4
Requirements of rationality

Requirements of Rationality[1]

You have arranged to go to a conference in Montreal on 16–18 October, some months ahead. You firmly intend to go to the whole conference. You write the appointment in your diary, and put it out of your mind for a while. A few weeks later a colleague asks you to a meeting in London on 17 October. You agree to go and intend to do so. For a while, till you discover your mistake, you intend to be in Montreal on 17 October and you also intend to be in London on that day. But, since you know a little geography, you believe you cannot be in both places on the same day. This is the sort of thing that happens to all of us occasionally. It is common, but nevertheless a lapse in our rationality. It is not entirely rational to have two intentions that you believe are incompatible. Rationality requires of you that you do not intend to do one thing and also intend to do another, when you believe you cannot do both.

This is typical of rationality. It requires certain things of you. It requires you not to have contradictory beliefs or intentions, not to intend something you believe you cannot do, to believe what obviously follows from something you believe, and so on. In this book, I shall need some examples of requirements of rationality, specified fairly precisely. The first part of this chapter lays some out. They are examples only; I am not trying to describe rationality completely. One purpose of the examples is to illustrate some general structural features of rationality.

I shall express requirements of rationality using schemata of the form:

> Rationality requires of *N* that *N F*s.

'*N*' stands for the name of a person. In informal statements I shall generally replace '*N*' with 'you'; I am thinking of you as a generic person. '*F*' stands for a verb phrase that specifies what is required of *N*. Generally I shall replace '*F*' with a phrase that itself contains some schematic letters. For example, in requirement (1) on page 73, '*F*' is replaced by 'does not believe that *p* and not-*p*'. So my various formulae are strictly requirement-schemata rather than individual requirements.

Whereas I can replace '*N*' with 'you' for friendliness, even in informal contexts I shall often have to use schematic letters to stand in for verb phrases. I shall have to say such things as 'If you *F* you *G*'. I apologize for this ugly necessity. Unfortunately, English does not contain a fully generic verb. 'Do' is not one, because it cannot stand in for verbs that describe a person's mental state, such as 'believe' or 'intend'. Since I shall be particularly concerned with mental states, I particularly cannot use 'do'. So I am stuck with the letters.

Two other informalities. First, I shall often replace the formula 'Rationality requires of *N* that *N F*s' with 'Rationality requires *N* to *F*'. The former expression has the advantage of making it explicit who is the owner, or bearer, of the requirement. But it is generally safe to use the latter expression, leaving ownership implicit. Second, I shall often call a requirement of rationality 'a rational requirement'.

Whether or not you are rational supervenes on – depends only on – the properties of your mind. I take this to be implied by the concept of rationality. It constrains what rationality can require of you. For one thing, it means that rationality cannot require you to perform a bodily action.[2] If rationality required you, say, to raise your arm, then, necessarily, if you failed to raise your arm you would be irrational. But you might fail to raise your arm because you were prevented from doing so by some external cause that does not affect your mind. For example, you might be momentarily paralysed, without knowing it. In that case, you might fail to raise your arm, yet nevertheless be rational.

From the constraint that rationality supervenes on the mind, it does not follow that rationality can require only mental properties of you. That is because some mental properties are necessary for some non-mental (or only partly mental) properties. For instance, raising your arm is an act, and you cannot perform an act unless your mind possesses some particular property. Different properties are required according to different theories of action. If you are to raise your arm, you might need to intend to raise it, or intend to try and raise it, or want to raise it, or something else, depending on whichever is the right theory of action. Now suppose you are necessarily irrational if your mind has this particular property. Then you are necessarily irrational if you raise your arm. Possibly, then, rationality might require you to *not* raise your arm.

The supervenience constraint implies this: if rationality requires you to *F*, there must be some property of your mind that is sufficient for your *F*ing. The property of not intending to raise your arm (or not intending to try and raise your arm, or whatever) is sufficient for your not raising your arm. So it is consistent with the supervenience

constraint for rationality to require you not to raise your arm. However, there is no property of your mind that is sufficient for your performing a bodily act, so rationality cannot require you to perform a bodily act.

Many authors' accounts of rationality violate the supervenience constraint. They imply that rationality requires you to perform particular bodily acts. I think these authors implicitly assume that external causes do not intervene to prevent your acts. But in a general account of rationality, I prefer not to make this contingent assumption. To avoid any risk of violating the constraint, all my examples of rational requirements will be on mental properties only.

How can we test whether some putative requirement is genuinely a requirement of rationality? It would be nice to have some general criterion to apply, or at least some general method. But I am sorry to say I do not have one. Several philosophers have argued that rational requirements must be somehow inherent in the nature of the mental states they are concerned with. I am sure they are right in some way. For instance, as people say, it is in the nature of beliefs that they aim at truth. Since two contradictory beliefs cannot both be true, that may explain why it is irrational to have contradictory beliefs. Similarly, your intentions seem not to be working properly if you intend an end but do not intend what you believe to be a necessary means to it. However, I do not know how this general idea can be worked out in detail, to provide a criterion for determining what rationality requires.

