

Incommensurable Values

John Broome

University of St Andrews

For *Well-Being and Morality: Essays of Honour of James Griffin*, edited by Roger Crisp and Brad Hooker, Oxford University Press

I have received penetrating comments on this paper from Ruth Chang, Ingmar Persson, Franz Prettenthaler, Jonathan Riley, John Skorupski and the editors of this volume. The comments are so penetrating, and the subject so difficult, that I been able to take account of only a few of them in this paper. They set me a programme for future work.

1. The idea of incommensurability

You might join the army or you might become a priest. Which would be better? Intuitively it seems this question may have no determinate answer; the values realized by these two careers seem to be so very different that they cannot be weighed against each other in a precise way. In some circumstances there will be a determinate answer to the question – for instance if you do not believe in God and like guns. But in more balanced circumstances, it will not be determinate which is the better option. This phenomenon of indeterminacy is often called the ‘incommensurability’ of values. It is often thought to be a central feature of ethical life.

James Griffin has taken up the subject of incommensurability several times during his career.¹ He is not enthusiastic about its importance, and even casts some doubt on its existence.² He points out, first, that what many philosophers have called ‘incommensurability’ is not really that at all, or ought not to be called that. ‘Incommensurability’ ought to be reserved for cases where alternatives are ‘incomparable’ as Griffin puts it, by which he means that they cannot be put in an order. When a philosopher says the value of free speech is incommensurable with the pleasure of eating pizza, she means that free speech is immeasurably more valuable than this pleasure. That is to say, free speech and this pleasure *can* be ordered, emphatically. A small amount of free speech is better than any large amount of pizza-pleasure. This is not really incommensurability but extreme commensurability. Griffin mentions other misuses of the term too. If we set the misuses aside, and concentrate on true incommensurability, he thinks it may be hard to find any.

So values are commensurable in Griffin's sense if alternatives can be ordered. It is not the values themselves that can be ordered. Griffin is not interested in ordering enjoyment and understanding – to take two of his examples of values. These values cannot be ordered – neither is better than the other – because they come in various amounts. Wren's enjoyment in designing St Paul's is better than understanding Article iv of Section 3.8 of Regulation 294 issued by DG65 of the European Union, whereas understanding evolution is better than enjoying a pizza. So we cannot order the values of enjoyment and understanding themselves. But we might be able to order events that realize these values, such as designing St Paul's or eating a pizza. Each event realizes one or more values to some degree. If the values are commensurable, that means they can be measured on the same scale. So a degree of one can be compared, as greater or less, with a degree of the other. Consequently the events can be ordered. We often need to compare other things besides events, so for generality I shall say that ‘options’ can be ordered, without specifying what sort of options. To claim that values are commensurable is to claim that options can be ordered.

Ordered with respect to what? Their goodness. Values are goods of different sorts, which contribute to the goodness of the options that realize them. To say that values are commensurable is to say that options are ordered by their goodness.

Griffin himself sometimes seems not to have recognized clearly enough that it is options rather than values themselves that need to be ordered. As part of his argument against incommensurable values, he mentions that we can and do compare values. ‘We

can and do compare pain and accomplishment. If the pain is great enough and the accomplishment slight enough, we should not consider this accomplishment worth the pain.³ But this is beside the point. Undoubtedly, *some* amount of pain-avoidance can be compared with *some* amount of accomplishment; indeed there may be no two values such that some amount of one cannot be compared with some amount of the other. Perhaps all values are commensurable in this sense. But we are not interested in comparing values themselves; we need to compare options that realize values. Options realize values to various degrees. A great accomplishment achieved with a little pain is better than painlessly accomplishing nothing. And painlessly accomplishing nothing is better than very painfully accomplishing very little. But if pain and accomplishment are to be thoroughly commensurable, *any* amount of pain-avoidance must be comparable with *any* amount of accomplishment, so that all options that realize pain-avoidance or accomplishment to any degree can be ordered by their goodness.

Still, despite this slip, Griffin does correctly identify the question of commensurability as a question of ordering options. Actually, the ordering of options is a more fundamental and more general matter than commensurability. To ask whether values are commensurable is implicitly to make a presumption about the way individual values combine together to determine the overall goodness of an option. It is to presume that the goodness of an option is determined by the various values it realizes, acting independently of each other. If that is so, and if these independent values are commensurable – so they can be measured on the same scale – then they can be added up to determine the overall value of the option. But actually the value of an option may be determined in a complex fashion by the interaction of values with each other and with other features of the option that are not themselves values. If values and other features interact like this, the notion of commensurability is obscure. But the question of ordering remains clear. However complex the determination of goodness, we can ask whether options are ordered by their goodness. This question of ordering is the subject of this paper, though I shall use the terms ‘commensurability’ and ‘incommensurability’ loosely to refer to it.

