

Empirical and Armchair Ethics

GREG BOGNAR

La Trobe University

In a recent paper, Michael Otsuka and Alex Voorhoeve present a novel argument against prioritarianism. The argument takes its starting point from empirical surveys on people's preferences in health care resource allocation problems. In this article, I first question whether the empirical findings support their argument, and then I make some general points about the use of 'empirical ethics' in ethical theory.

INTRODUCTION

In a recent paper, Michael Otsuka and Alex Voorhoeve present a novel argument against prioritarianism.¹ They argue that prioritarianism is inconsistent with a shift in moral judgement between certain intrapersonal and interpersonal resource allocation problems. The shift is supported by empirical surveys on people's preferences in health care resource allocation choices. In addition, the shift is justified by the separateness of persons: the idea that since people lead separate lives, some forms of balancing harms and benefits that are permissible within one person's life are impermissible between different persons.

Prioritarianism is the view that benefiting a person matters more the worse off that person is.² If two people can derive the same net benefit from your aid, but one of them is worse off than the other, then the benefit you can bestow on the worse off person has greater moral importance – and the worse off she is, the greater the moral importance. The benefit to the worse off person matters more, however, not because of her bad position relative to others, but because of her bad situation in absolute terms. The benefit would have the same moral importance even if there were no others. Prioritarianism is not a comparative view.

Egalitarian views, in contrast, are comparative: they are concerned with how people fare relative to others. While some forms of egalitarianism are widely seen as problematic, Otsuka and Voorhoeve also present a version of the view that, they claim, is consistent with the results of empirical studies, can account for the shift in moral judgement and respects the separateness of persons.

¹ Michael Otsuka and Alex Voorhoeve, 'Why It Matters that Some are Worse Off than Others: An Argument against the Priority View', *Philosophy and Public Affairs* 37 (2009), pp. 171–99. All otherwise unattributed page references are to this article.

² The canonical statement of the view is Derek Parfit, 'Equality or Priority?', *Bioethics*, ed. John Harris (Oxford, 2001), pp. 347–86.

In this article, my focus will be on the first part of the Otsuka–Voorhoeve argument that takes its building blocks from surveys on people’s preferences in health care resource allocation problems. I question whether these results support their argument, pointing out some of the pitfalls of using empirical results to draw conclusions in ethical theory. I will conclude that the empirical surveys do not in fact give more support to egalitarianism than to prioritarianism. But since Otsuka and Voorhoeve also formulate their argument in a way that does not rely on empirical results, my objections are not decisive. Nevertheless, I will address this formulation of the argument only indirectly, by questioning the role that empirical results play in their overall argument. Although I believe my discussion points to some problems in their favoured egalitarian view, I won’t pursue this issue here.

THE SHIFT

Otsuka and Voorhoeve begin their argument by introducing the following examples. First, suppose that a perfectly healthy young adult learns that she will develop one of the following two conditions with equal probabilities:

Slight impairment: a condition that renders it difficult for one to walk more than 2 km.

Very severe impairment: a condition that leaves one bedridden, save for the fact that one will be able to sit in a chair and be moved around in a wheelchair for part of the day if assisted by others.³

There are, however, preventive drugs for both of these conditions. Unfortunately, they have to be taken before either condition develops, and they cannot be taken at the same time. Moreover, each drug is ineffective against the other condition. The person therefore has to make a choice. On the one hand, she can take the drug against the slight impairment. If she subsequently develops this condition, she will be returned to full health, but if she develops the other condition, she will remain very severely impaired. On the other hand, she can take the drug against the very severe impairment. If she then develops the condition of slight impairment, she will remain slightly impaired; but if she develops the condition of very severe impairment, she will be restored to the following, less burdensome health state:

³ Otsuka and Voorhoeve adopt the health state descriptions from Erik Nord, Jose Luis Pinto, Jeff Richardson, Paul Menzel, and Peter Ubel, ‘Incorporating Societal Concerns for Fairness in Numerical Valuations of Health Programmes’, *Health Economics* 8 (1999), pp. 25–39.

		The condition that develops is the	
		slight impairment	very severe impairment
Choose the drug for the slight impairment		full health	very severe impairment
Choose the drug for the very severe impairment		slight impairment	severe impairment

Figure 1. The Intrapersonal and Interpersonal cases

Severe impairment: a condition in which one is no longer bedridden; rather, one is able to sit up on one’s own for the entire day but requires the assistance of others to move about.

