Debate

Double jeopardy, the equal value of lives and the veil of ignorance: a rejoinder to Harris

John McKie, Helga Kuhse, Jeff Richardson and Peter Singer   *Monash University, Clayton, Australia*

Abstract

Harris levels two main criticisms against our original defence of QALYs (Quality Adjusted Life Years). First, he rejects the assumption implicit in the QALY approach that not all lives are of equal value. Second, he rejects our appeal to Rawls's veil of ignorance test in support of the QALY method. In the present article we defend QALYs against Harris's criticisms. We argue that some of the conclusions Harris draws from our view that resources should be allocated on the basis of potential improvements in quality of life and quantity of life are erroneous, and that others lack the moral implications Harris claims for them. On the other hand, we defend our claim that a rational egoist, behind a veil of ignorance, could consistently choose to allocate life-saving resources in accordance with the QALY method, despite Harris's claim that a rational egoist would allocate randomly if there is no better than a 50% chance of being the recipient.

In a recent article in this journal we defended the use of QALYs as a basis for allocating health care resources.[1] John Harris criticised our efforts in the same issue.[2] Harris exposed some shortcomings in our defence of QALYs, and we are grateful for his response to our article. Needless to say, we do not agree with everything Harris says. Indeed, we disagree with his two central claims: that all lives are of equal value, and that Rawls's veil of ignorance test does not support the QALY procedure. We hope in this rejoinder to clarify these basic, underlying points on which there remains disagreement between Harris and ourselves, and also to point out some areas of agreement.

It is obvious that underlying several of Harris's criticisms is his contention that all lives are of equal value. One interpretation of this claim is the following: all individuals are equally entitled to life-saving treatment regardless of their quality of life or life expectancy, if their interest in continued life is equally as intense. However, Harris makes it clear that this is not what he means by the claim that all lives are of

equal value. "Singer *et al* wrongly attribute to me the view that the value of life depends upon the strength of the individual's preference for living. I don't believe this is true, nor even that if it were true, the strength of a wish to live could possibly be measured in any way that would make interpersonal comparisons meaningful or consistent".[2] (As Harris further points out, our misinterpretation of him on this point was excusable, in light of some of the things he has said.) Harris holds the view that it is sufficient simply that a person wants to go on living: all individuals are equally entitled to life-saving treatment regardless of their quality of life or life expectancy, and regardless of how *intense* their interest in continued life is, so long as they do have such an interest. After pointing out some consequences of the QALY method, he states: "I find these consequences unattractive and I am therefore drawn to the view that the value of life varies neither with its quality nor with its quantity, but is the value that the individual whose life it is places upon it. If she thinks that she would like her own life to stand equal competition with all other lives, then it should".[2]

But with this interpretation of the claim that all lives are of equal value we cannot agree. As we pointed out in our original article,[1] the logic of Harris's position would appear to require him to say that Otto is as entitled to the only available life-saving organ for transplant as Richard, even if Otto will live only two days in relatively poor health, whereas Richard will live for thirty years in perfect health, so long as Otto wants to live for those two days. Harris acknowledges that this is a problem with his view, but is not prepared to abandon his position, because he sees the supporter of QALYs as being vulnerable to an equivalent problem: "Singer *et al* are committed to valuing lives more, the more un-elapsed lifetime they are likely to contain. This they would (or should) maintain however small this temporal advantage is, so long as it is sufficient for the particular individual to derive some benefit from it, so long as it gives him or her *some* interest in continued life derived from that temporal advantage. Likewise for small advantages in quality of life. According to Harris, "both [sides] seem vulnerable when small differences are highlighted".[2]

Key words

However, we see these situations as quite different. It is true that someone with a slight QALY advantage – someone who stands to benefit only marginally (though still measurably) more in terms of quantity of life or quality of life improvement – will be a higher priority for treatment on the QALY approach (if all else is equal). But we see this as appropriate: it is a sign of the sensitivity of the QALY approach – its capacity to take account of even small benefits, if they are genuine, and there are no other differences between claimants on which to base a decision – ie all else is equal. By contrast, the "small differences" problem confronting Harris's position is not a "small differences" problem at all: it is a "big differences" problem. The two-day time advantage that Otto will gain, let us agree, is small. But this is not the problem. The problem is that the *difference* between what Otto will gain and what Richard will gain (thirty years) is huge (even setting aside their difference in quality of life), and still Harris's position can take no account of it. The "problem" confronting the QALY approach is one of sensitivity (ie it can be seen as a virtue), whereas the problem confronting Harris's view is one of insensitivity – it (sometimes) cannot take account of big differences in quality of life gains and quantity of life gains. It violates, not just in a minor way, but in a big way, the principle that we should allocate resources so as to achieve the most good.

