

RESEARCH ARTICLE

The Contested Value of Life

Søren Holm 

Centre for Social Ethics and Policy, University of Manchester, Manchester, UK

Email: soren.holm@manchester.ac.uk

Abstract

Putting a specific value on human life is important in many contexts and forms part of the basis for many political, administrative, commercial, and personal decisions. Sometimes, the value is set explicitly, sometimes even in monetary terms, but much more often, it is set implicitly through a decision that allows us to calculate the valuation of a life implicit in a certain rule or a certain resource allocation. We also value lives in what looks like a completely different way when we evaluate whether a particular life is being or has been lived well. Both of these ways of valuing are done from an outside or third-person perspective, but there is also a third way of valuing a life which is from the first-person perspective, and which essentially asks how much my life is worth to me. Is there any connection between these different ways of valuing life, and if so what is the connection between them? This paper provides an account of John Harris' analysis of the value of life and discusses whether it can bridge the gap between first-person and third-person evaluations of the value of life, and whether it can do so in a way that still allows for resource allocation decisions to be made in health care and other sectors.

Keywords: John Harris; QALY; resource allocation; value of life

Introduction

Putting a specific value on human life is important in many contexts and forms part of the basis for many political, administrative, commercial, and personal decisions. Sometimes, the value is set explicitly, sometimes even in monetary terms, but much more often, it is set implicitly through a decision that allows us to calculate the valuation of a life implicit in a certain rule or a certain resource allocation. In many contexts, the value that is set for a life is the value of a “statistical life,” when the actual persons who are going to be affected by a rule or resource allocation are unknown to the decision-maker and the effect on their lives is in the future. However, there are also contexts, especially in health care, where the persons whose lives are affected are present and identifiable.

We also value lives in what looks like a completely different way when we evaluate whether a particular life is being or has been lived well. This is the kind of evaluation recounted by Herodotus in his account of the discussion between Solon and Croesus, where we find the famous quote attributed to Solon:

If besides all this he ends his life well, then he is the one whom you seek, the one worthy to be called fortunate. But refrain from calling him fortunate before he dies; call him lucky.¹

Both of these ways of valuing are done from an outside or third-person perspective, but there is also a third way of valuing a life which is from the first-person perspective, and which essentially asks how much my life is worth to me.

Is there any connection between these different ways of valuing life, and if so what is the connection between them?

My life—my value

In Chapter 5 of his seminal book “The Value of Life,” John Harris discusses issues of resource allocation between two different patients who both need the same life-saving medical resource, but where scarcity means that only one of them can get it.² In his analysis and discussion of this issue, he develops a distinctive account of the value of life. According to Harris, the value of a life is determined by the person whose life it is. They are the only ones who can decide whether their life has value or not. And everyone who values their life values it equally:

All of us who wish to go on living have something that each of us values equally although for each it is different in character, for some a much richer prize than for others, and none of us know its true extent. This thing is of course ‘the rest of our lives’. So long as we do not know the date of our deaths then for each of us the ‘rest of our lives’ is of indefinite duration. Whether we are 17 or 70, in perfect health or suffering from a terminal disease we each have the rest of our lives to lead. So long as we each fervently wish to live out the rest of our lives, however long that turns out to be, then if we do not deserve to die, we each suffer the same injustice if our wishes are deliberately frustrated and we are cut off prematurely.³

It is not always clear whether Harris thinks that we suffer the same injustice if our life is ended because we all value our life equally, or whether the inference goes in the other direction and the equal value of life follows from the fact that there is equal injustice. For present purposes, this does not matter.

It follows from Harris’ account that neither age nor the quality of life matters when valuing a life. All valued lives have equal value, and in that sense, every life counts as one, and no life counts as more than one. He also argues that neither dependents nor family, social usefulness, or moral worth matter in relation to the value we should ascribe to a life when allocating life-saving resources.