I shall call this the *Intrapersonal Case*. It is illustrated in the figure. The rows represent the choice between the drugs, and the columns represent the condition that develops.

Second, consider what I shall call the *Interpersonal Case*. Suppose there is an even number of healthy people. You know that half of these people will develop the condition of the slight impairment, and the other half will develop the condition of the very severe impairment. Moreover, you (and you alone) know exactly which condition each person will develop. So you know which drug to give to each person. Even though you have enough drugs for all, you can provide drugs only to one of the two groups. You have to decide which group to aid.

The figure can stand as an illustration of the *Interpersonal Case* as well. You just have to read it differently. The rows again represent the choice between the drugs, but the columns now represent people, identified by the condition they will develop. Half of the people are assigned to the first column, and half of the people are assigned to the second column. In the *Intrapersonal Case*, the columns represent a probability distribution; in the *Interpersonal Case*, they represent a distribution of people.

Health economists have been studying people’s preferences about cases like these. They found that in the *Intrapersonal Case* respondents who imagine that they are faced with the same choice tend to be indifferent between the improvement from the slight impairment to full health and the improvement from the very severe impairment to the severe impairment. They judge the magnitude of the benefits equal. In other words, they tend to be indifferent between taking the two drugs. The utility gain that is associated with being restored to full health from the condition of slight impairment is equal to the utility gain that

is associated with moving from the condition of very severe impairment to the condition of severe impairment.

In the Interpersonal Case, respondents have different preferences. They strongly prefer treating those who would end up with the very severe impairment if they were not helped. Moreover, they prefer providing the drug for the condition of the very severe impairment even though they continue to regard the utility difference between full health and the slight impairment the same as the utility difference between the very severe and the severe impairments. Even though the utility gains would be equal, people prefer giving priority to those who would end up worse off.⁴

Otsuka and Voorhoeve ask you to imagine that you are a ‘morally motivated stranger’ who can provide one of the drugs to the healthy young adult in the Intrapersonal Case and to one of the groups in the Interpersonal Case. How should you choose?

In the Intrapersonal Case, you know that the person is indifferent between the two treatments, and there is no reason to think that her preferences do not adequately reflect her expected benefit. Otsuka and Voorhoeve argue that therefore it is reasonable for you to be indifferent between providing one or the other drug. In the Interpersonal Case, you know that each person is indifferent between the two treatments, and there is no reason to think that their preferences do not adequately reflect their expected benefit. In this case, however, it is *not* reasonable for you to be indifferent: ‘your only reasonable option is to provide the treatment to those who will develop the very severe impairment’ (p. 174).

Now we are in the position to formulate the argument against prioritarianism. For prioritarians, benefiting a person matters more the worse off that person is. In the Interpersonal Case, those who will develop the very severe impairment would end up worse off than those who will develop the slight impairment. The benefit you can bestow on these people – improving their condition from very severely impaired to severely impaired – has greater moral importance. You ought to provide them with the drug. Empirical results show that people agree. So far, prioritarians would concur. But since prioritarianism is not a comparative view, there should be no difference in the importance of the same benefit whether or not other people are present. Thus, for prioritarians, the same moral weights should apply in the Intrapersonal Case. As Otsuka and Voorhoeve put it, ‘the moral

⁴ See Nord et al., ‘Incorporating Societal Concerns for Fairness in Numerical Valuations of Health Programmes’. For a recent review of the empirical literature, see Koonal K. Shah, ‘Severity of Illness and Priority Setting in Healthcare: A Review of the Literature’, *Health Policy* 93 (2009), pp. 77–84.

importance of improving a person's condition from being very severely impaired to severely impaired in a one-person case does not change when we transform this into a multi-person case in which there are others who are better off (p. 177). But we have just seen that it is reasonable to provide the person in the Intrapersonal Case with the drug that maximizes her expected utility – and, as it happens, both drugs do, so there is reason to be indifferent. There is a *shift in moral judgement* between the Intrapersonal and Interpersonal Cases. Empirical findings seem to confirm this. Prioritarianism, however, cannot account for this shift.⁵

EMPIRICAL ETHICS

How can prioritarians respond to this argument? One obvious strategy is to deny that we can learn anything relevant from people's preferences in empirical surveys. Perhaps the preferences people reveal to researchers are irrational. But studies in health economics try to ensure that the respondents are adequately informed about the health states that they are asked to evaluate, care is taken that they have enough time to reflect carefully on the alternatives, and if the responses are inconsistent, the respondents might be asked to reassess them. Thus, the shift in judgement appears to be a *fact of common-sense morality* – a matter of our ordinary, general, but considered moral judgements. It is a fact that may help us decide whether prioritarianism or some other view is the more adequate theory of distributive justice. Unless you hold that common-sense moral judgements are wholly immaterial to this choice, there is a presumption in favour of the relevance of these findings.

