOME

II. FERNANDO RUDY

The term 'reason' is widely recognized to be ambiguous. For instance, we use it for things that explain an action done by a person, and explain it in a particular way that involves the person's rational faculty. So we often add a qualifier to identify the sense of 'reason' we are using: 'explanatory reason', 'motivating reason', 'normative reason' and so on. In my discussion of Fernando Rudy's paper I am concerned with normative reasons only. When I use 'reason' on its own, it refers to a normative reason.

I define a normative reason as follows:

A normative reason is either a pro toto reason or a subsidiary reason.

A pro toto reason for N to F is something that explains why N ought to F.

A subsidiary reason for N to F is something that plays a particular role in a particular form of explanation of why N ought to F or of why it is not the case that N ought to F.

This definition of a subsidiary reason leaves places open for specifying just what the particular form of explanation is, and what is the particular role played within it by a subsidiary reason. Different forms of explanation and different roles may give us different sorts of subsidiary reason. The most familiar sort is the *pro tanto* reason.

A pro tanto reason for N to F is something that plays a particular role in a weighing explanation of why N ought to F or of why it is not the case that N ought to F.

Here, I do not need to go into the details of a weighing explanation or of the role of a *pro tanto* reason within it (See Broome, 2013, section 4.3.).

These are definitions of properties: the property of being a *pro toto* reason, the property of being a subsidiary reason and the prop-

RESPONSES

erty of being a *pro tanto* reason. For example, the definition of a *pro toto* reason more fully spelt out is:

The property of being a *pro toto reason* for *N* to *F* is the property of being something that explains why *N* ought to *F*.

‘Normative’ is a technical term that has different meanings in different disciplines. In moral philosophy, it means ‘involving ought’. My definitions specify what is the precise involvement of ought in the property of being a normative reason. With this definition in place, I cannot make sense of a further question ‘Are reasons defined this way truly normative?’.

But this is the question that drives Rudy’s paper. Rudy takes his lead from Jonathan Dancy. Dancy doubts that a *pro toto* reason as I defined it can be normative. Rudy quotes him:

I do not see how he [Broome] can both say that the notion of an explanation why you ought to *F* is not normative, that the notion of the (*pro toto*) reason for you to *F* is normative, and that these are the same notion. (Dancy, 2015, 178)

But I never said that the notion of an explanation of why you ought to *F* is not normative. The notion of *an explanation* is not normative, because it does not involve ought. But the notion of *an explanation of why you ought to F* is normative, because it does involve ought. Similarly, the notion of change is not meteorological, but the notion of climate change is meteorological.

Dancy defines the notion of normativity differently from me, as having to do with *favouring*. I am willing to go along with his definition, because I can define favouring in a way that makes ‘favouring *N*’s *F*ing’ equivalent to ‘being a normative reason for *N* to *F*’. My definition is:

To favour *N*’s *F*ing is either to explain why *N* ought to *F* or to play a particular role in a particular form of explanation of why *N* ought to *F* or of why it is not the case that *N* ought to *F*.

JOHN BROOME

So I take Dancy's notion of normativity to be equivalent to mine. But Dancy denies this equivalence because he rejects my definition of favouring.

The first clause of my definition, put briefly, says that favouring is ought-making - in other words right-making. But Dancy denies that favouring can be right-making. Rudy presents Dancy's argument by comparing the two sentences:

Your having promised to take your friend to the airport *favours* taking your friend to the airport.

Your having promised to take your friend to the airport *makes it the case* that taking your friend to the airport is right.

Rudy points out that the two relations denoted by the italicized words in these sentences are not the same relation. Indeed they are not. But this does not show that the favouring relation is not the right-making relation. The first is the favouring relation, but the second is not the right-making relation. It is the making-it-the-case relation. Rudy's example shows only that the favouring relation is not the making-it-the-case relation.

The right-making relation is denoted by the italicized words in:

Your having promised to take your friend to the airport *makes it right* to take your friend to the airport.

Nothing stops this from being the favouring relation in the example.

I am not sure Rudy presents Dancy's argument accurately, because I am not sure what Dancy's argument is. He says:

The right-hand side of the right-making relation is the rightness so made. The relation involved is the general making-relation, the right hand side of which is in this case the rightness of the action (Dancy, 2015, 178).

In the first of these two sentences, Dancy insists that the reason (in the example, your having promised) stands in the right-making relation to the rightness of the act (the rightness of taking your friend to the airport). But this is not so; it stands in the right-making relation

RESPONSES

to the act itself (taking your friend to the airport). It stands in the general making relation to the rightness of the act, as Dancy says in the second sentence. I am not sure why he also says it stands in the right-making relation to the rightness of the act. When you make your friend happy, you stand in the happy-making relation to your friend and in the making relation to your friend's happiness. You do not stand in the happy-making relation to your friend's happiness.

Rudy ascribes a parallel mistake to me. I said at one point:

I agree that a reason is a consideration that counts in favour of something. But my definition goes further and specifies what is the relevant sort of counting in favour (Broome, 2013, p. 54).