I can therefore appeal only to our intuitions about rationality, and unfortunately people's intuitions vary. I know of no putative requirements that everyone will agree on. That does not matter very much for my purposes. My examples are meant only to illustrate the sort of requirements that rationality makes on us. If some of them turn out incorrect, or incorrectly formulated, others could serve in their place. So I shall not try very hard to defend the examples.

The first six are weak requirements, which I have selected in the hope that they will satisfy most people's intuition. I shall make some comments on each, by way of explanation and support. Some general conclusions will emerge. Requirements (7) and (8) are listed only to illustrate particular points that will come up in this book. I have no interest in supporting them.

Strictly, what are set out below are requirement-schemata rather than individual requirements. They contain various schematic letters. We obtain individual requirements by making appropriate substitutions for the letters.

Requirement (1). Rationality requires of *N* that *N* does not believe that *p* and not-*p*.

Rationality requires you not to believe a contradiction, that is. This is slightly contentious, because paraconsistent logicians think some contradictions are true, and it is rationally permissible to believe those ones. Paraconsistent logic is far too big a topic for me to broach here, and I shall say nothing against it.

It does allow me to raise a general point about requirements of rationality. This first requirement applies to you (a generic person), but it applies to you an objective standard that is external to you. The requirement applies whatever your own views may be about believing contradictions. It says that rationality requires you not to believe a contradiction even if you are, say, a paraconsistent logician who sees nothing wrong with doing so. Does rationality really require that much?

T. M. Scanlon argues it does not.[3] He distinguishes a broad and a narrow sense of 'rationality', and opts for the narrow sense. I am with him in that; my sense is narrow too. Indeed, I shall argue on page 97 that what he calls the broad sense – responding correctly to reasons – is not a sense of 'rationality' at all. But I think Scanlon specifies the narrow sense too narrowly. He believes that the requirements of rationality, understood narrowly, must issue from a standpoint that is subjective to the person. He wishes to call a person irrational only if she fails by her own standards. But we cannot eliminate objective standards altogether.

I could certainly narrow the first requirement, moving it in Scanlon's direction, by adding a clause. I could make it something like:

> Rationality requires of *N* that *N* does not believe that *p* and not-*p*, if *N* believes a contradiction cannot be true.

But this formula still applies an objective standard. It is by an objective standard that *p* and not-*p* counts as a contradiction. Suppose you believe a contradiction cannot be true, and yet you believe that grass is green and grass is not green. You would not meet the above requirement, yet you might not be failing by your own standards. You might not believe it is a contradiction that grass is green and grass is not green.

Because of that, presumably Scanlon would not count you as irrational, necessarily. To move the requirement further in Scanlon's direction, I could add this further clause: 'and if you believe that *p* and not-*p* is a contradiction'. But still this would not remove all objective standards from the requirement. Whether or not you count as *believing* something depends on the objectively-determined criteria

for belief. Even if, by objective standards, you believe *p* and not-*p*, you yourself might not believe you believe it. So you still might not be failing by your standards.

There is no point in adding more clauses. The conclusion is already plain: we cannot entirely eliminate objective standards from requirements of rationality. If we tried to, rationality would not require anything of anyone; nothing would be irrational. This is an application of Wittgenstein's lesson that we cannot judge whether a person is following a rule by her internal, subjective standards only; if we tried, there would be no rule.

So we cannot fully meet Scanlon's desideratum for a narrow sense of 'rationality'. But we can make our requirements of rationality more or less subjective, and there is a real question how far we should go in the subjective direction.

I find it plausible that my own formulation of the first requirement depends too much on objective standards. A paraconsistent logician fails my requirement if she believes a contradiction, and this seems harsh. It seems harsh to accuse her of irrationality, even if paraconsistent logic is incorrect. So I think it would be plausible to make this requirement more subjective.

On the other hand, we probably should count as irrational someone who believes it is not a contradiction that grass is green and grass is not green. As we move towards more and more outlandish views, somewhere we have to draw a line and count views beyond the line as irrational. I assume this line will be arbitrary to some extent, because I doubt that 'rational' has a very precise meaning. I shall therefore not worry about precisely where to draw it. Each of my requirements implicitly fixes it somewhere, and I recognize I may not have found the best place. I have probably erred too much in the direction of objective standards, just for the sake of simplicity. To make the requirements more subjective, I would need to add more clauses. You should feel free to add more if you think they are needed.

Each of my various requirements of rationality requires a particular thing of you. Rationality as a whole requires many things of you. But Scanlon seems to suggest that we could replace all these separate requirements by a general requirement that you should not fail by your own standards.[4] It would be something like this:

(c) Rationality requires of *N* that *N F*s if *N* believes she* ought to *F*.

However, a single formula like this cannot replace my specific formulae. For example, it could not replace the first requirement, because it does not adequately express the requirement not to believe a contradiction.

Here is why not. Suppose (c) was the only requirement of rationality, and consider a well-informed person who knows it is the only requirement. This person would never believe she ought not to believe a contradiction. If she ought not to believe a contradiction, that could only be because rationality requires her not to believe one; only rationality could be the source of that sort of ought. Being well-informed, she would know this. Therefore, she would believe she ought not to believe a contradiction only if she believed rationality requires her not to do so. But she does not believe rationality requires her not to do so. Instead she believes the formula (c), which tells her only that rationality requires her not to believe a contradiction *if* she believes she ought not to believe a contradiction. So this person will never satisfy the 'if' clause in (c) in this particular case. This means she will never fail to meet requirement (c), even if she believes a contradiction. Therefore, (c) does not correctly express the requirement of rationality not to believe a contradiction.