A more promising strategy is to adopt this presumption and focus on the shift in judgement on which the argument against prioritarianism relies. The argument would be undermined if it could be demonstrated that the survey results do not in fact show what Otsuka and Voorhoeve take them to show.

This strategy might begin with asking the following question. How do we know that people regard the utility difference between the slight impairment and full health equivalent to the utility difference between

⁵ Otsuka and Voorhoeve assume that in these examples all people have the same life expectancy, their overall utility is affected only by their health states, those who are in the same health state are at the same level of utility, and they remain in the health state they end up in for the rest of their lives. These strong assumptions raise a worry about the artificiality of considering the situation of the young adult in isolation. Shouldn't egalitarians take into account the fact that in real life the person in the Intrapersonal Case is not isolated and hence she would be worse off than others if she chose to take the drug against the slight impairment and ended up with the very severe impairment?

the very severe impairment and the severe impairment? Health economists have developed different preference elicitation methods for the measurement of health utilities. The best known is the *standard gamble*. In this method, respondents are presented with the following choice: they have to choose between living with a health condition, h , for a certain amount of time, and a treatment that either returns them to full health with probability p for the same amount of time or causes immediate death with probability $(1 - p)$. The value for p is varied until the respondent is indifferent between the two options. The utility of the health state is then set equal to this number – that is, $u(h) = p$.

Repeating this exercise for other health states allows us to construct a scale of health utilities. For instance, respondents may be asked to compare the condition of the slight impairment with a treatment that may return them to full health or cause death. Similar comparisons can be made with respect to the conditions of the severe and the very severe impairments. It may then be discovered that the difference in utilities between the health states associated with the slight impairment and full health equals the difference in utilities between the health states associated with the severe and the very severe impairments.⁶

But do these results allow you to infer that the benefit of being restored to full health from the slight impairment is equal to the benefit of improving one's condition from being very severely impaired to being severely impaired? Do these results allow you to infer that the differences in utilities correspond to differences in benefit? They do so only if you make a crucial assumption: that people are *risk-neutral* towards health states.⁷

⁶ Of course, in practice the best we can hope for are approximations, since there is likely to be considerable variation between individual valuations. The differences might be systematic depending on whether the respondents have experience of the health states that they are asked to evaluate. One well-known way that experience influences individual valuation is due to adaptation. Patients who have adapted to living with a chronic condition tend to judge the health state associated with it *less* bad than those who do not have the condition. There is also some evidence that people evaluate health states by using their own health as a reference point. For those who are currently in poor health or have past experience of illness, more severe conditions seem relatively less bad compared to their reference point. Given their own health as a baseline, however, less severe conditions seem relatively worse. That is, respondents who treat their own experience of poor health as a reference point use a narrower range of values to evaluate health states. One implication of this phenomenon is that not only the levels of health, but also changes in health are evaluated differently depending on a person's health experience. I set these complications aside in this article. For further discussion and empirical studies, see Paul Dolan, 'The Effect of Experience of Illness on Health State Valuations', *Journal of Clinical Epidemiology* 49 (1996), pp. 551–64 and Paul Menzel, Paul Dolan, Jeff Richardson, and Jan Abel Olsen, 'The Role of Adaptation to Disability and Disease in Health State Valuation: A Preliminary Normative Analysis', *Social Science & Medicine* 55 (2002), pp. 2149–58.