Actually, I shall be asking two questions. First, are options always ordered by their goodness? Second, if they are not, how important is that fact in ethics? Section 4 aims to answer the first question, and section 10 returns to it; sections 5 to 10 aim to answer the second. But before coming to the two substantive questions, I need to deal with some formal matters in section 2 and 3. I have suggested options might or might not be ordered by their goodness. What, exactly, is the difference?

2. *The standard configuration*

I find a particular formal device useful for answering this question; I call it a ‘standard configuration’.⁴ Take the career example again. You have a choice between joining the church and joining the army. In practice there is bound to be a great deal of uncertainty about how each of these careers will progress if you choose it. But for simplicity let

us ignore that; imagine each can be accurately predicted. So you have a choice between a particular career as a soldier, which is mapped out in detail, and a particular career as a priest, also mapped out. Now, imagine variations in the church career that make it better or worse. For instance, imagine the community becomes more religious, which makes for a better career, or less religious, which makes for a worse one. Think of a whole sequence of imaginary church careers, each a little bit better than the previous one in the sequence. I shall call this sequence a 'chain'. Next, fill in the spaces in the chain until we have fully continuous variations in goodness as we move along it. For instance, imagine continuous variations in the community's religiousness. Then our chain can be represented as a continuous line like this:

← worse better →
Possible careers as a priest

Each point on the line stands for a particular possible career. Each career on the line is better than every career that lies to its left. A standard configuration consists of a continuous chain like this, together with a single alternative that I call the 'standard'. In our example, the standard is the career available to you in the army. We hold that constant, while comparing it with all the various careers in the chain. In a standard configuration, the chain must be sufficiently long that the worst options in the chain are worse than the standard and the best options are better than the standard.

The choice of careers is only an example. Standard configurations can be drawn up in other contexts too. Suppose you are choosing between dinner and a film. You can hold the film constant as standard, and form an imaginary chain of dinners that varies continuously from grim to superb. In general, then, a chain is like this:

← worse better →
Possible options

As we move from left to right along a chain, two things are essential for my purposes. First, there is improvement all the way: every option is better than every option to its left. Second, the improvement is continuous. Joseph Raz thinks this may not always be possible.⁵ He suggests some values may not be continuously variable; they can only change in discrete steps. Myself, I can think of no values that are discrete like this. But even if there are some, we can still always change the goodness of an option in a continuous fashion, because options need not realize only one value. We can simply add or subtract amounts of some value that is not discrete. For instance, the service at dinner might be quicker or slower; that is a quality that can certainly vary continuously. I am sure there are at least some continuous values. So I see no difficulty in assuming the chain is continuous.

3. Definitions

Now we have the standard configuration, let us start to compare the options in the chain with the standard – in the example, the church careers with the army career. At the left of the line there are options that are worse than the standard; they form the ‘worse zone’, as I shall call it. At the right are options that are better than the standard; they form the ‘better zone’. What happens in between? There are three possible cases. There may be more than one option between these zones – between options that are better than the standard and those that are worse. There may be just one. Or there may be none. This third case is genuinely possible, but it is not important for this paper and I shall say no more about it.⁶ So we have two cases to consider.

In the first, where more than one option lies between the worse zone and the better zone, there must be many that do. If there are two options between those zones, any option in the chain that lies in between these two options must also lie between the zones. So there will be an intermediate *zone* of options between the worse and better zones. The cases we need to consider, then, are, first, where there is an intermediate zone and, second, where there is just a single intermediate option.

Take some intermediate option. Suppose it happens to be equally as good as the standard. Then any option better than it (to its right on the line) would be better than the standard, and any option worse than it (to its left) would be worse than the standard. It would therefore be the only intermediate option between those that are better than the standard and those that are worse. So if there is more than one intermediate option – an intermediate *zone* – no option within this zone can be equally as good as the standard.


In this argument, I assumed that if one option is better or worse than another, and the other is equally as good as a third, then the first is, respectively, better or worse than the third. This is part of the meaning of ‘equally as good as’. (It is an extension of transitivity.) It prevents any option in an intermediate zone from being equally as good as the standard.

But if there is only a single intermediate option, nothing of this sort prevents it from being equally as good as the standard. In this case, indeed, I think it must be equally as good as the standard. The slightest improvement or deterioration in this option would make it better or worse than the standard. If a choice is so finely balanced that the slightest change in one of the alternatives can make the difference between its being worse and its being better than the other, surely the alternatives are equally good.