We cannot replace the many specific requirements of rationality by a single one such as (c). So I shall continue to list a number of specific requirements.

Requirement (2). Rationality requires of *N* that *N* does not believe *p* and also believe not-*p*.

Rationality requires you not to have contradictory beliefs, that is. Paraconsistent logicians will not accept this requirement any more than the first, and I shall not try to defend it against their view.

A putative objection can be drawn from Gilbert Harman's *Change in View*. Everyone has some pairs of contradictory beliefs. Moreover, for most of us it is not worthwhile to weed out of all our contradictory beliefs, and surely rationality does not require us to do so. These points may seem inconsistent with requirement (2). But they are not. Together with requirement (2), they do imply that none of us satisfies all the requirements of rationality, so none of us is fully rational. But then none of us fallible creatures should expect to be fully rational, and we should not resent a requirement that implies we are not. Moreover, just because rationality requires us not to have any particular pairs of contradictory beliefs, it does not follow that rationality requires us to search out all our contradictory beliefs and get rid of them.

Requirement (3). Rationality requires of *N* that, if *N* believes *p* and *N* believe that if *p* then *q*, and if it matters to *N* whether *q*, then *N* believes *q*.

That is: rationality requires you to believe whatever follows by modus ponens from what you believe, if it matters to you.

A simplified version of this requirement would be:

(d) Rationality requires of *N* that, if *N* believes *p* and *N* believes that if *p* then *q*, then *N* believes *q*.

But this version is questionable. Harman rejects a requirement of this sort, on the grounds that rationality cannot require you to clutter your mind with all the trivial beliefs that follow from things you believe.[5] He does not actually reject (d) itself, because he recognizes that some beliefs do not take up space in your mind. For instance, no doubt you have believed for a long time that sheep do not live in the Atlantic Ocean, but this belief has probably not been taking up any space in your mind till now. Many beliefs are dispositions – dispositions to assent or do some other thing – and dispositions need not occupy space. For this reason, Harman applies his 'principle of clutter avoidance' to explicit beliefs only, and not all beliefs. Since (d) is about all beliefs, it does not fall foul of his principle.

But (d) is questionable even so. Robert Audi points out that you may be disposed to believe *q*, without actually believing it.[6] For example, you might be in a position to arrive at a belief in *q* through a quick process of reasoning, but not actually believe *q* because you have not done the reasoning. Plausibly, rationality requires you to be at least disposed to believe *q* if you believe *p* and you believe that if *p* then *q*, but perhaps it does not require you actually to believe *q* in those circumstances.

So (d) may be too strong, and I need a weaker requirement that is less questionable. One possibility would be:

> Rationality requires of *N* that, if *N* believes *p* and *N* believes that if *p* then *q*, then either *N* believes *q* or *N* is disposed to believe *q*.

This may indeed be a requirement of rationality, but I shall not make use of it as an example because it does not suit my purposes later in the book.

I arrive at my own requirement (3) by weakening (d) in a different way. Even if it is not always irrational to fail to believe what follows by modus ponens from what you believe, plainly it is sometimes irrational. For example, suppose that, in going home, you are determined to avoid a dangerous route. Suppose you do not believe South Street is dangerous, and you go home that way. There is nothing irrational in that. However, suppose you have beliefs whose contents imply by modus ponens that South Street is dangerous. Suppose you are disposed to believe South Street is dangerous, because you could easily come to believe it by modus-ponens reasoning from those beliefs. But you have not done the reasoning, and you therefore do not actually believe South Street is dangerous.

Plainly you are not entirely rational, and your irrationality consists at least partly in your not believing South Street is dangerous, when this follows by modus ponens from the contents of beliefs you have. Since it matters to you whether South Street is dangerous, given your determination to avoid a dangerous route, this is plainly irrational.

In general, when it matters to you whether *q*, and you believe *p* and you believe that if *p* then *q*, it is irrational to remain in the state of being merely disposed to believe *q*, rather than actually believing it. I therefore weaken (d) by adding the qualification about mattering to you. This qualification makes requirement (3) very plausible.[7]

The qualification is to be interpreted in a particular way. Mattering to you is to be understood as an attitude of yours. In my example, it is the attitude of being determined to avoid a dangerous route home. It is a feature of your state of mind. We also speak of mattering to you in an objective, external sense. For example, we say it matters to a child that she gets a good education, even if the child very much dislikes being educated. This is not my sense.

Requirement (4). Rationality requires of *N* that, if *N* intends that *e*, and if *N* believes that *e* will be so only if *m* is so, and if *N* believes *m* will be so only if she* intends that *m*, then *N* intends that *m*.

(The star attached to 'she' is to indicate that it is a reflexive pronoun; I shall explain what it means on page ; till then it is safe to ignore it.) This is a wide principle of instrumental rationality. It incorporates a rendering of Kant's remark:

> Who wills the end, wills (so far as reason has a decisive influence on his actions) also the means which are indispensably necessary and in his power,[8]

but it also goes further.