Rudy adds by way of elucidation: "The relevant sort of counting in favour is, of course, counting in favour of some action being the action one ought to do". But that was not what I meant. This is to specify the *object* of counting in favour, not the *sort* of counting in favour. And it is the wrong object, as Dancy himself explains in a remark quoted immediately afterwards by Rudy. The object of counting in favour is the act, not the rightness of the act. The sort of counting in favour I meant is specified in the definition of favouring I gave above. It contrasts with other sorts, such as the sort referred to in the sentence "The rough terrain favours the terrorists".

One aim of Rudy's paper is to defend me against Dancy's objection in the case of *pro toto* reasons. Rudy says that a *pro toto* reason makes it the case that an act is right, and because it does so, the *pro toto* reason favours the act. According to Rudy, the favouring relation is not the same as the right-making relation, but it can be explained by the right-making relation. I thank Rudy for his support, but I cannot take advantage of it because I think it just *is* the right-making relation.

Rudy correctly points out that his defence of me does not apply to *pro tanto* reasons, because a *pro tanto* reason to *F* does not make it the case that *F*ing is right. From this we can deduce that the right-making relation is not the whole of the favouring relation. The part I have been discussing is given by the first clause in my definition of

JOHN BROOME

favouring. Another part of it is given by the second clause. Roughly, a *pro tanto* reason favours an act by playing a particular role in making it the case that an act is right or that it is not right.

III. DANIEL FOGAL

Daniel Fogal provides us with a new resource for formulating and understanding requirements, including requirements of rationality. Each requirement has a limited domain of jurisdiction; in the case of requirements of rationality this is presumably the domain of all rational creatures. Furthermore not all requirements necessarily apply to everything within their jurisdiction; they may have a domain of application that is smaller than their jurisdiction. Fogal says, for instance, that the jurisdiction of New York law is the whole population of New York, but New York driving laws apply only to drivers in New York.

Requirements are surprisingly hard to formulate, and the idea of a domain of application is a very useful resource to have available in formulating them. As Fogal points out, I did not separate jurisdiction from application in *Rationality Through Reasoning*, and I could have benefited from doing so.

I found Fogal's argument convincing. I have just a few thoughts about it.

First, we should not be dogmatic about formulating requirements. With every requirement that is conditional in some way, there is the question of which conditions to put inside the scope of the requirement, which to put outside and which to put both inside and outside. We should not be dogmatic wide-scopers and insist on putting all the conditions inside, nor dogmatic narrow-scopers and insist on putting none of them inside. We need to consider each requirement and each condition individually.

Secondly, in doing so, we should try not to rely entirely on our intuitions. We may have intuitions about whom a requirement applies to or about the difference between complying with a requirement and evading it, but they are not very stable. For example, Fogal takes it for granted that New York driving laws apply only to drivers

RESPONSES

in New York, but I would not rely on his intuition about this. Suppose you have recently moved to New York from a state that permits turning right on red. Suppose you find you cannot shake your habit of turning right on red, and for that reason you do not drive in New York. You might see this as your way of complying with the law. At any rate, it is not clear to me that the law against turning right on red does not apply to you. The fact that you do not drive because of the law suggests to me that it does. If it clearly did not apply — for instance, if the law made an exception for drivers from states where turning right on red is allowed— then you would drive.

Now let us try out some particular examples of requirements of rationality. Start with the requirement not to have contradictory beliefs. Taking a particular instance, this might be formulated in a narrow-scope way:

Rationality requires of [all x : x is a rational creature] that x believes platypuses are mammals \succ x does not believe platypuses are not mammals.

If this is correct, there must also be another, converse requirement:

Rationality requires of [all x : x is a rational creature] that x believes platypuses are not mammals \succ x does not believe platypuses are mammals.

But that seems odd. It seems there is just one requirement not to have both beliefs. Narrow-scope requirements are supposed to capture the apparent asymmetry in some requirements, but this one is apparently symmetrical. So this wide-scope formulation seems more plausible:

Rationality requires of [all x : x is a rational creature] that (x believes platypuses are mammals \vee x believes platypuses are not mammals) \succ \neg (x believes platypuses are mammals \wedge x believes platypuses are not mammals).

I have inserted an application condition, as Fogal recommends for wide-scope formulae, but I have no intuitive attraction to it. In-

JOHN BROOME

tuitively, I would say than the requirement not to have these contradictory beliefs applies to all of us. If you are inclined by some evidence to believe platypuses are mammals and also inclined to believe platypuses are not mammals, you may suspend judgement and believe neither. You might do this precisely to avoid having contradictory beliefs. This suggests the requirement applies to you even when you have neither belief. But I would not rely on my intuition about this.

Means-end requirements are more symmetrical than Fogal supposes. Take his Grandma example, for which he proposes the wide-scope formula:

(1) Rationality requires of [all x : x is a rational agent] that $((x$ intends to visit Grandma $\wedge x$ believes that driving to Grandma's house is necessary to visit her) $\supset (x$ not intend to visit Grandma $\vee x$ not believe that driving to Grandma's house is necessary to visit her $\vee x$ intend to drive to Grandma's house)).