⁷ More precisely, the assumption is that people have a constant risk attitude towards health states. (For simplicity, I discuss constant risk-neutrality, which is also the standard

Setting aside the technical details, it is not difficult to see why. Consider the utility of money. Suppose a respondent is asked to state a preference between \$1 for sure and a gamble in which she might receive \$2 with probability p or \$0 with probability $(1 - p)$. Suppose she is indifferent between the sure prospect of receiving \$1 and the gamble when $p = 0.7$. Thus, $u(\$1) = 0.7$. Repeating this exercise over other gambles with similar 'prizes' yields an expected utility function that reflects the respondent's risk attitude towards money. In this example, this person is risk-averse: she values having \$1 for sure more than a gamble which has the same expected value.

Now consider health state utilities again. Suppose a person is asked to state a preference between living with the condition of the slight impairment and a gamble whose prizes are returning to full health with probability p and instant death with probability $(1 - p)$. Suppose further that the respondent is indifferent between the two options when $p = 0.7$. Thus, the utility of the health state associated with the slight impairment is set to 0.7. (For convenience, full health is usually assigned in practice the value of 1, and a health state which is just as bad as death the value of 0.)

What is the difference between the two cases? In the case of utility of money, the prizes of the gambles are *quantities*. Thus, when a person reports her preferences over the gambles, you are able to learn about her attitude towards risk by mapping her utilities to these quantities. In the case of utility of health states, however, the prizes of the gambles are not quantities, but *descriptions*. You cannot learn about the respondent's attitude towards risk, since there is no quantity to which you can map utilities. You need to make an assumption. The assumption that is usually made is that people are risk-neutral towards health states.

There are other preference elicitation methods that do not rely on this assumption. One of them is the *person trade-off* method. In this procedure, respondents are asked to make judgements of equivalence rather than indifference. In one version of the procedure, they are presented with two hypothetical health programmes. The first programme restores to full health a given number of people (for instance, 1,000) who all have the same health condition (for instance, they are very severely impaired). The second programme restores to full health n number of people who are all in a different health state (for instance, they are slightly impaired). The respondent has to determine

assumption.) The need for such an assumption is widely recognized. See, for instance, John Broome, 'Qalys', *Ethics Out of Economics* (Cambridge, 1999), pp. 196–213, and Graham Loomes and Lynda McKenzie, 'The Use of QALYs in Health Care Decision Making', *Social Science & Medicine* 28 (1989), pp. 299–308.

the value of n at which the two programmes are equally desirable. (n will be greater than 1,000, since the first programme cures a more severe health condition.) The badness (or *disutility*) of the health state that the second programme targets is determined by the ratio of the numbers in the two groups – in this case, it is $1,000/n$.

Just like the standard gamble, the person trade-off relies on a crucial assumption. Since the task of the respondents is to indicate the point at which the total benefits of the two hypothetical programmes are equal, it is assumed that their judgements are not affected by distributive considerations. Their task is to make judgements of the badness of health states. Utility measurement would therefore fail if, for instance, the respondents disregarded the differences between the severity of health conditions for the sake of giving an equal opportunity for treatment to all patients.

Let us return to the argument against prioritarianism. What causes the difficulty for this view is the shift in judgement between the Intrapersonal and Interpersonal Cases. The shift is supported by empirical studies. In the Interpersonal Case, the results were obtained by using the person trade-off method on a scale that consists of eight conditions from full health to death, including the slight, severe and very severe impairments. Each step on the scale represents an equal health improvement. Yet in the studies, respondents preferred moving one patient from a worse initial condition up one step to moving several other patients up one step from a better initial condition.⁸

⁸ The full scale consists of the following conditions: (1) full health; (2) slight impairment; (3) moderate impairment; (4) considerable impairment; (5) severe impairment; (6) very severe impairment; (7) complete disability; (8) dead. (For details, see the Appendix to Otsuka and Voorhoeve and Nord *et al.*, 'Incorporating Societal Concerns for Fairness in Numerical Valuations of Health Programmes'.) The reason differences between adjacent steps are considered to represent equal intervals is that even though initially the scale was developed using direct scaling methods, the results were later corroborated by other preference elicitation methods, including the standard gamble. When direct scaling methods are used, respondents are asked directly to evaluate health states at a high level of precision (e.g. on interval or ratio scales). These methods assume that respondents can make these comparisons, but the empirical evidence that they can do so consistently and in a manner that correlates with results obtained by indirect methods (including the standard gamble and the person trade-off) is notoriously inconclusive (see, e.g., Debra G. Froberg and Robert L. Kane, 'Methodology for Measuring Health-State Preferences – II: Scaling Methods', *Journal of Clinical Epidemiology* 42 (1989), pp. 459–71). Hence the fact that the health state utilities on the scale used in these studies are consistent with results obtained by the standard gamble is significant – especially given that many health economists regard the standard gamble as the benchmark preference elicitation method because it is based directly on the axioms of expected utility theory. See also Erik Nord, 'The Trade-Off between Severity of Illness and Treatment Effect in Cost-Value Analysis of Health Care', *Health Policy* 24 (1993), pp. 227–38 and Nord *et al.*, 'Incorporating Societal Concerns for Fairness in Numerical Valuations of Health Programmes', pp. 29–30 for a more detailed explanation of the way