In the same chapter where he develops his account of the value of life, Harris discusses the so-called “fair innings” argument, that is, the argument that whether or not someone has had a life long enough to allow them to experience all stages of life should matter when allocating life-saving resources. Those who have not had a fair innings should have priority over those who have. As Harris notes with his usual flair for turning a phrase:

It is sometimes said that it is a misfortune to grow old, but it is not nearly so great a misfortune as not to grow old.⁴

On this basis, he concludes that age indifference should be the norm, but that age can never the less matter in cases where we choose between people who have and have not had a fair innings:

It is for this reason that in the hopefully rare cases where we have to choose between candidates who differ only in this respect that we should choose to give as many people as possible the chance of a fair innings.⁵

Linking the value of life from a third-person perspective to the value of life from the first-person perspective and directly linking the injustice inherent in the ending of a life to that evaluation also neatly solves the ethical issues raised by voluntary euthanasia of competent persons, and of non-voluntary euthanasia of those in persistent vegetative states or similar conditions. If a competent person wants to have their life ended by euthanasia, and if they state that they no longer value their life, there is no injustice inherent in killing; and if a person has permanently lost the capacity to value their life, there is

again no injustice inherent in killing. The person who performs euthanasia in such circumstances does not wrong the person being killed, if Harris' account of the value of life is correct.

Resource allocation

John Harris' account of the value of life is, as described above developed in the context of a choice between two persons, both needing a medical resource to save their life. How does this account fare in the context of macro-level resource allocation, where the allocation decision is future directed and usually involves "statistical lives"?

As mentioned above, Harris is to some degree sympathetic to the fair innings argument in "The Value of Life," but he later changes his position, rejects the fair innings argument completely, and endorses a strict principle of age indifference.⁶ So, the position held by the mature Harris is that all lives that are valued by the persons who lead them have equal value.

This view has direct application to resource allocation situations where the only issue at stake is the number of lives saved (but see next section for a complication). Harris is, in many papers, scathing about "innumerate ethics," a term he borrows from Derek Parfit⁷ and argues consistently that if we have to choose between saving one life and saving more, we should choose to save more⁸. This view is, of course, also fully consistent with the analysis presented in his perhaps most famous paper "The survival lottery" which predates "The Value of Life" by 10 years⁹. He further argues that the same applies when we have a choice between saving one life now and more in the future, for instance in relation to setting priorities in public health.¹⁰ This analysis can easily be extended to resource allocation choices where we do not know the number of lives saved for certain for any of the options, but do know the probabilities of life-saving. In such cases, we should allocate the resources to the option that is likely to save the most lives.

Most resource allocation decisions in health care, and in many other contexts are however not "pure" in the sense that the only difference between the options is in the (likely) number of lives saved. Many resource allocation decisions do not involve life-saving at all, many involve a combination of life-savings and improvements of health states, and many involve a combination of life-savings and life extensions. What are the implications of Harris' view of the value of life for such "mixed" resource allocation decisions. Harris discusses this issue in another of his classical papers "QALYfying the value of life," taking as the point of departure the proposition put forward by Alan Williams that resource allocation decisions should maximize the number of Quality Adjusted Life Years (QALY) produced by the decisions.¹¹ Harris quotes Williams's description of the QALY measure:

The essence of a QALY is that it takes a year of healthy life expectancy to be worth one, but regards a year of unhealthy life expectancy as worth less than 1. Its precise value is lower the worse the quality of life of the unhealthy person (which is what the "quality adjusted" bit is all about). If being dead is worth zero, it is, in principle, possible for a QALY to be negative, i.e. for the quality of someone's life to be judged worse than being dead. The general idea is that a beneficial health care activity is one that generates a positive amount of QALYs, and that an efficient health care activity is one where the cost per QALY is as low as it can be. A high priority health care activity is one where the cost-per-QALY is low, and a low priority activity is one where cost-per-QALY is high.¹²

It is not surprising that Harris takes issue with the QALY, exactly because use of QALYs means that the lives of people are allocated unequal value. Saving some lives will produce more QALYs than saving other lives, and improving the quality of life of some people may produce enough QALYs to outweigh saving the lives of other people.

Harris also argues that QALY maximization is ageist, can be racist and sexist, and discriminates against the disabled because they face double jeopardy. If you are disabled and already have a quality of life less than 1 because of your disability, saving you will generate less QALYs than saving someone who is not disabled and can be brought back to a quality of life score of 1.