Macro-allocation

A second problem with Harris's equal consideration view is that it depends on knowing whether an individual "would like her own life to stand equal competition with all other lives". This may be of some use when choosing between two individuals, but when it comes to macro-allocation the situation is unclear. If we have to choose between funding one of two large-scale health-care programmes, it is impractical to survey even the direct beneficiaries, let alone the potential beneficiaries, to see how they value their lives. Presumably, on Harris's account we would just have to assume that *everyone* "would like their own life to stand in equal competition with all other lives". But this jeopardises the practicality of Harris's account. In effect, it requires us to fund everything; alternatively, it requires us to fund health programmes at random. We must assume that everyone is equally deserving of resources which will save their lives because in the macro-allocation context there is no way of knowing who "would like their own life to stand equal competition with all other lives". In brief, in the macro-allocation context, Harris's view offers no criterion of resource allocation at all.

Harris further claims that on the QALY account "all individuals would have lives of different value".[2] There is a sense in which this is true. If we *have* to make decisions about who will and who will not receive health care, due to limited resources, then some individuals will receive treatment and some will not. In this sense different lives will have different value. But decision-making based on QALYs is not unique in this respect: it is an unavoidable feature of decision-making in the context of health care shortages. On the other hand, there is a sense in which the claim that "all individuals would have lives of different value" on the QALY account is not true. The QALY approach is egalitarian in the sense that equal weight is given to the QALYs of all those potentially affected by an allocation decision; no one's QALYs count for more than anyone else's. In this sense all individuals have lives of equal value. We think it is quite appropriate to say that all individuals have lives of equal value in the latter sense, and see nothing wrong with saying that all individuals have lives of different value in the former sense – ie in saying that some individuals should get resources and others should not, in the context of resource shortages, depending on their capacity to benefit from them.

Harris also suggests that if we used QALYs to allocate resources there would always be a discrepancy between the estimated value of a person's life and its actual value, which in the case of life and death decisions would always be concealed: "If Singer's life were to be saved rather than Harris's because his quality of life and life expectancy were allegedly greater, no one would ever know whether or not my life was in *'fact'* the more valuable, because I did in the event (and *per imposibile*) live longer with a bigger smile on my face".[2] However, this merely reveals that human beings are not all-knowing. We have to make our decisions – even our life and death decisions – on the basis of the best information we have available, which is always limited. If we *have* to choose between saving Singer's life and saving Harris's life, then we should save the one who we *expect* will benefit most from having his life saved. The fact that we might make the wrong choice, and will never know, merely reflects our human fallibility. But facing the prospect of being wrong should not deter us from making hard decisions as best we can – it should not force us into making random choices.

Murder

With his equal consideration view in mind Harris also argues that on the QALY view the value of life would vary with the quality of an individual's life, "so that those with greater quality would have greater value and the wrong done by murder, for example, would vary with the quality of the victim's life, the greater the quality the greater the wrong done in ending it".[2] Again, in a sense this is true. If Otto is murdered with only two days to live the loss is less, from an impartial point of view, than if

Richard is murdered with thirty years left to live, if all else is equal. But we must be cautious about what implications we draw from this. It does not follow that the law should punish the murderer of Otto less severely than the murderer of Richard. If we want the institution of law to deter murder and other forms of violent assault, if we want to engender feelings of compassion and respect for life in the community, and so on, then it may be best for the law to treat the murderer of Otto no less leniently than the murderer of Richard. Of course, motives must be taken into account. If the murderer of Otto sincerely believed that he was acting in Otto's best interests by sparing him from two days of suffering, then it may be entirely appropriate for the law to treat him more leniently for this reason. On the other hand, if the murderer of Otto was merely taking advantage of Otto's defencelessness, then the law may treat this as a particularly heinous crime. In brief, we should not draw conclusions about the appropriateness or justifiability of punishment merely on the basis of the victim's quality of life or life expectancy.

Main objection

The other main objection Harris has to our defence of the QALY method concerns our use of the veil of ignorance – the device invented by Rawls to facilitate impartial choices. How would rational egoists behind a veil of ignorance choose if they were faced with the situation in which they each need life-saving treatment, but there is enough life-saving treatment for only one? We argued that to maximise the satisfaction of their own interests, rational egoists would have to choose a system that gives preference to saving life when it is most in the interests of the person whose life is saved.[1] In other words, if QALYs were an accurate way of measuring when life is most in one's interests, then rational egoists would choose to allocate in accordance with QALYs.

