Philosophy and Phenomenological Research
Vol. LXXV No. 1, July 2007
© 2007 International Phenomenological Society

Reply to Qizilbash*

JOHN BROOME
University of Oxford

In *Weighing Lives* I devoted some chapters to 'the intuition of neutrality.' This is the intuition that a person's existence is in itself neither a good thing nor a bad thing. I thought it deserved space because it is deeply embedded in much of the way we think.

A great many of the things we do will alter the future population of the world. Yet when we try to set a value on these things, we hardly ever take into account the value or disvalue of this alteration of the population. One striking example is provided by the UK's National Institute for Clinical Excellence. NICE's job is to evaluate medical treatments, and it chose at one time to evaluate treatment for infertility. In doing so, it took into account the benefit to parents of having a baby, but it refused to make any allowance for the goodness or badness of the baby's existence in the world, or of the existence of its prospective descendants. Why not? Since fertility treatment particularly aims to change the world's population, why did NICE ignore the direct value or disvalue of doing so? It must be because it is in the grip of the neutrality intuition.

I felt I should give this intuition as much credit as I could. As I have described it, it is just an inchoate idea, so I looked for a way of formulating it accurately and defensibly. I wanted to discover whatever truth is contained in it. In the end I failed, and I concluded that the intuition is mistaken.

In formulating it, we have to recognize it is subject to a limit. We do not always think intuitively that a person's existence has neutral value; we mostly think it is a bad thing for a person to exist if her life goes very badly—if, say, she has a short life full of suffering. So our intuition allows the value of a person's existence to depend on how well off the person is, at least to that extent. In formulating the intuition, we must therefore take the level of the person's wellbeing into account.

* I wrote this comment while supported by a Major Research Fellowship from the Leverhulme Trust. I thank the Trust for its generosity.

The intuition is that, for at least a range of possible levels of wellbeing, the existence of a person at that level is neither good nor bad. I call this 'the neutral range'.

Most people think the neutral range is wide. It does not include the level of a short life full of suffering, but it does include a wide range of levels that are better than this. The intuition is that, *as a general rule*, existence is neutral. So the intuition that needs formulating is this: there is a wide range of levels of wellbeing such that the existence of a person at one of these levels is neither a good thing nor a bad thing.

To begin the task of making this idea precise, compare two worlds. In one, a number of people live, each enjoying some level of wellbeing. Call this world *A*. In the other, just the same people live, and their levels of wellbeing are the just the same as in *A*, but there is an extra person—say Sam—too. Call this world *B*. Assume Sam's wellbeing is within the neutral range. In that case, how does the goodness of *B* stands relative to the goodness of *A*? That is what the question of formulating the intuition of neutrality comes down to.

Since Sam's existence is supposed to be neutral, we know *B* is neither better nor worse than *A*. There are then two possibilities. One is that *A* and *B* are equally good; the other is that they are not. So long as there is a neutral range and not just a single neutral level, it is easy to reject the first possibility. I did so in chapter 10 of *Weighing Lives*. The second possibility is that neither *A* nor *B* is better than the other, and nor are *A* and *B* equally good. In that case, I say they are *incommensurate* in value. Let us call this second possibility the 'incommensurateness view'. I explored it in chapter 12. In the end I concluded it does not provide an adequate formulation of the neutrality intuition. I made three objections to it, which I shall mention below. Since I could find no way to formulate the intuition satisfactorily, I finally abandoned it on page 208.

Mozaffar Qizilbash provides an interesting new defence of the neutrality intuition, using the idea of parity. He suggests that, when *B* is in the neutral range, *A* and *B* are not merely incommensurate, but also *on a par*. If that is so, he claims to have a response to each of my three objections to the incommensurateness view.

What does he mean by 'on a par', exactly? He says (p. 134) 'the distinctive mark' of parity is this: 'If *x* and *y* are on a par, then a significant improvement (worsening) of one of the options makes it better (worse) than the other.' To say that *A* and *B* are on a par therefore implies that any world *C* that is significantly better than *B* is not on a par with *A*.

This tells us something about the neutral range if we adopt Qizilbash's suggestion. Suppose *C* is another world just like *B* in that it

contains all the same people as *A* does, at the same levels of wellbeing, and also Sam. But suppose Sam has significantly higher wellbeing in *C* than she does in *B*. Since *C* is significantly better than *B*, and *B* is on a par with *A*, the distinctive mark of parity implies that *C* is not on a par with *A*. According to Qizilbash, then, it is not in the neutral range. The conclusion is that levels of wellbeing within the neutral range are none of them significantly different from each other. The difference in wellbeing between the top and bottom of the neutral range is not significant.

However, as I described the neutrality intuition, the neutral range is wide. That is the idea that is embedded in much of our thinking. Take NICE and fertility treatment again. NICE does not count the existence of a baby that results from fertility treatment as contributing to the value of the treatment. Its attitude is not much affected by the level of wellbeing that the baby will enjoy. No doubt NICE would oppose any treatment that would lead to the existence of a baby whose life would be short and full of suffering. But it would not attach value to the existence of a baby who would live comfortably to the age of 65, and nor would it attach value to the existence of a baby who would live to 100 and enjoy a varied, exciting and successful life. So on NICE's application of the neutrality intuition, those two levels of wellbeing both fall within the neutral range. Yet one is much better than the other. It is very significantly better.

This means that Qizilbash and I are talking past each other. I would not be much concerned with a neutral range that was insignificantly wide. The problem that concerns me is that our intuition suggests there is a wide neutral range. Qizilbash's response (pp. 141-2) to my first objection to the incommensurateness view depends on taking neutrality as parity. That does not meet the problem that concerns me, since it implies the neutral range is insignificantly wide.