Let us add some details. Your car is broken and is being fixed. You therefore cannot drive to Grandma's house and do not intend to. But you believe that driving to Grandma's house is a necessary means of visiting her, and for this reason you do not intend to visit her. You deliberately lack the intention of visiting Grandma in order to comply with the means-end requirement. It is very plausible that this requirement applies to you, yet according to formula (1) it does not. To account for your behaviour we would need this different requirement:

(2) Rationality requires of [all x : x is a rational agent] that $((x$ does not intend to drive to Grandma's house $\wedge x$ believes that driving to Grandma's house is necessary to visit her) $\supset (x$ not intend to visit Grandma $\vee x$ not believe that driving to Grandma's house is necessary to visit her $\vee x$ intend to drive to Grandma's house)).

But I doubt there are really two requirements. I would say there is just one, which is symmetrical between intending the end and not intending the means. It may be formulated like this:

RESPONSES

(3) Rationality requires of [all x : x is a rational agent] that ((x believes that driving to Grandma's house is necessary to visit her) \supset (x not intend to visit Grandma \vee x not believe that driving to Grandma's house is necessary to visit her \vee x intend to drive to Grandma's house)).

I have not included either an intention or a non-intention among the application conditions. Whereas I formulated the platypus example in a way that does not accord with my intuition about application, in this case I have gone with my intuition. I think this means-end requirement applies to you whether or not you have the appropriate intention or non-intention. On the other hand, I have included the means-end belief as an application condition.¹ But again, I would not rely on my intuition.

I have an objection to all of these Grandma formulations that is more definite than intuition. Remember first that they are synchronic requirements. For example, the third formula, more fully spelt out is:

Rationality requires of [all x : x is a rational agent] that ((x believes at t that driving to Grandma's house is necessary to visit her) \supset (x not intend at t to visit Grandma \vee x not believe at t that driving to Grandma's house is necessary to visit her \vee x intend at t to drive to Grandma's house)).

Notice next that, if you have an attitude at a time t then, in a temporally relative sense of necessity, you necessarily at t have that attitude at t . From the standpoint of t that attitude cannot be changed. This is a metaphysical truth, so the necessity in question is metaphysical. It is metaphysically impossible, relative to a time t , for you to have at t any attitude other than ones you do have.

Each formula says that, in certain conditions, you are required to satisfy a disjunction. But the application conditions ensure that, in each case, one or more of the disjuncts is metaphysically impossible. So the only way you can satisfy the requirement is by satisfying one of the other disjuncts.

¹ In this case, I agree with Errol Lord (2014, p. 460).

JOHN BROOME

This implies you are required to satisfy one of the other disjuncts. This conclusion follows from a quasi-logical principle I call ‘necessary detachment’ (Broome, 2013, pp. 123-5). Unlike the quasi-logical principles that I mention below in my response the Worsnip, this one is not subject to damaging counterexamples. Indeed, it is proved in Broome (2013, p. 124), given the semantics I assumed there. I believe the proof would extend to Fogal’s system too.

The result is that what is required in the three Grandma formulations can be collapsed into smaller disjuncts by means of necessary detachment. Requirement (1) becomes:

(1)‘ Rationality requires of [all x : x is a rational agent] that ((x intends to visit Grandma \wedge x believes that driving to Grandma’s house is necessary to visit her) \supset (x intend to drive to Grandma’s house)).

That is to say, it reduces to a narrow-scope requirement. I think Fogal would not like this result. He might prefer requirement (3), which becomes:

(3)‘ Rationality requires of [all x : x is a rational agent] that ((x believes that driving to Grandma’s house is necessary to visit her) \supset (x not intend to visit Grandma \vee x intend to drive to Grandma’s house)).

This is less narrow-scope. However, it is still narrow-scope enough to suffer from some of the difficulties of narrow scope that Fogal describes.

My objection might be dealt with in various ways. For one, we could make all the conditions of application into disjunctions. But this seems ad hoc. If they can be varied in this way, we might wonder why application conditions are needed in the first place. For another, we could make the requirements diachronic. We could make the time of the attitudes specified in the application conditions earlier than the time of the attitudes that are required. But then you would be required to satisfy the requirement at a time when you may not satisfy the application conditions, and that seems strange. At any rate, there is a problem to be solved.

RESPONSES

I now want to take up an issue that arises between me on the one hand and both Fogal and Alex Worsnip on the other. This is a matter of terminology. It is not so much a criticism as a plea. I think they would help the progress of philosophy if they would adopt a less misleading terminology. Both Fogal and Worsnip make a distinction between ‘substantive rationality’ and ‘structural rationality’. I make the same distinction using different terms ‘normativity’ and ‘rationality’. I think mine are less misleading.

On the face of it, it should not matter what terms we use. But if philosophers treat normativity and rationality as though they are merely two species in the same genus —two sorts of rationality— bad mistakes can result. They may think the difference is unimportant and slip without noticing from one to the other. Many philosophers are indeed badly confused between rationality and normativity, and this is partly because of a terminological confusion. This confusion is damaging to philosophy.²

Here is an example of the damage done by confusing the terminology of normativity and rationality. David Hume’s statement “‘Tis not contrary to reason to prefer the destruction of the world to the scratching of my finger’ (Hume, 1739-40, Book 2, Part 3, Section 3) is one of his most quoted remarks and also one of the most often reviled. Allan Wood (2013) calls it ‘silly’; Michael Smith (2004) calls it ‘grotesque’. These authors understand Hume to be saying: ‘I can have no reason not to prefer the destruction of the world to the scratching of my finger’ (This is Wood’s (2013) paraphrase). That would indeed be silly or grotesque. But it is not what Hume says, nor what he means.