Harris objects to this defence of the QALY method on several grounds. First, he identifies what he takes to be a fallacy in our reasoning: "a rational egoist would surely only give preference to saving the life of the person with the highest interest in continued existence, when he is also *most likely to be that person*. If he has no better than a 50/50 chance of being that person, it cannot be in his interest to prioritise the life of such a person because he cannot know that he won't in fact be worsening his own chances; the same goes of course for non life-saving gains".[2] But Harris has made an elementary error here. Whether treatment is offered according to QALY principles or randomly, since it is only offered to one patient, there is a 50% chance of being the patient who receives no treatment and dies. Since this is common to both methods of allocation, the rational egoist could disregard it. But

similarly, on both methods there is a 50% chance of being the patient whose life is saved. Then what the rational egoist is offered if treatment is allocated according to QALY principles is a 50% chance of getting a big utility gain (having your life saved when you will have a long life of good quality). On the other hand what the rational egoist is offered if treatment is allocated randomly is a 25% chance of the same big utility gain and a 25% chance of a small utility gain (life-saving treatment when you will die within a few days anyway). All we have to do is substitute numbers for the various outcomes and it is entirely obvious what the rational egoist could consistently choose. For example, if we rate living a long life of good quality as having 10 units of value, living for a few days as having 1 unit of value, and dying immediately as having 0 units of value, then the expected utility of the rational egoist when treatment is allocated according to QALY principles, but before the veil of ignorance has been lifted, is 50% of ten, that is, five. On the other hand the expected utility of the rational egoist if treatment is allocated randomly is 25% of $10 + 25\%$ of 1, that is, 2·75. While the rational egoist may not always precisely follow the arithmetical predictions of expected utility, the point here is that behind a veil of ignorance he or she could consistently distinguish between the outcomes of different treatments on the basis of length of life and not simply in terms of the probability of survival.

Of course, this assumes that it makes sense to talk about some people's interest in continued life being low while that of others is high, and this is just what Harris rejects.[2] He concedes that it makes sense for an individual, when comparing alternative futures, to prefer one to another – it makes sense to say that an individual has a greater interest in continued life under one set of conditions compared with another. But he rejects the idea that it makes sense when comparing the alternative futures of different people. We agree with Harris that it is not possible directly to compare one person's interest in living with another. But this does not mean that such comparisons are meaningless, nor that they are avoidable in practice. After all, it is not possible directly to compare one person's pain with another, but this does not mean that interpersonal comparisons of pain severity are meaningless. The circumstance may arise where we have to decide which of two individuals is in the most pain – for example, when there has been an accident and we have only enough pain-killer for one. In this situation we do the best we can to make a judgment based on the criteria of pain severity with which we are all familiar: the injuries of the two individuals involved, their responses to those injuries, how bad they say their pain is, and so on. If one person's pain rated high on these indicators and the other's low, and all other things were equal, we would *not* flip a coin. Even if we might be mistaken, it would be most responsible to rely on these

admittedly imperfect indicators, especially if the gap were large, and give the pain-killer to the person we believed was in most pain (if the severity of the pain does not mean that the pain-killer would be ineffective, etc).

It is also important to distinguish clearly between the interests of the rational egoist who must choose from behind a veil of ignorance, and the interests of the individual who is saved once the veil is lifted. A rational egoist behind a veil of ignorance will have a greater interest in living a long life of good quality than in living a short life of poor quality. But once the veil is lifted whichever individual is saved may express the same interest in life as other individuals, and through all his or her actions indicate the same interest. However, resource allocation decisions must be made before individuals know which treatment they may need and the attraction of the veil of ignorance is that it facilitates impartial choices; it requires us to compare and rank our interest in continuing to live without a prior commitment to some particular alternative. Under these circumstances a rational egoist may well state or reveal that he or she has a greater interest in living in one health state rather than another. It is for this reason that we can state that an "interest in continued life" in one health state may be greater than in another: individuals rank their interest in health states in this way when choice is necessary, and when their choice is impartial.