So we know that the judgements in the Intrapersonal Case were corroborated by studies using the standard gamble, and the judgements in the Interpersonal Case are based on studies using the person trade-off method. We also know that each procedure makes a crucial assumption about people's preferences. Most importantly, they make *different* assumptions. And this raises a worry about the shift in moral judgement: rather than a fact of common-sense morality, it is likely to be *an artefact of different preference elicitation methods*.

The import of this argument is not that health state utility measurement is inconsistent. After all, health state utilities are approximations, and all preference elicitation methods make simplifying assumptions one way or another. But Otsuka and Voorhoeve's shift in judgement is obtained by combining the results of two different procedures. What the discrepancy is likely to indicate is not a difference in judgement, but a difference in measurement.

There is, however, an argument that can be made in defence of the claim that the shift in judgement is no mere artefact. For there is an interpretation of the person trade-off method on which it is not a health-state utility measure at all. Recognizing that people's preferences are bound to be affected by distributive considerations, you can make a virtue out of necessity by interpreting the responses obtained in person trade-off studies as expressions of *social value judgements* of alternative resource uses. And indeed, this is how the responses to person trade-off questions are interpreted in the studies I have been discussing. As Erik Nord justifies this interpretation, 'there is no need for utilities for health states in informing resource allocation decisions if society's valuation of different improvements in health can be measured directly'.⁹

On this interpretation, there is an important difference between the values that are elicited by a procedure like the standard gamble on the one hand, and the person trade-off method, on the other. In introducing the Intrapersonal Case, I said that respondents tended

the scale was derived. (I thank Alex Voorhoeve for clarification of the way they interpreted these studies, and Erik Nord for discussion on the development of the scale.)

⁹ Erik Nord, *Cost-Value Analysis in Health Care: Making Sense out of QALYs* (Cambridge, 1999), pp. 115–16. But it must be noted that this is a minority view. For instance, the widely used burden of disease measure, originally developed by the World Bank and the World Health Organization, uses a person trade-off protocol to assign 'disability weights' to different health conditions. These weights are not intended to reflect distributive considerations; they represent the burden that different health conditions impose on a person relative to full health. For details, see Christopher J. L. Murray, 'Rethinking DALYs', *The Global Burden of Disease: A Comprehensive Assessment of Mortality and Disability from Diseases, Injuries, and Risk Factors in 1990 and Projected to 2020*, ed. C. J. L. Murray and A. D. Lopez (Harvard School of Public Health on behalf of the WHO and the World Bank, 1996), pp. 1–98.

to be indifferent between the improvement from the slight condition and the improvement from the very severe condition. Hence, they were indifferent between the drugs. But, strictly speaking, these claims are only *inferred*. For respondents in studies using common preference elicitation methods are asked to evaluate health states – that is, utility *levels* associated with different conditions – rather than *changes* in health states. It is a further assumption that the value of a health improvement is simply the difference between the utility associated with the health state that obtains before the improvement and the utility associated with the health state that obtains after the improvement. There is no direct measurement of the value of the improvement itself.

In the person trade-off method, in contrast, improvements are measured directly. Health state utilities, if they are needed at all, are used only to ensure that the magnitudes of the health improvements are equal between adjacent health states on scales of conditions like the one I have been discussing (see note 8). The assumption is that when people make their judgements, they care about more than just the magnitude of the improvement in health, and these additional factors would be ignored if the value of an improvement was regarded simply as the difference of the utilities of health levels. The hypothesis that health economists test by using person trade-off questions is that all things considered it is not only the magnitude of the health improvement that matters to people in health care resource allocation.