Indeed, I think we may *define* a single intermediate option to be as good as the standard. I propose this as a definition of ‘equally as good as’: one option *A* is equally as good as another *B* if and only if any possible option that is better or worse than *A* is also, respectively, better or worse than *B*, and any possible option that is better or worse than *B* is also, respectively, better or worse than *A*. (This definition does not state explicitly that *A* is neither better nor worse than *B*, because that is implied by the rest.) This is an acceptable definition because it meets the formal condition that ‘equally as good as’ must satisfy as a semantic matter: it must be an *equivalence*

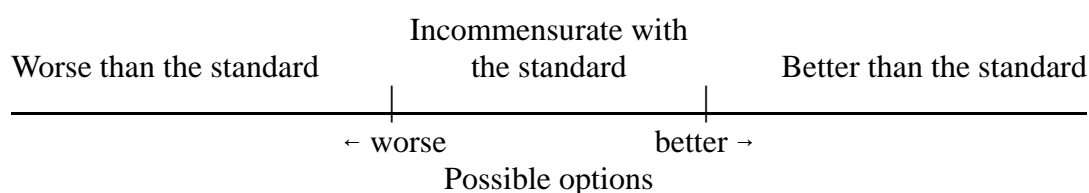
relation – it must be transitive, reflexive and symmetric.

The definition has the effect that, when there is just a single intermediate option, it must be equally as good as the standard (except in some peculiar cases). The picture is, then:



Let us call this 'the equality case'.

Next, let us define one option *A* to be 'incommensurate' with another *B* if and only if it is not the case that *A* is better than *B*, and not the case that *B* is better than *A*, and not the case that *A* and *B* are equally good. We have found that, if there is an intermediate zone, all the options in it are incommensurate with the standard. Let us call this the 'incommensurate case'. Its picture is:



To sum up: we only have one primitive relation to deal with: 'better than'. 'Equally as good as' and 'incommensurate with' are defined in terms of that. The difference between them is fixed by the structure of the 'better than' relation. This relation picks out the zone of options that are better than the standard and the zone of options that are worse than the standard. If it leaves only one option in a chain intermediate between these zones, this option is equally as good as the standard. If it leaves an intermediate zone, the options in this zone are incommensurate with the standard. The difference between equality and incommensurateness is simply whether there is one intermediate option or many. There is no further, intuitive distinction between the two.⁷

My definition of 'equally as good as' may seem unsatisfactory, because it may seem to deny a significant possibility. When there is only a single intermediate option between the worse and better zones, there may seem to be two possibilities for it: it might be equally as good as the standard or it might not be. When I defined it to be equally as good as the standard, I was denying the second possibility.

However, I do not think this is unsatisfactory; I think it is correct. I see no genuine second possibility. What real difference could there be between a case where this intermediate option is equally as good as the standard and a case where it is not? Presumably, the difference would have to show up in some practical way: it would have to make some difference to your decision making, if you ever had to make a

choice between this option and the standard. Now, it certainly makes a practical difference whether there is a single intermediate option or a whole intermediate zone; I shall explain why in section 8. But when we are anyway dealing with just a single intermediate option, I see no real difference. Various things would be true if you had to make the choice between the intermediate option and the standard: it is not the case that you ought to choose the intermediate option, nor that you ought to choose the standard; choosing either would not be wrong; it is not determinate which you ought to choose; and so on. But all of these things are true simply because neither alternative is better than the other; none of them indicates a difference between their being equally good and their not being equally good. I shall say more about this in section 7.

4. Is there incommensurability?

Having set up the definitions, we can turn to substantive questions. The first question I mentioned at the end of section 1 is: are options necessarily ordered by their goodness? As I have now formalized it, this question has become: does the incommensurate case really exist, or is it only a technical possibility? Let 'commensurabilism' be the view that it does not exist. We can now see that commensurabilism is implausible. It is very stringent because it requires a sharp division between options that are worse than the standard and those that are better. There is at most one intermediate option. But this is implausible in, for instance, the careers example. It is implausible that a tiny change in the conditions of a church career could make all the difference between this career's being worse than soldiering and its being better. This example seems to have an intermediate zone; it seems to fit the incommensurate case.

At first, one might think commensurabilism could be supported by the epistemicist theory of vagueness. Take a vague predicate such as 'red'. Most people believe there are borderline cases between things that are red and things that are not red. Of these borderline cases, we cannot say definitely that they are red, nor that they are not red. But *epistemicists* about vagueness believe there is a sharp boundary between red and not red, despite appearances. All coloured things are either red or not, and for apparently borderline cases we are simply ignorant which they are. Although epistemicism is intuitively unattractive, there are arguments in its favour; epistemicists believe their theory is the only way to avoid certain logical problems that afflict other theories of vagueness.⁸ So if epistemicism implied commensurabilism, it would give commensurabilism support.