Kant made a slight error; he should have said '… the means *he believes to be* indispensable necessary …', and that is corrected in my formula.

More importantly, my formula does not apply only to means you believe are *indispensably* necessary. If it did, it would rarely have much application. You will rarely believe that a particular means is indispensably necessary to an end you intend; you will nearly always recognize there is more than one way to achieve your end. If you intend to get milk, no doubt you believe you could buy some from a shop, but you should also recognize that you could break into your neighbour's house and steal some. This is such a bad means you will give it no attention. Nevertheless, you do not believe that buying milk from a shop is an indispensably necessary means to getting some. So Kant's formula does not apply to you. On the other hand, you

probably do believe you will not get milk unless you buy it from a shop. Therefore my formula probably does apply to you.

My formula is very far from encompassing the whole of instrumental rationality. It has nothing to say when you believe you have a real choice between alternative means. Nevertheless, it is a wide-ranging requirement of rationality.

It does not even require explicitly that *m* is a means to *e*. *m* might be a consequence of *e*, for example. Of course, rationality does not require you to intend every consequence of something you intend. So my formula would be far too broad if it weren't for the last condition, that you believe *m* will not be so unless you intend it. You believe *m* is 'in your power' as Kant expressed it. Because Kant imposed the condition that (you believe – he should have said) the means should be in your power, he could have dispensed with the condition that (you believe – he should have said) it is indispensably necessary. That is what I have done.

That last condition is anyway necessary. Even if I had required *m* to be a necessary means to *e*, it would have been necessary. Think about this putative requirement:

(e) Rationality requires of *N* that, if *N* intends that *e* and if *N* believes that *m* is a necessary means to *e*, then *N* intends that *m*.

This is incorrect. Suppose you believe the means will happen anyway, whether or not you intend it. Suppose you intend to fly to Venice tomorrow and believe a necessary means of your doing so is that you wake at 0600 tomorrow. But you know you are woken at 0600 every morning by the braying of your neighbour's donkey. You do not need to intend to wake at 0600, because you believe that will happen anyway, without your intending it.[9]

That is one example where (e) fails. But it is perhaps not a very interesting objection to (e), because in this example the necessary means is not an action of yours. So I shall mention a more subtle type of example, which comes from Frances Kamm.[10] Kamm calls it 'triple effect'. Suppose you intend something, and believe some action of yours is a necessary means to it. But suppose you believe this means will be a side effect of something else you intend. Then you have to do it, but you need not intend it.

Here is an example, which is not Kamm's own. You are a doctor and you intend to relieve one of your patients' pain by giving her morphine. You know that you will have to give her so much morphine to relieve her pain that, as a side effect, you will kill her. But you do not intend to kill her. (We have to accept the doctrine of double effect if this example is to work.) Meanwhile, you also intend to admit a new patient to your hospital. Since there are no spare beds, and you

cannot move living patients out of the hospital, the only way you can admit a new patient is by killing an existing one. But that is all right. You do not need to intend to kill one, since you will be killing one anyway, as a side effect of your other intention.

The upshot is that (e) needs to be adjusted by adding the condition that you believe *m* will not be so unless you intend it. Once that condition is added, the condition that you believe *m* is a necessary means to *e* becomes redundant.

In some ways, requirement (4) in the practical sphere is parallel to requirement (3) in the theoretical sphere. Requirement (3) as I stated it contains a condition that it matters to you whether the proposition *q* is true. I explained that, without this condition, the requirement would be questionable. But requirement (4) contains no parallel condition about mattering. When you intend an end, it automatically matters to you what else you need consequently to intend, at least as I mean 'matters to you'. So no mattering condition needs to be added.

Requirement (5). Rationality requires of *N* that, if *N* believes she* ought that *p*, and if *N* believes *p* will be so if and only if she* intends that *p*, then *N* intends that *p*.

Rationality requires you to intend what you believe you ought.

The condition that you believe *p* will not be so unless you intend it appears in requirement (5) for the same reason as it appears in requirement (4).

Requirement (5) makes the bridge between theory and practice. We have many beliefs about what we ought to do. They are often formed by deliberation, and sometimes by explicit reasoning. Reasoning about what we ought to do is not practical reasoning. I take the cognitivist view that normative propositions are propositions like any other. We can believe them in the way we believe other propositions, and our reasoning with our normative beliefs is just like reasoning with other beliefs. It is theoretical reasoning, therefore. However, the point of doing it, and the point of having normative beliefs at all, is to influence what we do. This influence is governed by rationality, and the requirement of rationality that governs it is requirement (5).

Requirement (5) contains the clause 'if you believe you ought to *F*'. The second 'you' in this clause is a reflexive pronoun. I now need to explain what that means and why it is necessary.

In indirect speech, each of the personal pronouns 'you', 'they', 'she', and so on may have either of two different meanings: a reflexive or a nonreflexive one. A pronoun with a reflexive meaning is the representation in indirect speech of a first-person pronoun, 'I', 'me', or 'my', in direct speech. For example, if Hans says 'I shall go to the conference', a correct report of his utterance would be 'Hans says he

will go to the conference', where 'he' has a reflexive meaning.