This raises another problem for comparing empirical results that were arrived at by different methodologies. In examples like the Intrapersonal Case, it is implicitly assumed that what matters to respondents are the differences of the magnitudes of health improvements. The preference elicitation procedures used, including the standard gamble, leave no room for respondents to express judgements about what matters to them *all things considered*. In studies using the person trade-off method, in contrast, the preference elicitation procedure allows respondents to express judgements about how different improvements with given magnitudes matter to them.

Let me illustrate this point on the Otsuka–Voorhoeve argument. Their case against prioritarianism can be constructed from the following building blocks:

- (i) On the scale of health conditions that includes the slight, severe and very severe impairments, the utility differences of being moved from the slight impairment to full health and being moved from the very severe impairment to the severe impairment are equal. These improvements represent benefits with the same magnitude, as established by health state utility measurement.

- (ii) However, when people are asked about their judgements about the social value of improving the health of patients in different conditions, as illustrated by the Interpersonal Case, they hold that these benefits *matter differently*. In particular, benefiting people who are in a worse condition matters more.
- (iii) But when only one person is concerned, a morally motivated stranger has no reason to hold that these benefits matter differently. As the Intrapersonal Case illustrates, she has reason to share the person's indifference. And since the only relevant difference between the Intrapersonal and Interpersonal Cases is the presence of others, prioritarianism cannot countenance this shift in moral judgement.

Notice, though, that only (i) and (ii) are supported by empirical studies by Otsuka and Voorhoeve. The judgements in (ii) are based on person trade-off studies. The judgement in (i) is based on studies using standard health state utility measurement methods. The judgement in (iii) is taken as a direct implication of the judgement in (i).

But, as it happens, we do have empirical results that are relevant to (iii). In more recent work, researchers directly asked respondents how they value improvements on the disability scale that was used in the studies discussed above. In one version of the question, respondents were told to think of two equally large groups of people who are all very severely impaired, but surgery can improve their condition. Patients in the first group can be restored to the slight impairment; patients in the second group can be restored to the following health state:

Considerable impairment: a condition in which one can move about with difficulties at home, but needs assistance on stairs and outdoors.

Respondents were asked to regard the improvement from the very severe impairment to the slight impairment *twice as large* as the improvement from the very severe impairment to the considerable impairment.¹⁰ Their task was to imagine first that they belong to the first group and then that they belong to the second group, and judge their strength of desire for surgery as patients in both groups. Both surgeries were assumed to have the same mortality risk. But in spite of the differences in the magnitudes of the benefits, almost two-thirds of the respondents reported 'much the same desire' for both of the surgeries, and fewer respondents reported a 'much stronger desire' than only a 'somewhat stronger desire' for the surgery that would restore them to the slight impairment. That is, a larger health gain

¹⁰ In terms of the ranks on the scale, described in note 8, the first group can be moved from step 6 to step 2, and the second group can be moved from step 6 to step 4.

was not regarded proportionally more valuable. Considering these two patient groups, and imagining that they could belong to either of them, the larger benefit does not seem to *matter more* to many people – or, more generally, the way benefits matter to people is not proportional to their size. As the researchers conclude, ‘our overall impression is that a majority of respondents seem to deviate considerably from being health benefit maximisers’.¹¹

Remember that as a morally motivated stranger, it is reasonable for you to provide to the healthy young adult the treatment that she prefers in the Intrapersonal Case. It is reasonable for you to let her preferences determine whether you give her the drug that prevents the slight impairment or the drug that prevents the very severe impairment. But this could be understood in two different ways. It could mean that you should let her expected utility determine your choice. Or it could mean that you should let her judgement about *how these benefits matter* determine your choice. Prioritarians can point out that it is the latter judgement that morally motivated strangers should take into account. And there is some empirical evidence that suggests that people’s judgements in this respect are not incompatible with prioritarianism.¹²

ARMCHAIR ETHICS

Otsuka and Voorhoeve would insist that prioritarians are not off the hook. For the shift in moral judgement is not just an alleged fact of common-sense morality – it is also justified on moral grounds. There is a moral reason to treat the Intrapersonal and Interpersonal Cases differently. As Otsuka and Voorhoeve put it, ‘whether or not the stranger should maximize expected utility in the [Intrapersonal

¹¹ Erik Nord, Anja Undrum Enge and Veronica Gundersen, ‘QALYs: Is the Value of Treatment Proportional to the Size of the Health Gain?’, *Health Economics* 19 (2010), pp. 596–607, at 606. The experiment that is summarized here was only one of four to study the hypothesis that distributive considerations enter into *personal* valuations of health gains as well.