The sentence itself and all the surrounding text makes it plain that the word ‘reason’ in Hume’s statement refers to the faculty of reason. ‘Rationality’ is a more modern synonym for ‘reason’ in this sense. Hume’s view about rationality is, as he put it, that ‘Reason is the discovery of truth or falsehood’. His remark about finger-scratching simply draws the obvious consequence that reason has nothing to say about preferences. This may or may not be true, but it is definitely not silly or grotesque.

² For the following argument in more detail, see Broome (forthcoming).

JOHN BROOME

Hume is most certainly not using the mass noun 'reason' in its normative sense. It is a bad mistake to suppose he is. This mistake arises from the ambiguity of 'reason' between a rational and a normative sense. We need to reduce this ambiguity. We can do so by making use of the terms we have available, and keeping their meanings distinct as far as possible. If we use 'rationality' to apply to normativity we perpetuate the ambiguity rather than closing it off. We endow 'rationality' with the very same ambiguity as 'reason' has.

Of the four philosophical terms, 'substantively rational', 'structurally rational', 'normative' and 'rational', only 'rational' comes from ordinary English with its original meaning. 'Substantively rational' and 'structurally rational' are obviously artificial. The term 'normative' is also artificial, and unfortunately it has been given different meanings in different disciplines and even in different branches of philosophy. I use it in a sense that is well established in moral philosophy and the philosophy of normativity. It means 'involving ought' or —to make it clearer in this context— 'involving reasons or ought'. Fogal and Worsnip use 'substantively rational' with a meaning that is normative in this sense. Fogal says substantive rationality is a matter of being 'justified or reasonable' and Worsnip that it is 'responding to reasons'. These are normative terms. Fogal and Worsnip use 'structurally rational' for 'rational' with its original meaning.

The normative sense of 'rational' is relatively recent, and the etymological basis of it is weak. True, 'rational' is etymologically an adjective that corresponds to the noun 'reason', so it could in principle be used for 'involving reason' in its normative sense. But that is not its actual origin. It developed in English to mean 'involving reason' in the sense of a faculty. In this original, non-normative sense, the *Oxford English Dictionary* dates it to 1398. It was at first applied only to people and their minds, with the meaning 'having the faculty of reason'. Later, after two hundred years, the dictionary shows 'rational' being applied to a wide range of things, initially including severity, worship, hopes, and sayings. These uses are derivative from the original one. For example, 'a rational severity' means something like 'a severity that might be possessed by a person who possesses the rational faculty'. Derivative uses retain the connection to the rational faculty, but they open the way for an interpretation in terms of normativity. For

RESPONSES

instance, when Daniel Defoe in *Moll Flanders* (1722) writes ‘It did not seem Rational that we would chose to remain here at the Expence and Peril of Life’, a normative interpretation is hard to avoid.

IV. ALEX WORSNIP

I am very glad to find that Worsnip and I are on the same team, as he says. He and I and others who play for this team would do well to agree on a common terminology. I have just proposed we should adopt mine, and I shall use it here. I use ‘rational’ in its original sense, and ‘normative’ to mean ‘involving reason or ought’. I shall go so far as to use ‘Normativity requires you to *F*’ as equivalent to ‘You ought to *F*’. This creates a convenient verbal parallel between requirements of rationality and requirements of normativity.

Our team’s distinctive opinion can be put like this: we distinguish between the requirements of normativity and the requirements of rationality. Members of the opposing team think that the requirements of normativity are the same as the requirements of rationality: that you ought to *F* if and only if rationality requires you to *F*. I call their view the ‘requirement identity theory’.

I have a ‘quick objection’, as I call it, to the requirement identity theory. What rationality requires of you supervenes on your mind, but what you ought does not supervene on your mind. Therefore the requirement identity theory is false.³

I take it to be a conceptual feature of rationality that, when it is a property of a person, it is a mental property of hers. This means that what rationality requires of you supervenes on your mind in two distinct senses. The first is this. Take two situations in which the properties of your mind (excluding rationality) are the same. The properties of the external world may be different. Then, necessarily, rationality requires of you in one situation that you *F* if and only if it requires of you in the other situation that you *F*.

The second sense is this. If rationality requires of you that you *F*, then your *F*ing supervenes on your mind. To spell this out: take two

³ The argument that follows is set out in more detail in Broome (forthcoming).

JOHN BROOME

situations in which the properties of your mind are the same; then, necessarily, you *F* in one if and only if you *F* in the other. Briefly: requirements of rationality are requirements on your mind.

The quick objection is that what you ought does not supervene on your mind in both these senses. One way to oppose the quick objection is to argue that what you ought does supervene on your mind in both these senses. Worsnip thinks this defence is plausible. I do not, and I think his examples of it are unsuccessful.