But epistemicism does not imply commensurabilism. Consider the predicates 'better than soldiering' and 'worse than soldiering'. Epistemicism implies these predicates have sharp borderlines. A tiny improvement in the career of a priest takes us across the borderline from a career that is worse than soldiering to one that is not worse than soldiering, and from one that is not better than soldiering to one that is better than soldiering. In the incommensurate case, epistemicism implies that the borderline

between the worse zone and the incommensurate zone is sharp, and so is the borderline between the incommensurate zone and the better zone. But it does not rule out the existence of an incommensurate zone. Only if these borderlines had to coincide would this be ruled out. Epistemicism does not imply they coincide, so it gives no support to commensurabilism.

Given that, I think we may safely reject commensurabilism because of its sheer implausibility. I shall go on to the second of the two questions mentioned at the end of section 1: given that commensurabilism is false, and options are not necessarily ordered by their goodness, what importance does that have for ethics? I can now express this question more precisely: how ethically significant is the difference between the equality case and the incommensurate case? Sections 5 and 6 consider two arguments about this that have to do with the form of the betterness relation. Sections 7 to 10 consider the importance of the difference for practical decision making.

5. Rough equality

Griffin does not support commensurabilism as I have defined it. For example, he would accept there may be a range of careers as a priest that are intermediate between those that are better than soldiering and those that are worse. So he accepts there are incommensurate cases. However, he seems to think they are all, or at least mostly, cases of 'rough equality'. We might call Griffin a 'rough commensurabilist'. Evidently he thinks the difference between equality and incommensurateness is not very important.

What does 'rough equality' mean exactly? Griffin presents an analysis,⁹ but we do not need to analyse the concept far in order to see that rough commensurabilism is implausible. So long as 'roughly equal' has something like its ordinary meaning, this must be true: two options that are each roughly equal in goodness to a third cannot differ very much from each other in their goodness. They do not have to be roughly equal in goodness to each other; that would follow if rough equality was a transitive relation, but (as Griffin points out) it is not. Nevertheless, they plainly cannot differ very much in their goodness. Griffin is evidently assuming no incommensurate zone can be very wide. But there seems to be no reason why it should not be. We would generally expect the zone to be narrow when we are comparing similar types of thing, and wide when we are comparing very different things. Between joining the Northumberland Rifles and the Fifth Lancers, the zone may be narrow, but between the church and the army it may be very wide. A mediocre career in the church may not be definitely worse than a career in the army, and a successful career in the church may not be definitely better than a career in the army. But a successful career in the church is much better than a mediocre one. It follows that not all careers in the incommensurate zone are roughly equally as good as a career in the army.

To make this argument, I have presumed we have some scale of goodness. It need only be a rough scale, just enough to give sense to such expressions as 'much better',

'do not differ much in their goodness' and so on. The idea of rough equality itself presumes a rough scale, so I have taken one for granted in arguing against it.

Commensurabilism is very implausible because it presumes there is no intermediate zone at all, but at most one intermediate option. Rough commensurabilism is less implausible, but implausible nonetheless because it presumes the intermediate zone is narrow. If we are to believe it, at least we need from Griffin an argument why this zone should be narrow.

6. *Vagueness*

Griffin does make one point that I think is intended as such an argument. He claims that incommensurability is vagueness.¹⁰ When two options are incommensurate, we cannot say either is better than the other. This leaves open several possibilities. Griffin's view is that, although we cannot assert 'This one is better', we also cannot deny it, and although we cannot assert 'That one is better', we cannot deny that either.

Take the career example again. Griffin believes the predicates 'better than soldiering' and 'worse than soldiering' are vague. Unlike the epistemicists, he believes this means there is no sharp borderline between priest-careers that are better than soldiering and those that are not, nor between priest-careers that are worse than soldiering and those that are not. More than that, he thinks these borderlines overlap and extend across the whole incommensurate zone. The consequence is that, although we cannot assert of a career within this zone that it is better than soldiering, we cannot deny it either. Also, although we cannot assert it is worse than soldiering, we cannot deny that either. The incommensurate zone is a zone of vagueness.

Joseph Raz disagrees. He thinks that, of a career in the indeterminate zone, we can deny it is better than soldiering and deny it is worse than soldiering too.¹¹ The issue between these views is a fiddly question of logic. As it happens, I think Griffin is right and Raz wrong, but since the arguments are fiddly I shall not present them here.¹²

The real difference of opinion between Griffin and Raz is not the fiddly one about logic, but that Raz thinks incommensurability is important, and Griffin thinks it is not. They both seem to take it for granted that, if incommensurability was vagueness, it would not be important. That is why they disagree about whether it is indeed vagueness.