According to Harris, allocating in order to maximize QALY production is thus fundamentally unethical because it does not treat people as equals. He states forcefully that:

Because my own life would be better and even of more value to me if I were healthier, fitter, had more money, more friends, more lovers, more children, more life expectancy, more everything I want, it does not follow that others are entitled to decide that because I lack some or all of these things I am less entitled to health care resources, or less worthy to receive those resources, than are others, or that those resources would somehow be wasted on me.¹³

But where does this leave the healthcare planner who has to allocate a fixed budget? Harris argues that the planner's problem can be solved by making more resources available, and that society should do so:

QALYs encourage the idea- that the task for health economics is to find more efficient ways of doing the wrong thing - in this case sacrificing the lives of patients who could be saved. All people concerned with health care should have as their priority defensive medicine: defending their patients against unjust and lethal policies, and guarding themselves against devices that tend to disguise the immorality of what they are asked to do.

[...]

It is implausible to suppose that we cannot deploy vastly greater resources than we do at present to save the lives of all those in immediate mortal danger. It should be only in exceptional circumstances - unforeseen and massive disasters for example - that we cannot achieve this. However, in such circumstances our first duty is to try to save the maximum number of lives possible. This is because, since each person's life is valuable, and since we are committed to treating each person with the same concern and respect that we show to any, we must preserve the lives of as many individuals as we can. To fail to do so would be to value at zero the lives and fundamental interests of those extra people we could, but do not, save. Where we cannot save all, we should select those who are not to be saved in a way that shows no unjust preference. We should be very clear that the obligation to save as many lives as possible is *not the obligation to save as many lives as we can cheaply or economically save*. Among the sorts of disasters that force us to choose between lives, is not the disaster of overspending a limited health care budget!¹⁴

Is this a realistic prescription? Perhaps, if what we have in mind is a rich country and a situation where the health care budget can be increased significantly without taking into consideration any other spending. But, resource allocation has to happen in low- and middle-income countries as well, and the healthcare sector is not the only sector where life-saving occurs under budgetary constraints. So, the Harris solution is not a solution that can actually be implemented in any currently existing society.

A further problem for the Harris position is that whereas it might be true that we all value our lives and that we value them equally, it is also true that each of us values interventions that can make our lives better. If I am in pain I value analgesics. If I have arthrosis of the hip, I value a hip replacement. If I have severe mobility problems, I value an electric scooter. And, I value these things even in cases where I know that their only effect is to improve my quality of life. A health care system that denied equal access to any of these interventions would plausibly be just as unethical as one that denied equal access to life-saving treatments. But, Harris' position denies that there is any metric for comparing life-saving and other beneficial effects of health care interventions. It is thus unable to provide any action guidance in relation to the many situations where the resource allocation options are mixed, except perhaps in those situations where one option dominates the other in relation to both life-savings and increase in welfare independent of the life-savings.

One way out of this problem, which seems to be implicit in some of Harris' argument is to claim that life-saving is more important than any other activity. Harris, for instance, raises what seems to be a mainly rhetorical point:

Instead of attempting to measure the value of people's lives and select which are worth saving, any rubric for resource allocation should examine the national budget afresh to see whether there are any headings of expenditure that are more important to the community than rescuing citizens in mortal danger. For only if all other claims on funding are plausibly more important than that, is it true that resources for life-saving are limited.¹⁵

The view that life-saving is more important to the community than any other headings of expenditure, does however not follow from Harris' analysis of the value of life, but is a separate claim that needs justification. Harris' analysis may show that the value of life is equal for anyone who values their life, but it does not provide any "magnitude" for that value, which allows us to determine the strength of a claim based on that value. The value cannot be infinite, since it would then be trivially true that the number of life-savings would not matter since $n \times \infty = \infty$. It is intuitively plausible that the value must be high, e.g., because the intentional killing of another, that is, murder is seen as one of the most serious crimes in most societies. But, it is not obvious that it always lexically outranks all other values. Even a high value can be outweighed, and there are plausibly other things that a society spends money on that are also of high value. Children, for instance, presumably have an equal and very important claim to an adequate education, and a society which did not allocate enough resources to education to discharge its responsibilities to children could not plausibly claim that that was justified by marginal increases in life-saving. Or to put the point more generally, if the value of a life is finite, then there will be a number of non-life-saving acts preserving or promoting other values that add up to one life-saving. We might want to discount the aggregation of other trivial values and rule out that even a large number of alleviations of minor headaches can outweigh a life-saving, but we cannot discount the aggregation of non-trivial values.