Worsnip argues that a plausible defence can be mounted on the basis of the view that what you ought is determined by those of your reasons that you have evidence of (These may or may not be all your reasons, he says). Let us accept this view, at least for the sake of argument. I shall call it the 'evidential account of reasons'. He gives two examples of its application. I shall concentrate on the first, which is the better one.

Worsnip asks us to accept for the sake of argument that the fact that 99% of scientific experts agree that human activities have contributed to climate change is a decisive reason for believing human activities have contributed to climate change. Let us accept that and let us also assume you have evidence of this fact. Then according to the evidential account, you ought to believe that human activities have contributed to climate change.

This requirement of normativity—that you ought to believe that human activities have contributed to climate change—supervenes on your mind in the second sense. This is the merit of Worsnip's first, epistemological example. Epistemological requirements of normativity are requirements on the mind. This explains why epistemological normativity resembles rationality more closely than moral normativity does, as Worsnip notes. It is why his first example is better than his second one.

However, even in the first example, the requirement of normativity does not supervene on your mind in the first sense. Compare two situations in which your mind has exactly the same properties but the external world is different. The first is the one described by Worsnip. In the other, it is not the case that 99% of scientific experts agree that human activities have contributed to climate change. Your mind is the same because you are deluded about scientists. In

RESPONSES

this situation, the reason Worsnip mentions for believing that human activities have contributed to climate change does not exist. We may suppose it is not the case that you ought to believe that human activities have contributed to climate change. In one situation you ought to believe this; in the other it is not the case that you ought to believe it. Yet your mind is the same in both situations. Therefore, what you ought does not supervene on your mind. The quick objection remains in place.

The evidential account of reasons is therefore not enough to ensure that the requirements of normativity supervene on the mind. Reasons themselves will have to be properties of the mind. Subjectivism is one way to ensure they are; subjectivism takes reasons to be ordinary properties of the mind such as beliefs and desires. But that is not the way of the evidential account. Instead, the evidential account can adopt a strong sort of externalism about the mind.⁴ In effect, it can take a view of your mind that is expansive enough to encompass your reasons within it. For example, it could identify your evidence with what you know, and take this to be a state of your mind (see Williamson, 2000). On that supposition, your mind is not the same in the two situations of the climate change example. In one situation you know that 99% of expert scientists agree that human activities have contributed to climate change. In the other you do not know this because it is not true. Your mind is different because what you know is different. The example is then consistent with the claim that what you ought supervenes on the mind.

I do not find either subjectivism or this strong sort of externalism plausible. Moreover, neither is enough to establish the claim that what you ought supervenes on the mind. This does not follow even if reasons are properties of the mind. What you ought will normally depend on the probabilities of various possible outcomes, so we need at least an account of how probabilities supervene on the mind. The evidential account of reasons needs to be supplemented by an evidential account of probabilities.

⁴ Interestingly, Fogal, when he too argues that normativity may supervene on the mind, eschews this way of proceeding. I think he may have to fall back on subjectivism.

JOHN BROOME

This turns out to pose a particular difficulty for the evidential account. This account assumes that what you ought is determined by your evidence, but when your evidence is sparse, it may leave probabilities very indeterminate. This leads to problems that I examined in *Rationality Through Reasoning* (Broome, 2013, section 3.4). I found no solution to them within an evidential account. Instead, I found they pushed me towards a Bayesian subjectivism about probabilities. I am intuitively inclined to believe that what you ought is determined by your evidence, but I worry that this view may be unsustainable. Worsnip reports that I see the possibility that normativity supervenes on the mind as a threat; I think this is what he is referring to. But the threat comes from Bayesian subjectivism, not from the evidential account.

Nevertheless, for the reasons I have given, I remain fairly confident about rejecting the claim that what you ought supervenes on the mind. But I may be wrong about that. If I am, I can once again stand next to Worsnip in playing for our team. Worsnip rightly points out that what you ought is not necessarily the same as what rationality requires of you, even if they both supervene on the mind. The quick objection fails, but that does not mean our opponents' case is proven. Worsnip's examples serve well for making this point. We have only to take on board for the sake of argument enough further assumptions to make the examples consistent with the claim that what you ought supervenes on your mind. They clearly show that, even so, what you ought is different from what rationality requires of you.

In part 2 of his paper, Worsnip turns to reasoning, and questions my claim that the correctness of reasoning is a matter of rationality. He is right to question it. I do not understand the relation between rationality and the correctness of reasoning as well as I should.

My account of the correctness of reasoning is this. I take reasoning to be an activity of yours that involves following a rule. The reasoning is correct if and only if the rule is correct, and a correct rule is one that corresponds to a basing permission of rationality. On the other hand, I take rationality to consist, at least to a large extent, in a system of requirements of rationality. I do not understand the rela-

RESPONSES

tion between basing permissions of rationality and requirements of rationality as well as I should.

In *Rationality Through Reasoning*, I made at least one mistake about this. I said (Broome, 2013, p. 190) that a basing permission is nothing other than the negation of a basing prohibition, which (Broome, 2013, p. 186) is just a requirement with a negative content. For instance, to say you are permitted to believe q on the basis of believing p and believing that if p then q is simply to say that you are not required not to believe q on the basis of believing p and believing that if p then q . This is a direct connection between a permission and a requirement, but I now see it is mistaken.