¹² There might be, however, a question about how far the correspondence between people’s judgments and prioritarianism goes. Most of the relevant studies focus on ‘concern for severity’ – that is, the concern for giving priority to those whose health condition is worse. But perhaps the concern for severity is narrower than the prioritarian concern for the worse off, hence the empirical results cannot ultimately be used to support the prioritarian view. A recent study, however, found that similar results are obtained when the better and worse off patient groups are identified in non-health terms as being ‘advantaged’ and ‘disadvantaged’. In fact, the study found that there is a stronger preference for providing treatment to the worse off group when it is described in more general terms. These results may suggest that the concern for severity can indeed be interpreted as a general prioritarian concern for the worse off. (Colin Green, ‘Investigating Public Preferences on “Severity of Health” as a Relevant Condition for Setting Healthcare Priorities’, *Social Science & Medicine* 68 (2009), pp. 2247–55.)

Case], a shift of weighting when we move to the Interpersonal Case can be resisted only on pain of denying the moral significance of the separateness of persons'. The *separateness of persons*, in the sense Otsuka and Voorhoeve use the term, entails that 'some forms of balancing benefits and burdens that are permitted when these accrue to a single person are impermissible in cases where these benefits and burdens accrue to different people'.¹³ That is, benefits and burdens have different moral importance depending on whether they have to be traded off for one person or they have to be traded off between different persons. There are cases when it is permissible to trade off benefits and losses for one person, but it is impermissible to trade off the same benefits and losses between different people – as, according to Otsuka and Voorhoeve, it is illustrated by the different moral judgements in the Intrapersonal and Interpersonal Cases. Prioritarianism, however, is insensitive to this difference. Thus it cannot account for the separateness of persons.

In this article, I won't assess Otsuka and Voorhoeve's claims about the separateness of persons. What I am interested in is the role that the empirical results play in their overall argument against prioritarianism and in defence of egalitarianism. For, you might think, it is a bit odd to use these results as an illustration of the argument, but at the same time claim that they are ultimately irrelevant to their overall case.

To be sure, an illustration is just that – a way to introduce and motivate an ethical argument. If you have strong, independent reasons to accept the argument, then you have good reasons to reject the illustration and argue that the underlying judgements are ethically unacceptable. Even if my interpretation of the empirical results – that they are more compatible with prioritarianism than egalitarianism – is correct, Otsuka and Voorhoeve can still claim that this does not take away from the force of their overall argument. If people's preferences in health care resource allocation problems are incompatible with egalitarianism, so much the worse for those preferences.

But remember that the starting point of the argument is a *presumption in favour of the relevance of empirical ethics*. People's preferences in health care resource allocation problems are taken to reveal a fact of common-sense morality – a fact that may help us decide whether some prioritarian or egalitarian view is the more adequate theory of distributive justice. An illustration of an ethical argument that appeals to common-sense morality should surely provide *some* support for that argument, even if the judgement of common-sense

¹³ Both quotes are from p. 179, with capitals added to case names.

morality should ultimately be rejected when it conflicts with a sound ethical principle or reflects some form of social prejudice.¹⁴

In other words, the presumption should be considered *defeasible*. If some judgement of common-sense morality conflicts with our best ethical theory, then, other things being equal, we should favour the theory. I don't think this idea is especially controversial. I believe it is in fact widely shared. In a conflict between empirical and armchair ethics, the latter is able to defeat the former.

Suppose, then, that my interpretation of the empirical results is correct: the discrepancy in people's preferences between the Intrapersonal Case and the Interpersonal Case does not indicate a shift in moral judgement. It's simply an artefact of different preference elicitation methods. The moral judgements of the respondents in these cases support a prioritarian rather than an egalitarian view. So the question is this: is the justification of the shift in moral judgement by Otsuka and Voorhoeve that appeals to the separateness of persons able to defeat the presumption in favour of the relevance of these results?