But if incommensurability is vagueness, why should that make it unimportant? Raz says it would merely be an instance of the general indeterminacy of our language.¹³ That is true; at least, vagueness is a common feature of our concepts. But why should it make it unimportant? If 'better than' is vague, I see no reason why that should not be a significant feature of ethics. Vagueness in general is not unimportant in ethics. The vagueness of 'person' is a very important feature of several ethical problems.

I think Griffin believes that if the incommensurate zone is a zone of vagueness, then it cannot be very wide, and that would make it unimportant. It would support his contention that options within the zone are roughly equal to the standard in goodness. But vagueness can be very extensive. The Southern Ocean is a vague expanse of water

whose boundaries with other oceans are indeterminate to the extent of a thousand miles or so. So the claim that incommensurability is vagueness does not support the claim that it is rough equality.

I think the question of vagueness is a red herring when we are interested in the importance of incommensurability. It does not provide an argument why the incommensurate zone must be narrow, and it does not suggest independently that incommensurability must be unimportant.

The real question about importance is whether incommensurability makes an important difference to practical decision making. In the next section I shall turn away from the formal questions of vagueness and rough equality to consider decision making directly. In doing this, I shall assume incommensurate zones exist, and that they need not be small. We have found no reason for doubting either of these assumptions.

7. Practical decision making

What implications does incommensurability have for how we ought to act? What ought we to do when faced with a choice between two options that are incommensurate in their goodness? To smooth the way for answering this question, let us make the assumption that the goodness of the options is the only consideration relevant to the choice. This is to assume *teleology*; teleology is the view that how one ought to act is determined by the goodness of the available options. Nonteleological theories deny this. For instance, some nonteleological theories hold there are rules of right action – deontic rules – that are not determined entirely by the goodness of the options. To get a clear view of the importance of incommensurability, let us set nonteleological theories aside.

Granted teleology, when you are faced with a choice between two options, if one is better than the other, you ought to choose the better one. It would be wrong to choose the other. But if the two options are either incommensurate or equally good, what then? We cannot say that you ought to choose one, nor that you ought to choose the other. Choosing either would not be wrong. It is not determinate which you ought to choose. Reason, as it were, leaves you on your own. You must simply decide without the guidance of reason.

This is mysterious in one way. Rational creatures have to be able to decide without the guidance of reason, and it is mysterious how we can do that. Buridan's ass could not manage it; how can we? It is not simply a matter of choosing which we like, since we may not like either of the options any more than the other. So how do we choose? One might be tempted to use this mystery as an argument against incommensurability. One might argue that reason simply cannot leave us in the lurch like this; there must always be something we ought to do.¹⁴ But if this was a valid argument, it would work just as effectively against equality of goodness as against incommensurability. Since equality of goodness is surely possible, the argument must fail. It leaves us with the mystery of how we choose. But I have no need to solve this mystery, because it applies

equally to the incommensurate case and the equality case; it does not distinguish them.

If you are faced with a choice between two options, and they are equally good, it does not matter which you choose. If they are incommensurate, it does matter. People often identify this as the practical importance of the distinction: one type of choice matters and the other does not. But actually this remark is a mere tautology. To say a choice between two options does not matter – literally, and not simply that it does not matter much – is simply to say the options are equally good, and to say it does matter is simply to say they are not equally good.

I conclude that when we think about an isolated choice between two options, there is nothing in practical decision making that is different between the incommensurate case and the equality case. To find a difference we need to look further. In Section 3, I distinguished the cases by means of a structural feature of betterness that does not show up in an isolated choice like this. Using the device of the standard configuration, I made the distinction depend on whether there is an intermediate zone or a single intermediate option, between options that are better than the standard and those that are worse, or just one. So this structural feature is where we should look for a practical difference between the cases. The difference must depend on there being an intermediate zone rather than a single intermediate option.

8. *A puzzle*

Here indeed there is a significant difference. Decision making in the incommensurate case, but not the equality case, can lead to a puzzle. Suppose two careers are open to you: a career in the army and a good career as a priest. Suppose they are incommensurate in their goodness. Then choosing either would not be wrong. You have to choose without the guidance of reason, and suppose you choose the army: you commit yourself to the army career, and give up the chance of a good career in the church. In doing so you are doing nothing wrong. But then suppose another opportunity comes up to join the church, this time in much worse conditions. You now face a choice between the army or a much less good career as a priest. Suppose these two, also, are incommensurate. Choosing either would not be wrong. You have to choose without the guidance of reason. Suppose this time you choose the church. Once again you do nothing wrong. But though you have not acted wrongly in either of your choices, the effect of the two together is that you end up with a much worse career in the church than you could have had. Surely rationality should be able to protect you from this sort of bad result; surely there is something irrational in what you have done. Yet apparently neither of your decisions was irrational. This is puzzling.¹⁵

No such bad result could have emerged had you only been choosing amongst options that are equally good. Suppose you have a choice between *A* and *B*, and the two are equally good. Suppose you choose *B*. Next suppose you have a choice between *B* and *C*, which are also equally good, and you choose *C*. Neither of your choices is contrary to reason. But in this case you are no worse off with *C* than you would have been with *A*. *A* and *C* must be equally good because the relation 'equally as good as'

is transitive. On the other hand, the relation 'incommensurate with' is not. So we have identified a practical difficulty that arises in the incommensurate case but not in the equality case. This certainly makes the difference significant.