The idea that a permission is the negation of a requirement is not consistent with my adoption (Broome, 2013, p. 116) of what I called the ‘source sense’ of ‘requirement’. An important feature of the source sense is that you are not necessarily required to do everything that is a necessary condition for doing something you are required to do. On the other hand, you are clearly not permitted not to do anything that is a necessary condition for doing something you are required to do. Take an example from Broome (2013, p. 122). Suppose the law requires you not to camp on the streets. I denied it follows that the law requires you, in the source sense, not to camp on the streets on a Thursday. (Why I denied it does not matter here.) But it clearly does follow that the law does not permit you to camp on the streets on a Thursday.

So in this respect I lack a proper account of the relation between permissions, which ground the correctness of reasoning, and requirements, which constitute the heart of rationality.

I lack a proper account in another respect too.⁵ I take it that correct reasoning is a means of improving your rationality. To do this, it must bring you to satisfy a requirement of rationality that you would otherwise not satisfy—that is the only way of improving your rationality. Yet reasoning is made correct by basing permissions of rationality, so how does it manage to bring you to satisfy a requirement of rationality?

⁵ The argument that follows is set out in more detail in Broome (2015).

JOHN BROOME

Only because to each permission there is a corresponding requirement. For example, the *modus ponens* permission:

Rationality permits you that, if you believe at some time that p and if you believe at some time that if p then q , then you believe at some time that q on the basis of those beliefs.

corresponds to the *modus ponens* requirement:

Rationality requires of you that, if you believe at t that p and if you believe at t that (if p then q) and if you care whether q , then you believe q .

Because of the correspondence, reasoning made correct by the *modus ponens* permission can bring you to satisfy the *modus ponens* requirement.

However, there are differences between what is permitted and what is required. In this case there are three differences. First, the permission mentions basing and the requirement does not. Second, the requirement is synchronic and the permission is diachronic; the three beliefs it mentions do not all have to exist at the same time. Third, the requirement contains a 'caring' clause and the permission does not.

The first two differences are systematic. Every permission that supports correct reasoning is a diachronic basing permission, whereas the corresponding requirement is synchronic and not concerned with basing. These differences fit neatly into my account of correct reasoning. But the third difference is particular to the *modus ponens* example. Other requirements do not have a caring clause. But they often do have special features that do not appear in their corresponding permissions. I can find no general principle for determining what these features are. I have no systematic way of deriving a permission from its corresponding requirement. So in this respect too, I do not have a proper account of the relation between permissions and requirements.

Rationality Through Reasoning therefore makes this confession (Broome, 2013, p. 258):

RESPONSES

It is a serious gap in this book's argument that I cannot explain in general how a permission is derived from a requirement. I could seem to be identifying basing permissions simply by the condition that they validate reasoning that is correct.

Not only do I see a problem in connecting permissions with requirements; I must also confess that I am not sure how to formulate basing permissions of rationality correctly. Indeed, partly as a result of Worsnip's criticisms, I am pretty sure that the formulations that appear in *Rationality Through Reasoning* are not correct. I now favour a formulation in which the permission governs a conditional proposition, such as the formulation of the *modus ponens* permission above.

My formulations need more probing, and this one may be incorrect. However, I am not convinced by Worsnip's arguments against them. In criticizing my formulation of the *modus ponens* permission, he relies on questionable quasi-logical principles. For instance, look at what he says about this claim:

Rationality permits you that if you believe that it is raining, and you believe that it is not raining, then you believe that pink elephants will invade China on the basis of your belief that it is raining and your belief that it is not raining.

'If ... then' is to be interpreted as a material conditional. Worsnip says that this claim must be true, but that on the other hand it is not correct to reason from the first two beliefs to the third, so permissions like this cannot be what makes reasoning correct.

Why does he say this claim must be true? Because the permission governs a conditional proposition, and it is definitely permissible for you to make the conditional's antecedent false. So he implicitly assumes that if rationality permits (not p) it follows that rationality permits (if p then q). But this is one of those questionable quasi-logical principles. This one runs up against Alf Ross's famous paradox (Ross, 1941). Suppose it is permissible for you to post a particular letter. This principle implies that it is permissible

JOHN BROOME

for you, if you do not post the letter, to burn it. This inference is questionable.

Worsnip thinks that the paradox can easily be dealt with by adopting a standard contextualist semantics for modals. But I think this semantics simply replicates the old errors of deontic logic, which was discredited decades ago. For so-called 'deontic modals' I recommend instead a neighbourhood semantics that does not generate these questionable principles and their resulting paradoxes (Broome, 2016). At any rate, I cannot be convinced by arguments that depend on such principles.

Worsnip does appear to think it is always correct to reason by *modus ponens*. Whatever we think about rationality, we should surely not think that correct reasoning could be rationally impermissible. So at least rationality permits you to reason by *modus ponens*; I do not think Worsnip could deny this. He cannot be opposed to all general permissions of rationality, therefore. He must recognize rational permissions to reason in particular ways. The only question is how to formulate these permissions.