Reflexive pronouns can appear in other contexts besides indirect speech. Such a context is created by a verb that ascribes an attitude to a person, where the verb's subject designates the person, and its complement designates the object of the attitude. The complement constitutes the context where a reflexive pronoun can appear. If a pronoun appears within the complement and refers to the subject of the attitude-verb, it may be reflexive. For example, in 'Hans intends to bury his donkey' the complement 'to bury his donkey' is a context where a reflexive pronoun may appear. Since 'his' appears in this context and refers to Hans, it may be reflexive.

If it is reflexive, it means that the donkey belongs to Hans, and also that Hans is able to describe his intention correctly using a first-personal pronoun. It implies Hans is in a position to say, correctly, 'I intend to bury my donkey'. A reflexive pronoun indicates a first-personal attitude.

But a pronoun in this sort of context is not necessarily reflexive. In 'Hans intends to bury his donkey', 'his' might be nonreflexive. If it is, it means only that the donkey belongs to Hans. This nonreflexive pronoun would be appropriate if Hans does not know that the donkey he intends to bury is his. To describe his own intention in that case, he might say 'I intend to bury that donkey', pointing to the donkey. We might say 'Hans intends to bury Hans's donkey' or 'Hans intends to bury his donkey', using 'his' nonreflexively.

When a pronoun appears in a context where it can be reflexive, it is very rare for it to be nonreflexive. To avoid ambiguity, I shall make sure it never happens anywhere in this book outside examples that concern donkeys. In this book outside those examples, any pronoun that can be reflexive is reflexive.

In the clause 'if you believe you ought to *F*', the second 'you' is reflexive. So if you believe you ought to *F*, that means you are in a position to say 'I ought to *F*'. Your belief is first-personal. This is essential for requirement (5); the requirement would not be correct if that 'you' was nonreflexive. Suppose you believe the owner of the dead donkey ought to bury it, and suppose you are the owner of that donkey but do not know you are. Then even if you do not intend to bury the donkey, you might be entirely rational. Rationality requires you to intend to do only what you first-personally believe you ought to do.[11]

Requirements (4) and(5) contain reflexive pronouns in other places too. In the condition 'if you believe you will not *F* unless you intend to *F*', the second and third 'you's are reflexive. So are the two 'your's in requirement (4)'s clause 'you believe your *F*ing is a necessary means to your *G*ing', It is easy to see that these beliefs must be first-

personal too, if the requirements are to be correct.

It is often said that practical reason is essentially first-personal. This might mean various things, and not all of them are true. But at least this is true: some of the rational requirements that connect a person's intentions with her beliefs apply to first-personal beliefs only. Requirements (4) and (5) are examples.

Requirement (6). Rationality requires of *N* that, if *N* believes *p* at a time t_1, and if between t_1 and a later time t_2 *N* does not consider whether *p*, then *N* believes *p* at t_2.

Rationality requires you to drop a belief only after consideration. I include this as a plausible example of an intertemporal requirement; all my other examples are requirements on your mental states at a particular time. The formula says 'consider' rather than 'reconsider' because I am allowing for the possibility that your original belief in *p* was acquired without consideration.

Requirement (7). Rationality requires of *N* that, if *N* prefers *a* to *b* and *N* prefers *b* to *c*, then *N* prefers *a* to *c*.

Requirement (8). If *p*, *q* and *r* are mutually contrary propositions such that (*p* or *q* or *r*) is necessarily true, rationality requires of *N* that the degree to which *N* believes *p*, the degree to which *N* believes *q*, and the degree to which *N* believes *r* together add up to one.

I mention these as typical examples of the putative requirements that are found in decision theory and Bayesian probability theory. I do not wish to defend them.

Very many more requirements of rationality could be added to my eight, but I am stopping now. A significant feature of my examples is that most of them require particular relationships to hold among your propositional attitudes; they do not require you to have any particular individual attitude. This is typical of rational requirements. Typically, 'rationality requires' governs a logically compound statement that describes relations among your attitudes. It has a *wide scope*, as I shall sometimes put it.

Not all requirements of rationality are of this sort. My requirement (1) is an exception; it requires you to have the individual attitude of not believing a contradiction. But exceptions are rare.

The wide scope of rational requirements is very plausible. I think everyone will agree that whether or not it is rational for a person to have a particular attitude almost always depends not just on what attitude it is, but also on other features of the person. It may be irrational for me to believe something that you believe quite rationally, because this belief does not cohere properly with my other

beliefs, whereas it coheres with yours. So when we say an attitude of a person is rational or irrational, we generally mean it is rational or irrational in relation to others of her attitudes. It can also happen that an attitude of yours is rational in relation to some of your attitudes and irrational in relation to others of your attitudes.

'In relation to' is an inexact notion. My formulae for rational requirements make the relative nature of the requirements explicit, and they say precisely what sort of relation among your attitudes is required.

Nevertheless, although we recognize the relative nature of requirements of rationality, we very commonly speak of rationality in a way that is not explicitly relative. We commonly say that a single attitude is rational or irrational. It is natural to assume we are entitled to these nonrelative statements, even if rationality is fundamentally a relative matter.