The puzzle is that rational decision making leaves you worse off than you could have been. But what is puzzling about that? The same thing can happen in ways that are not puzzling at all. For example, you might rationally make a decision that could be expected to lead to a good result, but by sheer bad luck it might lead to a bad result. The example I described can also be interpreted in a way that is not puzzling. You might just change your mind between one decision and the next. Having chosen the army, you might decide that was a mistake, and make the best of a bad job by taking up the only church career that is still available. Then you would not have acted irrationally. Nor is it puzzling that you end up worse off than you could have been.

The puzzle only arises if, when you make the second choice of a poor church career rather than the army, you do not at the same time repudiate your first choice of the army rather than a good church career. If at one single time you are willing to endorse both decisions, then you are certainly in a puzzling condition.

9. Bayesianism

Compare a similar puzzle that arises in a different context. Suppose you are at a horse race, knowing absolutely nothing about the horses. You have no information whatsoever about which is likely to win; indeed you do not even know the number of entries in the race. You find the bookies offering even odds on Gorgeous to win. Having no reason not to, you place a £10 bet on Gorgeous. In that, you are doing nothing wrong. Next you find even odds are also offered on Screaming Angel to win, and having no reason not to do this either, you place a £10 bet on it. Again, you are doing nothing wrong. Finally, you find even odds offered on Intravenous to win, and you place a £10 bet on Intravenous too, once more doing nothing wrong. By these three bets, you have guaranteed yourself a loss of at least £10. You have accepted what is known as a 'Dutch book' against yourself. This is puzzling, since you have done nothing wrong. (Let us abstract from the fun of the game and assume your only object is to make money.)

Once again, there are ways this can happen that are not at all puzzling. For instance, you might change your mind about the horses' chances as the afternoon progresses. Having bet on one horse, you might decide another is more likely to win and bet on that one, resigning yourself to losing on the first, and hoping to recoup some of your losses. But you are in a puzzling condition if you are willing to endorse your bets on the three horses at the same time.

Here is a Bayesian solution to the puzzle. Your bets are in fact irrational. True, none is individually irrational, because you have no reason not to bet on each of the horses individually. But each bet you place on a horse implies you have a particular degree of belief in that horse's winning. When you bet on Gorgeous at even odds, that implies you believe to some degree more than a half that it will win. You have no reason either

to have this belief or not to have it, so it is rational for you to make the bet or not to make it. But you should only make it if in fact you do have this degree of belief in the horse's winning. Similarly, you should only bet on Screaming Angel if you believe to a degree more than a half that Screaming Angel will win, and you should only bet on Intravenous if you believe to a degree more than a half that Intravenous will win. Now, it is irrational to believe to a degree more than a half that Gorgeous will win, and to a degree more than a half that Screaming Angel will win, and to a degree more than a half that Intravenous will win. So your three bets together are irrational.

Let me put this Bayesian argument in better order. Bayesians think, first, that reason imposes particular constraints on the combinations of bets you should make. These are constraints of consistency, and they are laid down in the 'axioms' of Bayesian theory. The grounds for the axioms are that, if you do not stick to them, you will be vulnerable to traps such as Dutch books.

Bayesians next claim that bets imply particular degrees of belief. For instance, betting on a horse at particular odds implies a particular degree of belief in its winning. They then demonstrate mathematically that, if your betting is consistent according to the axioms, it implies a consistent pattern of beliefs. Otherwise it implies an inconsistent pattern. 'Consistent' here means conforming to laws of probability. For instance, your degrees of belief in various horses winning a race must not add up to more than one. This consistency of belief gives further support to the consistency requirements of the axioms: bets that are consistent according to the axioms can be interpreted in terms of consistent beliefs.

The upshot is that, if you are to be rational, you must act according to some probability beliefs or other. This is so whether or not there are rational grounds for some beliefs rather than others. If there are none – for instance if you know nothing about the horses – you should still have beliefs. As it is often put, you should have subjective probabilities and act on them. This is so even if there are no objective probabilities, so reason does not determine which subjective probabilities you should have. Reason still determines that you should have some.