My approach to formulating them is as follows. The particular way of reasoning in this case is to follow the *modus ponens* rule. I assume that this is a correct way to reason because the rule is a correct one to reason by. I do not think that could be doubted. Then I assume that a rule that is correct to reason by is one that 'corresponds' to a permission of rationality. But the corresponding permission, as I formulate it, is not directly a permission to reason in a particular way. It is a permission to be in the particular state that would result from reasoning in that way.

Possibly this will turn out to be a mistake. It may be that permissions of rationality should be more directly formulated as permissions of reasoning. Worsnip may be right about that. He suggests that they may not be permissions of rationality at all, but *sui generis* principles of correct reasoning, independent of rationality. But there he goes too far. I do not see how it could be so; correct reasoning is of course permitted by rationality. But permissions of reasoning may be more remote from synchronic requirements of rationality than my basing permissions are.

RESPONSES

The problem will be to understand permissions of reasoning as explaining why correct reasoning is correct. Since they apply to nothing else besides processes of reasoning, they may simply record which processes of reasoning are correct. I take it that the correctness of reasoning is explained somehow by its role in promoting our rationality, which is to say in satisfying requirements of rationality. Basing permissions serve better as explanations because they are more distant from reasoning itself, and they have clear connections with requirements of rationality, even if I am not able to work out those connections properly. If we are to recognize permissions of reasoning instead, they will also need to be connected with requirements of rationality.

V. CARLOS NUÑEZ

People who believe they ought to do something often, because they have this belief, intend to do that thing. Why and how does happen? Call this 'the motivation question'. One purpose of *Rationality Through Reasoning* is to answer it in an illuminating fashion. The first part of my answer is that rationality requires you to intend to do what you believe you ought to do, so long as you also believe it is up to you whether or not you do it. We are rational to a large extent, so we often satisfy this requirement. The second part of the answer is to explain why and how we are rational. The explanation is that, for one thing, our subpersonal processes make us rational to an extent and, for another, when our subpersonal process let us down we have a means of improving our rationality. The means is reasoning.

One merit of this answer to the motivation question is that it explains how we can motivate ourselves by our own actions to intend to do what we believe we ought to do. Reasoning is something we do. A second merit is that, in my account of rationality and reasoning, a belief that you ought to do something is just a belief; it is not some different attitude. Noncognitivism is the view that a belief like this, with a normative content, is not really a belief, but some other attitude that embodies motivation in some way. One argument for

JOHN BROOME

noncognitivism is that it offers an answer to the motivation question. By providing a better answer, I aimed to remove one support from noncognitivism.

That is a thin sketch of the argument. Carlos Nuñez asks some questions about it. The first is this. Is the connection I affirm between believing you ought to do something and intending to do it necessary or rational? (This is not how Nuñez puts it. He uses the notion of 'rational necessity', which I find misleading). The answer is 'both'. My statement of the connection is this.

Enkrasia. Necessarily, rationality requires of you that, if you believe you ought to *F* and you believe it is up to you whether or not you *F*, you intend to *F*.⁶

I take this to be a conceptual truth; that is why it is necessary. *Enkrasia* is a conceptual feature of ought. Ralph Wedgwood (2007, chapter 4) takes his similar formula, Normative Judgement Internalism, to determine the meaning of 'ought'. I do not go that far, for a reason explained in Broome (2013, pp. 24-5). But I do take *enkrasia* to be a conceptual feature of ought.

Nuñez next asks 'why is it that, often, when you believe you ought to do something, you reason your way to an intention to do it?'. My answer is, first, that we do not do this particularly often. More often, when you figure out that you ought to do something, you automatically find yourself intending to do it. You think you ought to drink some water and already you find yourself on the way to drink it. Subpersonal processes do this for us.

But we do indeed sometime arrive at the intention by reasoning, and Nuñez is right that I said little about why. I explained how we do it when we do, but not why we bother in the first place. Here is my answer now. Rationality is our plight. Automatic processes make us rational to a high degree. But it is also our plight that we do reasoning from time to time to enhance our rationality. Sometimes it is just that our mind wanders, but often our reason-

⁶ Broome (2013, p. 171). This is a simplified version. I have added 'necessarily' to make explicit what was only implicit in the book.

RESPONSES

ing is focused. If you find yourself interested in a subject, you may cast around to remember what you already believe about it, and draw further conclusions by theoretical reasoning. If you intend an end, at some point you will search out your relevant beliefs about means to it, and may come to intend a means through instrumental reasoning. If you are wondering what to do on a particular occasion, you may engage in some theoretical reasoning to figure out what you ought to do, and then on that basis do some enkratic reasoning to decide to do it.

This explanation of why we do enkratic reasoning (as I call it) is just an explanation of why we do reasoning in general. It has nothing to do with the particular nature or content of the belief you reason with. It appeals to a disposition to reason.

Nuñez correctly recognizes that this is the sort of answer I will give and raises some objections to it. He first suggests it is a ‘thin’ explanation of why we reason. He interprets me as saying we just have a natural disposition to reason from believing we ought to do something to intending to do it. He calls this the ‘enkratic reasoning disposition’. But actually it appeals to the much wider disposition to reason, which is part of our rationality. We are rational creatures. This is not a thin explanation of our behaviour.