In particular, it is natural to think that narrow-scope requirements of rationality, which apply to a single attitude, may be derived from the wide-scope requirements in particular circumstances. This sort of derivation is called *detachment*. Various putative patterns of inference appear to make detachment possible. I now turn to examining some of these patterns.

The validity of an inference is always bound up with the semantics of its terms. So in examining validity I shall also be probing the meaning of 'rationality requires'. As well as probing it, to some extent I shall be sharpening up the specific meaning I give this term. But I shall stick close to its intuitive meaning.

For convenience in this discussion, I am going to give the terms 'rational' and 'irrational' more rigid meanings than they have in real life. By 'rational' I mean fully rational, and by 'irrational' I mean not fully rational.

Here is one putative detachment inference. You might think at first that, from

(f) Rationality requires of you that, if you *G*, you *F*

and

(g) You *G*

it would be valid to infer

(h) Rationality requires of you that you *F*.

This would allow a narrow-scope requirement that you *F* to be derived from a wide-scope requirement that you *F* if you *G*. I shall this *contingent detachment* because it is premised on the contingent

proposition that you *G*.

Here is an example of contingent detachment. According to requirement (3), rationality requires of you that you believe *q* if you believe *p* and you believe that if *p* then *q*, and if it matters to you whether *q*. Suppose you do actually believe *p*, and you believe that if *p* then *q*, and it does matter to you whether *q*. Then according to this pattern of inference, we could derive the conclusion that rationality requires of you that you believe *q*.

You might at first think that 'Rationality requires of you that you *F*' means 'If you are rational, you *F*', or equivalently 'You are irrational if you do not *F*'. If 'rationality requires' had this meaning, contingent detachment would be valid. (f) would mean that, if you *G* but do not *F* you are irrational. Given (g), that you *G*, it would follow that if you do not *F* you are irrational. That is to say, under this interpretation, that rationality requires of you that you *F*. So we have the conclusion (h).

But actually 'Rationality requires of you that you *F*' cannot mean 'If you are rational, you *F*'. To show that rationality requires you to *F*, it is not sufficient to show you are irrational if you do not *F*. Take a person who is irrational anyway; perhaps she has some pair of contradictory beliefs. Whatever *F*ing may be, it is vacuously true of this person that she is irrational if she does not *F*. If 'rationality requires' had this meaning, rationality would require absolutely everything of this person. Indeed, since in practice all of us are in her position – none of us rational in every respect – it would require absolutely everything of all of us.

So contingent detachment cannot draw support from this interpretation of 'rationality requires'. In any case, contingent detachment is intuitively unattractive, at least if requirement (3) is correct. Suppose you believe some propositions *p* and *q*, and you believe (if *p* then *r*) and (if *q* then not *r*). You might have beliefs of this form, though you would probably be irrational if you did. Suppose also it matters to you whether *r*. Then requirement (3), together with contingent detachment, implies that rationality requires you to believe *r*. It also implies that rationality requires you to believe not *r*.

These two conclusions are not credible together. It is not credible that rationality should require you to believe a proposition and also require you to believe its negation. Quite apart from anything else, this conclusion would be extremely hard to reconcile with requirement (2) of rationality, which is that rationality requires you not both to believe a proposition and believe its negation. And requirement (2) is very plausible.

The incredible conclusion I derived from contingent detachment results from applying it to requirement (3) of rationality. Requirement (3) may not be exactly correct, but something like it must be correct,

and any similar requirement would have a similar incredible implications when combined with contingent detachment. So we should reject contingent detachment.

We might next try interpreting 'Rationality requires of you that you *F*' as meaning that, necessarily, if you are rational, you *F*. More colloquially: you must *F* if you are to be rational. Put another way: *F*ing is a necessary condition for you to be rational. This interpretation defines 'rationality requires of you' in terms of 'you are rational'. It treats the property of being rational as analytically prior to the notion of rational requirement.

There is a grammatical justification for this interpretation. In one of its senses 'rationality' is just the name of the property of being rational. To have rationality means to be rational. So in this sense, 'Rationality requires of you that you *F*' means the same as 'Being rational requires of you that you *F*'. That in turn can be taken to mean 'Your *F*ing is necessary if you are to be rational' or 'Necessarily, if you are rational you *F*.

It is worth putting this interpretation slightly more formally, in terms of the standard semantics for modal logic. A proposition is a subset of the set of all possible worlds; intuitively it is the set of worlds where the proposition is true. The proposition 'You are rational' is the set of worlds where you are rational. If and only if you *F* at every one of these worlds, rationality requires of you that you *F*.

The proposition that you are rational may be relative to the actual world, because the conditions for your being rational may depend on how things actually are. Technically, there is a function from worlds to sets of worlds. For each world w, this function picks out the set of worlds where you are rational if w is the actual world.

A further formal point for those who are interested in such things. Suppose we were to generalize the interpretation by giving a meaning to all sentences of the form 'Rationality requires of you that *p*', for any proposition *p*. Suppose we took these sentences to mean that, necessarily, if you are rational, *p*. This would give 'rationality requires of you that' exactly the logic of the operator '*O*' in standard deontic logic. I have been restricting 'rationality requires of you that' to govern propositions of the specific form that you *F*, and I shall continue to do so. With this restriction, 'rationality requires of you that' conforms to standard deontic logic restricted to propositions of this form.