Consider again the Interpersonal Case. You have two groups of people: one group will develop the condition of slight impairment, and the other, equally large group of people will develop the condition of very severe impairment. You have to decide whether you help the people in the first group to avoid the slight impairment, or you improve the condition of the people in the second group from being very severely impaired to being severely impaired. Empirical results show that people prefer helping the second group. Prioritarianism entails the same, since the benefit of improving one's condition from the very severe to the severe impairment has greater moral weight than the benefit of restoring the health of a person who would end up with the slight impairment, even if the size of these benefits is equal. Moreover, egalitarianism has the same implication: you should help the group that would end up in the worse situation, since this way you are able to minimize the resulting inequality between the two groups.

Otsuka and Voorhoeve do not question the empirical results. In fact, they explicitly endorse them. They use them as part of the illustration of their case. Evidently, these results are not 'defeated' by their claim about the importance of the separateness of persons. In this case, egalitarians make the same judgement as respondents in empirical studies.

¹⁴ For a take of empirical researchers on this issue, see Paul Menzel, Marthe R. Gold, Erik Nord, Jose-Louis Pinto-Prades, Jeff Richardson, and Peter Ubel, 'Toward a Broader View of Values in Cost-Effectiveness Analysis of Health', *Hastings Center Report* 29 (1999), pp. 7–15. See also Rebecca L. Walker and Andrew W. Siegel, 'Morality and the Limits of Societal Values in Health Care Allocation', *Health Economics* 11 (2002), pp. 265–73, for a more sceptical view.

Consider now the Intrapersonal Case. You have to decide whether to provide a drug that prevents the condition of slight impairment or a drug that prevents the condition of very severe impairment when a person has an equal chance to develop either condition. No one else is affected by this choice. Prioritarians would argue that other things being equal, it is more important to provide the person with the drug that prevents the very severe impairment. Even though she will end up with either the slight or the severe impairment, it is morally more important that she avoids the very severe impairment than the slight impairment. This is morally more important even if the benefit of improving her condition from the very severe impairment to the severe impairment and the benefit of avoiding the slight impairment are the same size.

On my interpretation of the results, respondents in empirical studies do not disagree: they hold that avoiding the very severe impairment matters more than avoiding the slight impairment. This judgement does not support the case of Otsuka and Voorhoeve against prioritarianism. It cannot motivate their argument. Of course, Otsuka and Voorhoeve do not actually rely on these results. They simply *infer* what people's judgements would be in this case from valuations of health state utility levels. But we know from other studies that when people evaluate health state utility *changes*, their judgements tend to be different.

Now Otsuka and Voorhoeve might argue that their independent argument from the separateness of persons entails that we should not accept these judgements of common-sense morality. Their defence of egalitarianism defeats them. Therefore, we have sound ethical reasons to reject the judgements in the Intrapersonal Case.

But it's hard to see what the basis of this rejection would be. Their case for egalitarianism has nothing to do with how we should judge the case. Recall that you are asked to consider the case of this healthy young adult in complete isolation from others. No one else is affected by your choice. But egalitarian views are comparative: they are concerned with how people fare relative to others. Hence egalitarianism has no implication for the Intrapersonal Case. On the assumptions that are made about the case, egalitarianism does not tell you which drug to choose. So it cannot 'defeat' common-sense moral judgements in such one-person cases.

Otsuka and Voorhoeve's justification of the shift in moral judgement appeals to the separateness of persons. The shift concerns the difference between the Intrapersonal and Interpersonal Cases. Prioritarians deny that there is such a shift, and people's preferences in health care resource allocation problems seem to be closer to their view. Given the presumption about its relevance, common-sense

morality helps defend prioritarianism from the Otsuka–Voorhoeve argument.

CONCLUSION

In this article, I have focused on Otsuka and Voorhoeve's use of empirical studies in support of their case against prioritarianism. I have argued that these studies do not show what they take them to show. Their mistake is to rely on an inference that is due to an artefact of different preference elicitation methods. Once studies with more consistent methodologies are considered, it appears that empirical surveys do not in fact give more support to egalitarianism than to prioritarianism.

Since Otsuka and Voorhoeve also provide an argument that does not rely on empirical results, my objections are not decisive. But I also claimed that if empirical surveys are taken to reveal the judgements of common-sense morality, prioritarian views receive more support from them than egalitarian views. Of course, that support should ultimately have no moral weight if there are sufficient reasons to accept egalitarianism on independent theoretical grounds or if further empirical studies yield different results. But surely, until then the support should have some moral weight.

greg.bognar@latrobe.edu.au