This would not have been a successful explanation of why we reason enkratically if I had not situated enkrasia within rationality as a whole. No doubt seeing this sort of explanation coming, Nuñez complains that ‘we still don’t know why it is irrational to believe you ought to do something and yet not intend to do it’. The theory of reasoning is not the right place in my book to look for an answer to this question about rationality. The rationality of enkrasia gets a lot of discussion earlier in the book. Since he asked, I explained above that I take it to be a conceptual truth. This point is set out in Broome (2013, pp. 24-5), and there is more discussion in Broome (2013, section 9.5).

Nuñez also complains that I do ‘not show how such a disposition emerges, in an expected manner, from the natural interplay of the attitudes involved’. Well, the disposition to reason does not emerge from the natural interplay of the attitudes involved, since it is a general disposition that covers reasoning with attitudes of various

JOHN BROOME

different sorts. The disposition to reason enkratically is part of this general disposition. It is a part of the general disposition to reason just because enkrasia is a requirement of rationality. And this requirement of rationality does arise naturally from the natural interplay of the attitudes involved, since it is a conceptual feature of ought that rationality requires you to intend to do something if you believe you ought to do it.

So what is the complaint? In the following pages of Nuñez's paper he deploys the Humean dogma that (to edit Hume a bit) 'belief alone can never be a motive to any action of the will' (cf Hume, 1739-40, Book 2, Part 3, Section 3). What does this mean, exactly? No Humean should deny that a belief can cause an intention; in the complex soup of the mind, 'any thing may produce any thing' (Hume, 1739-40, Book 1, Part 3, Section 15). In *Rationality Through Reasoning* (Broome, 2013, pp. 293-4) I gave an example that evidently did not impress Nuñez, but he should be impressed by this remark of Hume's own. The dogma therefore cannot be that belief alone cannot cause an intention. It has to be that belief alone cannot cause an intention in a particular way that is indicated by the word 'motive'.

Presumably Nuñez takes reasoning to be an example of this particular way. And he is presumably interested in correct reasoning rather than reasoning in general. Since reasoning is a process that goes from premise-attitudes to a conclusion-attitude, for the case of reasoning we can express the dogma like this: if the conclusion-attitude of correct reasoning is an intention, not all the premise-attitudes can be beliefs. I think this is Nuñez's claim.

Why should we believe it? Nuñez says it is not part of the 'characteristic functional role' of a belief to cause an intention by reasoning. This is an odd way of putting what he means to say. 'Characteristic functional role' is causal language. The particular causal process in question is reasoning, and reasoning is a rule-governed process. It would be more natural to ask whether there are rules of inference that take you from beliefs to an intention. Indeed, I think this is what Nuñez means, because near the end of his paper he formulates the question this way:

RESPONSES

The reason you often intend to do what you believe you ought to do is that you infer such intentions from such beliefs. But what is it about these beliefs that would sustain such inferences?

My book answers this question, starting with the presentation of enkrasia in chapter 2. Enkrasia is a requirement of rationality. Later I consider at length how reasoning is related to rationality. I argue that correct reasoning is made correct by basing permissions of rationality, and I consider (admittedly not conclusively as I said in my comments on Worsnip) the connection between basing permissions and requirements of rationality. On the basis of enkrasia, I formulate the enkratic permission, which I take to validate enkratic reasoning.

These arguments are not incontrovertible. Nuñez could have confronted them, but he seems not to have recognized them.

VI. REFERENCES

- Broome J (2013), *Rationality Through Reasoning*, Wiley-Blackwell.
- (2015), 'Synchronic requirements and diachronic permissions', *Canadian Journal of Philosophy*, 45 (2015), pp. 630-46.
- (2016), 'A linguistic turn in the philosophy of normativity?', *Analytic Philosophy*, 57, pp. 1-14.
- (forthcoming), 'Practical reason: rationality or normativity but not both' in *The Routledge Handbook of Practical Reason*, edited by Ruth Change and Kurt Sylvan, Routledge.
- Dancy J (2015), 'Reasons for Broome' in *Weighing and Reasoning*, edited by Iwao Hirose and Andrew Reisner, Oxford University Press, pp. 177-88.
- Hume D, (1739-40), *A Treatise of Human Nature*, London.
- Lord E (2014), 'The real symmetry problem(s) for wide-scope accounts of rationality', *Philosophical Studies*, 170, pp. 443-64.
- Ross A (1941), 'Imperatives and logic', *Theoria*, 7 (1941), pp. 53-71.

JOHN BROOME

- Smith M (2004), 'Humean rationality' in *The Oxford Handbook of Rationality*, edited by Alfred R. Mele and Piers Rawling, Oxford University Press, pp. 75-92.
- Wedgwood R (2007), *The Nature of Normativity*, Oxford University Press.
- Williamson T (2000), *Knowledge and its Limits*, Oxford University Press.
- Wood AW (2013), 'Kant on practical reason' in *Kant on Practical Justification: Interpretive Essays*, edited by Mark Timmons and Sorin Baiasu, Oxford University Press, pp. 57-86.