10. Reasons and intentions

We might take a similar Bayesian line about the career example. Neither of your decisions – to join the army rather than to take up a good career as a priest, and to take up the less good career as a priest rather than join the army – is irrational on its own, but they are irrational in combination. Your decisions should be consistent with each other in particular ways, because otherwise you may find yourself in traps such as the one that actually caught you. Just as Bayesians interpret betting decisions as implying beliefs, we can interpret your decisions as implying values. Your first choice in the example implies that you value a career in the army at least as highly as a good career in the church. On the other hand, your second choice implies that you value a less good career in the church at least as highly as a career in the army. So these two decisions imply inconsistent values. Consistent decisions will imply a consistent value

system. In this case, the objective facts of goodness do not determine what values you ought to adopt. They do not determine whether you should value an army career higher than a good church career, or a less good church career higher than an army career. So reason does not determine what your subjective values should be. But it does determine that you should have some and act according to them.

This Bayesian line is commensurabilist in a sense. Values may be objectively incommensurable, but subjectively they should not be. Objective values may leave room for subjective choices, but to be act consistently you need to settle on subjective values for yourself, and these must have no incommensurability. Moreover, this is a sort of commensurabilism that many people might find credible even if they agree with my claim in Section 4 that commensurabilism is implausible. They might think objective commensurabilism is implausible, but the need for consistency in decisions gives grounds for subjective commensurabilism. Whatever the objective facts, we need to create fully commensurable value-systems for ourselves. This conclusion would have made a pleasant ending for a paper in honour of James Griffin.

However, I am sorry to say I shall have to end with doubts about it. There is a good objection to the Bayesian line. When you make your decisions in the career example, there is no reason to think you are reflecting your values. You have to decide between one career and another. If there are no grounds for your decision, you have to make it all the same, without grounds. Consequently, the fact that you make it does not show you have grounds.¹⁶ In particular, your decision need not imply a subjective system of values that gives it grounds.

The same objection may be raised against the Bayesian line on probability. You have to choose whether or not to bet on Gorgeous. Since you have no information relevant to its chances of winning, you have no grounds on which to base this decision. You have to decide, so you decide without grounds. The fact that you decide to bet does not show you have grounds for doing so. In particular, it does not show you believe to a degree greater than a half that Gorgeous will win.

Bayesians might not be troubled by this objection. To avoid a Dutch book, you must certainly act *as if* you have a consistent pattern of beliefs. Whether you actually have one or not might not particularly concern a Bayesian. If the Bayesian favours a behaviourist or functionalist notion of belief, she will think you do indeed have consistent beliefs, because acting consistently is a criterion for having consistent beliefs. But Bayesians are more concerned with the rationality of your behaviour than with the state of your mind, and they may be perfectly content so long as you act as if you have consistent beliefs.

But I am not satisfied by this casual response. Having bet on Gorgeous and Screaming Angel, Bayesians think you ought not to bet on Intravenous, unless you have repudiated your previous bets. So they think you have some reason not to bet on Intravenous. But what could that reason be? It can only be that you believe, to a degree more than a half, that Intravenous will not win. Your belief must be substantive enough to generate a reason for you. We may have a notion of belief that entitles us to attribute to you particular beliefs on the basis of your previous bets. But could these beliefs be reason-giving for you? Could your previous bets give you a reason not to

bet on Intravenous? I doubt it. You have no information about the horses that could justify you in having particular beliefs and acting on them. Your placing a bet gives you no information either. So how could placing a bet make it the case that you now have a belief, and that this belief gives you reason to act in a particular way? I doubt it could.

Similarly, I doubt that your decision to join the army, in preference to a good church career, necessarily makes it the case that you value the army career more highly than the church career. I recognize that decision making can sometimes stimulate you to form values, and it may even be a way of forming your values, but I doubt it is necessarily so. You have to act, whether or not you form values that determine how you should act. Consequently, I do not see how your first decision gives you a reason not to make your second one (not to choose a poor church career rather than the army). True, it is common nowadays to think that decisions are themselves reason-giving.¹⁷ For example, if you have decided on some end, that gives you a reason to take an appropriate means to this end. In a similar way, your decision to join the army rather than take a good church career could give you a reason not to take a worse church career rather than join the army. But I think this common opinion is incorrect. If you have no reason to pursue some end, even if you have decided to do so, you have no reason to take some means to this end. I do not believe decisions create reasons.