On this interpretation, contingent detachment is invalid. From (f) and (g), we can derive that you are irrational if you do not *F*, but we cannot derive that necessarily you are irrational if you do not *F*. So we cannot derive (h).

On the other hand, this interpretation does validate a different pattern of inference. From

(i) Rationality requires of you that, if you *G*, you *F*

and

(j) Rationality requires of you that you *G*,

it permits us to derive

(k) Rationality requires of you that you *F*.

Let us call this *rational detachment*. It is valid under this second interpretation for the following reason. Necessarily, if you do not *F*, there are two possibilities: either you do not *G*, or you do *G* but do not *F*. And the premises tell us that, necessarily, in either case you are irrational. So, necessarily, if you do not *F* you are irrational. That is to say, under the interpretation, that rationality requires of you that you *F*.

Rational detachment is intuitively unattractive. Suppose the premises (i) and (j) are both true. According to rational detachment, rationality requires of you that you *F*. But now suppose that – irrationally – you do not *G*. Then intuitively there is nothing irrational about your not *F*ing. Premise (i) says you are required to *F* if you *G*, but since you do not *G*, you are not failing to satisfy this requirement by not *F*ing. Your failure is in not satisfying the requirement expressed in premise (j), but this failure consists in your not *G*ing. It has nothing to do with your not *F*ing.

To be sure, necessarily, if you do not *F* you are irrational. But it does not follow that your not *F*ing would be itself irrational.[12] Your not *F*ing may simply indicate that you are irrational, and your actual irrationality might be elsewhere. That is so in my example: it is your not *G*ing that is irrational. We have to recognize a difference between 'If you do not *F* you are irrational' and 'Your not *F*ing would be irrational'. Given this difference, it intuitively most natural to take 'Rationality requires of you that you *F*' to be equivalent to the latter rather than the former. This makes the inference from (i) and (j) to (k) intuitively unattractive.

Intuitively, things that are necessary if you are to be rational are not all of them things that rationality requires of you. Here is an example of a different type. Necessarily, if you are rational you are alive. But it would be odd to say that rationality requires of you that you are alive.

I explained that grammar can justify us in interpreting 'Rationality requires of you that you *F*' to mean that, necessarily, if you are rational you *F*. It takes 'rationality' to be the name of the property of

being rational. Analogously, 'stoutness' is the name of the property of being stout, and 'Stoutness requires of you that your have a girth of more than 90 cm' means that, necessarily, if you are stout you have a girth of more than 90 cm. But actually, this grammatical construction, in which the name of a property is followed by 'requires', is rather strained. 'Stoutness requires' has to be read that way because 'stoutness' has no other meaning than the property of being stout. But 'rationality' has another meaning; it is also the name of a code. In the construction 'rationality requires', 'rationality' is more naturally read with this other meaning, because the construction is rather strained if it is read as the name of a property.

By a *code* I mean a collection of rules. The code requires you to satisfy the rules. 'Rationality requires of you that you *F*' is more naturally read to mean that the code of rationality contains the rule that you *F*. The rules in the code may be called 'the requirements of rationality'.

Different codes have different structures. Perhaps there are some codes such that whatever is necessary if you are to satisfy all the rules of the code is also among those rules. But for most codes, some things that are necessary if you are to satisfy all the rules of the code are not themselves among those rules. For example, to satisfy all the rules of chess, you must play chess. But it is not among the rules of chess that you play chess.

Rationality belongs to the majority of codes in this respect. Some things are necessary if you are to satisfy all the requirements of rationality without themselves being among those requirements. For example, your being alive is necessary if you are to satisfy the requirements of rationality, but it is not a requirement of rationality that you are alive. So if we read 'rationality' in 'rationality requires' as the name of a code rather than the name of a property, rational detachment is not valid. Moreover, this it the natural way to read it, because the alternative interpretation of 'rationality requires' is rather strained. That explains why rational detachment is intuitively unattractive.

I shall stick to the intuitive interpretation, and take rational detachment to be invalid.

If you are rational – if you have the *property* of rationality – I take that to mean you satisfy all the requirements of rationality. So I define the property of rationality in terms of requirements of rationality; the requirements are analytically prior to the property.[13]

Here is my proposal described in terms of possible worlds. There are a number of rational requirements on you. What each requires of you is a proposition of the form: that you *F*. It is the set of worlds

where you *F*. What rationality requires of you therefore constitute a set of propositions – a set of sets of worlds, that is to say. Call this the set of required propositions. The set of worlds where all the required propositions are true is the proposition that you are rational. It follows that, if rationality requires of you that you *F*, then you *F* at all worlds where you are rational. However, even if you *G* at all worlds where you are rational, the proposition that you *G* is not necessarily required of you.

The set of required propositions may be relative to the actual world: how things actually are may affect what rationality requires of you. Technically, there is a function from worlds to sets of sets of worlds. For each world w, this function picks out the set of propositions that are required of you at w.