Slave-owning society

Harris has also objected to the use of Rawls's veil of ignorance on the ground that it does not guarantee the justice of arrangements that would be decided upon by rational egoists. In support of this he offers the example of rational egoists choosing a slave-owning society, gambling on being a member of the large number of slave-owners who enjoy living luxurious lives, rather than one of the small number of wretched slaves. In our original article we rejected this example. Since the difference between being a citizen in a non-slave-owning society and being a slave-owner in a slave-owning society is not so great as the difference between being a slave in a slave-owning society and being a citizen in a non-slave-owning society, at least if these two societies are otherwise much like ours, it would not be worth taking a chance of ending up as a slave. It would not be rational. Moreover, in the real world, a ratio of one slave to every ten free people would certainly not be enough to make the lives of the slave-owners wonderfully luxurious. And again, the more slaves there are the less rational it would be to take the risk of ending up a slave. It would seem then that if human nature stays roughly the same as it is now, rational egoists behind a veil of ignorance would not choose to allow slavery. This is the argument we put forward in our previous article for rejecting Harris's

example of the slave-owning society.[1] But Harris in turn rejects our argument, and cites the example of Aristotle against it: "Aristotle is notorious for having supported slavery, and in particular for having defended the right of the victors in battle to enslave the defeated, although he could not for sure and certain have known that he, or the Greeks, would never be defeated in battle by non-Greeks, indeed he knew that the contrary was the case.... Greek rational egoists with human nature not a million miles from our own, might well have supported slavery from behind a veil of ignorance, believing in the superiority of their own nature, or believing that those whose nature was not superior would deserve to be slaves".[2]

Twofold response

Our response to this is twofold. First, if Aristotle endorsed slavery in the knowledge that he might be disadvantaging himself – in the knowledge that he was making his own situation worse – then he was not choosing as a rational egoist. For the reasons given above, the more the slave-owning society is otherwise like our own, the more slaves there must be to ensure a luxurious life for the slave-owners and the more wretched the lives of the slaves must seem when compared to our own, and thus the more irrational it would be to choose slavery. In the real world, rational egoists would only support slavery if they knew for certain that they would be among the masters and not the slaves. Secondly, we are not convinced that Aristotle in fact supported slavery in the knowledge that he could end up a slave, or that if this transpired it would be right. Aristotle appears to have supported the enslavement of captives in war only in the case of inferior peoples: "It is part of nature's plan that the art of war, of which hunting is a part, should be a way of acquiring property; and that it must be used both against wild beasts and against such men as are by nature intended to be ruled over but refuse; for that is the kind of warfare which is by nature right".[3] It is unlikely that Aristotle considered himself "by nature intended to be ruled over" – for example, someone who "participates in the reasoning faculty so far as to understand but not so as to possess it".[3] Like most advocates of slavery, it seems that Aristotle supported slavery only for a group of which he did not consider himself a member.

We end by noting a point on which all parties can agree. Harris dramatically points out how a simple policy of QALY maximisation might have adverse consequences for society as a whole: "Where people are frightened not only of suffering illness, or of possessing genes which will likely shorten their life expectancy, or are already coping with the deleterious effects of these; but are also frightened of the effect that others' knowledge of these disadvantages will have on their standing in the community, their access to other dimensions of its care, including

health care, and to rescue and other services, then this is surely likely to have a divisive and corrosive effect on the sense of community".[2] We acknowledged this in our original article.[1] We mentioned that there is more to social welfare than health-related QALYs, and suggested that allocating health care so as to produce the largest possible number of QALYs might not always be the right thing to do. The QALY method is a form of cost-utility analysis, and the overriding goal is to maximise utility, not QALYs. Like cost-benefit analysis, it is a tool of health care allocation, but it can be overridden in the interest of promoting social welfare, just as cost-benefit analysis can. So if it could be shown that allocating health care according to the QALY method would, in some particular case, have "a divisive and corrosive effect on the sense of community" (sufficient to outweigh any benefits its application might be expected to have), then we would be prepared to restrict its application in that case. Unlike Harris, however, we do not want to abandon the QALY method altogether because of this possibility. Rather, we hold that there is a presumption in favour of allocating according to QALYs, except where it can be shown that it conflicts with other concerns – particularly with the wider goal of maximising social welfare.

John McKie is Research Assistant at the Centre for Human Bioethics, Monash University, Clayton, Victoria, Australia. Peter Singer is Deputy Director, and Helga Kuhse Director, of the centre. Jeff Richardson is Director of the National Centre for Health Program Evaluation, Monash University, Fairfield Hospital, Fairfield, Victoria, Australia.

References

1 Singer P, McKie J, Kuhse H and Richardson J. Double jeopardy and the use of QALYs in health care allocation. *Journal of Medical Ethics* 1995; **21**: 144–50.
2 Harris J. Double jeopardy and the veil of ignorance – a reply. *Journal of Medical Ethics* 1995; **21**: 151–7.
3 Aristotle. *The politics*. Translated by Sinclair, TA. Harmondsworth: Penguin Books, 1962: book I.