So I doubt the Bayesian idea. But since it was a solution to a puzzle, I need an alternative. The one I have to offer is only a suggestion, which will have to be worked out properly another time. Take the means-end connection again. Deciding on an end gives you no reason to take a means to it, but in a different way it does commit you to taking a means. Deciding on an end is itself deciding to take some means to it. Deciding to go to the cinema is itself deciding to take some means of getting there. To put it another way: intending an end is intending some means to it. So the movement from end to means is not mediated through reasons but through intention. The difference is important. You can intend a means to an end without its being the case that you ought to take a means to this end, and even without your believing you ought to. Moreover, there is absolutely nothing wrong with changing your mind and deciding not to (unless you have already invested some effort). On the other hand, if you had a reason to take the means, there would be something wrong with changing your mind.

Similarly, if you decide to join the army rather than take a good job as a priest, that in itself is plausibly a decision not to take a worse job as a priest rather than join the army. You can change your mind, but not without repudiating your earlier decision. So I suggest the connection between the two decisions is mediated by intention rather than reason. One intention implies another. The structure of intentions – what intention is implied by what – and the nature of this implication still need to be worked out.

The Bayesian solution implies that subjective values must be commensurable, because it assumes only commensurable values can resolve the puzzle I mentioned. My alternative does not require even subjective values to be commensurable. The puzzle can be resolved by the structure of intentions, rather than by the structure of values. I suggest incommensurability of values is indeed an important phenomenon in ethics, which cannot be dissolved away by the Bayesian method.

Notes

1. See his 'Are there incommensurable values', his *Well-Being*, Chapter V, and his 'Incommensurability: what's the problem?'. The latter is Griffin's most recent piece on the subject, and my main source for the views I attribute to him.
2. Actually, he only doubts the incommensurability of what he calls 'prudential values'. He thinks moral values may well be incommensurable. (See 'Incommensurability: what's the problem?'.) But Griffin's 'moral values' are more strictly moral rules such as: 'Do not deliberately kill the innocent'. I take value to be synonymous with goodness, and since these rules are not about goodness, I do not count them as values at all.
3. *Well-Being*, p. 80.
4. See my 'Is incommensurability vagueness?' for more details.
5. *The Morality of Freedom*, p. 326.
6. One way it can arise is if one value lexically dominates another. This means that any increase in one value, however small, is better than any increase in the other, however big. Suppose water and whisky are both good, but any increase in the amount of whisky is better than any increase in the amount of water. Suppose the standard is one litre of whisky, and the chain is composed of various amount of whisky together with one litre of water. Then any option in the chain with less than a litre of whisky is worse than the standard. Any option in the chain with a litre or more of whisky is better than the standard. The chain contains no options in between.
7. Raz (*The Morality of Freedom*, p. 326) takes the existence of an intermediate zone as a central feature of incommensurability, though not its definition. (I did the same in my 'Choice and value in economics'.) My account of incommensurability agrees with Raz's in many ways.
8. See Timothy Williamson, *Vagueness*.
9. *Well-Being*, pp. 96–8.
10. *Well-Being*, p. 96.
11. *The Morality of Freedom*, p. 324.
12. They are set out in my 'Is incommensurability vagueness?'.
13. *The Morality of Freedom*, p. 324.
14. Raz develops this argument and responds to it in 'Incommensurability and agency'.
15. The same point is made by Ruth Chang in 'Incommensurability, incomparability and value'.
16. This point is made in my 'Choice and value in economics', and by Raz in *The Morality of Freedom*, pp. 338–40.
17. For instance, see Edward McClennen, *Rationality and Dynamic Choice*.

References

- Broome, John, 'Choice and value in economics', *Oxford Economic Papers*, 30 (1978), pp. 313–33.
- Broome, John, 'Is incommensurability vagueness?', in Ruth Chang (ed.), *Incommensurability, Incomparability and Practical Reason*, Harvard University Press, 1997.
- Chang, Ruth, 'Incommensurability, incomparability and value: a framework', in Ruth Chang (ed.), *Incommensurability, Incomparability and Practical Reason*, Harvard University Press, 1997.
- Griffin, James, 'Are there incommensurable values?', *Philosophy and Public Affairs*, 7 (1977), pp. 39–59.
- Griffin, James, 'Incommensurability: what's the problem?', in Ruth Chang (ed.), *Incommensurability, Incomparability and Practical Reason*, Harvard University Press, 1997.
- Griffin, James, *Well-Being: Its Meaning, Measurement and Moral Importance*, Oxford University Press, 1986.
- McClellenn, Edward F., *Rationality and Dynamic Choice*, Cambridge University Press, 1990.
- Raz, Joseph, 'Incommensurability and agency', in Ruth Chang (ed.), *Incommensurability, Incomparability and Practical Reason*, Harvard University Press, 1997.
- Raz, Joseph, *The Morality of Freedom*, Oxford University Press, 1986.
- Williamson, Timothy, *Vagueness*, Routledge, 1994.