I shall come to my second objection soon. My third was that when neutrality is formulated as incommensurateness, it can be 'greedy' as I put it, and greedy neutrality is not intuitively neutrality at all. I used the example of global warming (*Weighing Lives*, pp. 203-7). Global warming will kill perhaps a hundred million people; that is a very bad effect it will have. It will also alter the future population of the world to an unpredictable degree; according to the neutrality intuition, that is a neutral effect. Now consider those two effects of global warming taken together. One is bad; the other neutral. Intuitively, the net effect should be bad. Yet on the incommensurateness view it may turn out neutral. I showed on p. 170 how incommensurateness allows a neutral effect to swallow up a bad effect and neutralize it. Even though the change in the world's population is supposed to be neutral in value, it

can cancel out a bad thing such as all that killing, and make the combined effect neutral. Global warming might not be bad after all.

That is not how neutrality should intuitively behave. When we make judgements of value, we use our neutrality intuition to let us ignore the effects of our actions on the world's population. That is what NICE does and it is what we do when we conclude that global warming is a bad thing because it will kill so many people. But if neutrality is incommensurateness, we must not ignore those neutral effects. Although neutral, they can affect the value of what we do, by swallowing up the goodness or badness of other effects.

In response, Qizilbash says (p.144) 'If the addition of a person (or people) is on a par with the status quo, it cannot justify any *significant* sacrifice of other values', but he thinks it may justify an insignificant sacrifice. He means it would not be contrary to intuition to have a neutral effect swallow up insignificant good or bad effects. He may be right. But in the example of global warming, the bad effect that is swallowed up is very significant. It is the early deaths of many people. Each one might lose, say, half her life. A sort of neutrality that can swallow up such badness is not intuitively neutral. So the idea of parity does not overcome my third objection to the incommensurateness view.

My second objection is more technical. I claimed that the dyadic predicate 'is better than' is vague. I then argued that its vagueness is incompatible with the incommensurateness view. My argument was found on something I call 'the collapsing principle'. Qizilbash responds firstly by asking why I think 'is better than' is vague, and secondly by arguing that I should not accept the collapsing principle because it can be turned against my own views.

Why do I think 'is better than' is vague? Only because I assume virtually every predicate outside mathematics and physics is vague, and value predicates are vaguer than most. I thought it obvious. Still, Qizilbash asks for an argument. I cannot give a theoretical argument, because vagueness is an intuitive phenomenon. Theories of vagueness take the existence of vagueness for granted, and aim to account for it. Indeed, they often deny some of its intuitive features, because those features tend to lead to paradoxes. So to argue for the vagueness of 'is better than', all I can do is prompt your intuition.

Suppose you have a moderately good chardonnay. Suppose you also have an extensive range of sauvignon blancs, running from very bad to very good. Imagine the range has been so finely graduated by skilful blending that you simply cannot tell the difference between one wine and the next in the range. The best in the range are better than your chardonnay. The worst are not better than your chardonnay. Is there is a sharp boundary between those that are better and those that are not?

Intuitively not. So intuitively, 'better than' is vague. *Weighing Lives* dealt with the quality of lives rather than the quality of wines, but vagueness in that context seems even more intuitively obvious.

Finally, the collapsing principle. Qizilbash does not directly attack this principle, so I shall not directly defend it. Instead, he argues it can be used against my own theory; I shall respond to that argument. He illustrates my theory in his figure 5. Go back to my worlds A, B and C. (A is Qizilbash's $\mathbf{w}^{e}$; B and C are his $(\mathbf{w}^{e}, \mu)$ for different values of μ; μ is the extra person Sam's level of wellbeing.) Among worlds like B and C that contain one more person than A does, I believe there is a single vague borderline between worlds that are better than A and worlds that are worse than A. It is also the borderline between worlds that are better than A and those that are not better than A, and between worlds that are not worse than A and those that are worse than A.

Qizilbash points out (p. 149) 'The logic of asymmetry [the collapsing principle] would imply that the point on the edge of the vague zone is not in the vague zone but in the better zone'. He is perfectly right. The collapsing principle implies that the vague zone does not contain its boundary points; it is *open* in the mathematical sense. Qizilbash seems to think it follows that the entire zone collapses to nothing, but that does not follow.

I showed (*Weighing Lives*, pp. 173-5) that the collapsing principle makes the incommensurateness view incompatible with vagueness. Given the collapsing principle, the incommensurateness view is incompatible with the existence of a vague borderline between worlds that are better than A and worlds that are not better than A. Since this borderline is undoubtedly vague, if the collapsing principle is true, the incommensurateness view is false.

In response, Qizilbash points out (p. 150) that the collapsing principle rules out *second-order* vagueness in my theory: the boundary of the vague zone cannot itself be vague according to the collapsing principle.

He is right, but it does not bother me much. Undoubtedly, the borderline between worlds that are better than A and worlds that are not better than A is vague. That means 'better than' has first-order vagueness. But I am not convinced it has second-order vagueness. Our intuition in favour of second-order vagueness arises from the general thought that, when we are dealing with a vague predicate, there cannot be any sharp borderline anywhere. So there cannot be a sharp borderline at the edge of the zone of first-order vagueness. But there has to be a sharp borderline somewhere. At the far right of Qizilbash's figure 5 there are worlds that are better than A (that is, $\mathbf{w}^{e}$), whose betterness is not infected by vagueness of any order. In the middle, there are

worlds whose betterness, relative to A, is infected by vagueness of some order. There must be a sharp borderline between those that are infected and those that are not. Since there has to be a sharp borderline, I do not see why there should not be one at the edge of the zone of first-order vagueness. If so, there is no second-order vagueness.