This new interpretation of 'rationality requires' does not validate either contingent or rational detachment. However, it does not itself imply these inference patterns are not valid. That depends on the structure of the code of rationality. The code might have various structural features. These will appear in this semantic structure in the form of conditions on the set of required propositions.

For example, suppose the code has the structural feature that rationality requires you to *F* whenever it both requires you to *G* and requires you to *F* if you *G*. Then, if the set of required propositions contains the proposition that you *G* and the proposition that you *F* if you *G*, it contains the proposition that you *F*. If this is so, rational detachment is valid. But I have already explained that rational detachment is not intuitively attractive, so we should not expect the set of rational requirements to have this particular structural feature. Nor should we expect its structure to validate contingent detachment, since I explained that contingent detachment is also intuitively unattractive.

A more plausible condition is that rationality does not place contradictory requirements on you. That is to say, if the set of required propositions contains the proposition that you *F*, it does not contain the proposition that you *G*. Almost as plausible is the condition that what rationality requires of you is consistent. That is to say, the proposition that you are rational is not empty.

The semantical structure I have given validates a pattern of inference I have not mentioned yet. From

Rationality requires of you that, if you *G*, you *F*
and Necessarily you *G*
derive Rationality requires of you that you *F*.

Let us call this *necessary detachment*. The following argument shows it is valid in my semantics. Suppose you *G* necessarily; then you *G* in every world. The proposition that, if you *G*, you *F* is the set of worlds

in which either you *F* or you do not *G*. But there are no worlds in which you do not *G*, so this is just the set of worlds in which you *F*. That is to say, the proposition that, if you *G*, you *F* is the same as the proposition that you *F*. Therefore, if rationality requires of you that, if you *G*, you *F*, it requires of you that you *F*.

Is necessary detachment intuitively attractive? To settle that question, I would need an example of this inference, but I cannot find one. I could find one if I adopted a very relaxed notion of necessity. Suppose we count a true proposition as necessary if nothing can now be done that would make it false. This would make necessity relative to a time, and it would make events that occur at a particular time and states that exist at a particular time necessary relative to any later time. This is the relaxed notion I have in mind.

Suppose you believe a proposition p at a time t_1 and you do not consider whether p between t_1 and a later time t_2. This conjunctive fact is necessary relative to t_2, in the relaxed sense. Then requirement (6) of rationality, together with necessary detachment, would yield the conclusion that, relative to t_2, rationality requires you to believe p at t_2.

This example implies a sort of bootstrapping:[14] just because you believe something at one time, it is supposed to follow that rationality requires you to believe it at a later time. This is implausible. I think it tells us to reject necessary detachment under the very relaxed, time-relative interpretation of necessity. But we can treat this as a reason to reject that interpretation of necessity rather than necessary detachment itself. That is how I shall treat it; I shall reject the relaxed notion of necessity.

If I adopt a less relaxed notion, I cannot find an example of a mental state you are necessarily in. So, since my semantics implies it, and since I can find no intuitive objection to it, I am willing to let the inference pattern of necessary detachment stand.

From all this discussion of detachment I draw the conclusion that contingent and rational detachment are both invalid. Necessary detachment may be valid, but I know no examples of it. I therefore continue to think that requirements of rationality have a wide scope. I started to examine detachment because if potentially offered a way to derive narrow-scope requirements from wide-scope ones. But I conclude it does not.

My interpretation gives logical priority to the notion of rational requirement. That is why I started my account of rationality by listing rational requirements. But under any interpretation, it is a good idea to concentrate on requirements because they pick out necessary conditions for rationality. Under any interpretation, if rationality

requires you to *F*, then *F*ing is a necessary condition for your rationality. Discovering necessary conditions is a good way to delimit the notion of rationality. It would also be useful to discover sufficient conditions, but those are extremely hard to find. I do not know any.

Some authors ask, not what rationality requires of you, but what would be rational for you. To say that *F*ing would be rational for you is to specify neither a necessary nor a sufficient condition for your rationality. It is to say it would be possible for you to *F* and be rational. That is: your not *F*ing is not a necessary condition for your rationality. To say what is not a necessary condition is generally less useful than to say what is a necessary condition.

Notes

1 This chapter owes a lot to comments I have received from Niko Kolodny and James Morauta, and to Kolodny's paper 'Why be rational?'.
2 I take this point from Andrew Reisner's *Conflicts of Normativity*.
3 *What We Owe to Each Other*, pp. 25–30.
4 *What We Owe to Each Other*, p. 25.
5 *Change in View*, p. 12.
6 'Dispositional beliefs and dispositions to believe'.
7 The qualification is what Harman in *Change in View*, p. 55, calls an 'interest condition'; in adding it I am following Harman.
8 *Groundwork of the Metaphysic of Morals*, pp. 80–1.
9 I think Wlodek Rabinowicz first made this point to me.
10 'The doctrine of triple effect'.
11 Christian Piller pointed this out to me. See his 'Normative practical reasoning', p. 202.
12 Niko Kolodny and James Morauta have both separately made this point to me.
13 Niko Kolodny recommended this order of priority to me.
14 I believe the term 'bootstrapping' was made popular by Michael Bratman in *Intention, Plans and Practical Reason*, pp. 24–7.