A problem in counting instances of life-saving

A problem occurs if we think that life-saving is important and that it is *pro tanto*, a right-making feature of an act that saves a particular number of lives or saves more lives than another available act. The problem is that it is not obvious that we can count instances of life-saving in a non-arbitrary way, and thereby uphold a clear distinction between life-saving and life-extension (or death-postponement).¹⁶ Barring the possibility of true resurrection from the dead, a person can only die once, but their life can be saved many times, and saved many times within a short period of time.

Let us define a life-saving act as any act that is necessary and sufficient to prevent a person from dying at a specific time, where "necessary" is defined in terms of practical and not logical necessity. Wading into a pond to save a child that is drowning will then count as a life-saving act as will initiating ventilator treatment of a patient with complete respiratory failure. Hopefully, the child will not need life-saving again in the immediate future, and it will have learned an important lesson and stay clear of ponds. The patient with complete respiratory failure will, however, need a series of consecutive life-saving acts performed by the ventilator. This may be obscured by the fact that ventilators are now mechanical and will continue to ventilate as long as we do not turn them off, but before mechanical ventilators were invented during the polio epidemics in the early 1950, ventilation was done by hand by human ventilators (Figure 1).¹⁷

We might still think that the human ventilator only performs one life-saving act, with that one act consisting of many distinct bodily actions, but what happens when the human ventilator's shift is at an end, and they are replaced by another human ventilator? That other human actor must consciously decide to start ventilating. They would obviously be morally delinquent if they did not start to ventilate, but that does not show that they do not have to decide to do so. But, that seems to imply that a separate act token has been initiated, although it is (hopefully) of the same act type, that is, a life-saving act of sustained manual ventilation, as the act immediately preceding it. The child with polio and respiratory failure thus has its life saved numerous times each day.

The same counting issue will occur if we claim that it is the number of life extensions that matter, but it can be circumvented if we focus on the length of life extension. However, on Harris' account what matters is life-saving and not life-extension since the value of a life is independent of its length.



Figure 1. Hand ventilation during the polio epidemic in Copenhagen 1952.

The value of a life well lived

As noted in the introduction, there is what initially seems to be a completely separate way to talk about the value of a life when we consider whether a life has been or is being lived well. The answer to this question clearly admits of degrees. Few people live a perfect life, many live their lives well, and a few live truly bad lives. Harris' analysis of the value of life shows that this conception of the value of life is orthogonal and independent of the question of how we should value a life when considering life-saving. We have no reason to save better lives than to save lives that are not going too well. But it might nevertheless be important to consider what it means for a life to be lived well, and who has priority in making that assessment, the person themselves, or a third-person observer. The quote above that Herodotus ascribes to Solon implies the primacy of the third-person perspective. If we can only truly assess a life when the person whose life it was is dead, then the third person perspective is the only option, if we discount the possibility of an afterlife from which the person can evaluate their own, past earthly life. On the other hand, there are good arguments against giving the third-person perspective primacy when we consider the life of the still-living. The person living a life will often be in a position to correct an outside evaluation of their life, because they are in a better position to know what their life goals are, and to what extent those goals have been or are being fulfilled. Unless you know something about my life of which I am ignorant, it would be astonishing hubris to claim that my life is going badly, when I think it is going well. It may well be the case that there are ways in which my life could go better, and that you can identify some of them, but that does not necessarily show that it is not going well.

So far, I have implicitly assumed that a life goes well if the person in question achieves their life goals, and that the only dispute is about who is best placed to evaluate this. However, if those goals are unethical, we may want to dispute them. The person may think they lead a successful life, gradually fulfilling their goal of genocide, but from the outside, we can see that this is not a life well lived, simply because it is lived in the pursuit of perverse and unethical goals. There may thus be circumstances where the moral evaluation of a life is best performed from a third-person perspective, although even here, there are probably distinctions to be made. Someone may lead their life according to a moral code that is not mine, and that I find incomprehensible, but unless that code is positively immoral, they may still be best placed to evaluate whether their life is lived well.

Conclusion

John Harris' analysis of the value of life provides an important counterpoint to positions that claim that the value of life differs widely, and that these differences should form the basis for resource allocation decisions and other decisions in which we have to choose who to save. He is clearly right that seen from the first-person perspective everyone who values their life values it equally, and right that this equality must matter when we decide who to save. His position does, however, create a set of very significant problems for resource allocation decisions in health care and more generally. Even if we thought that we can enumerate instances of life-saving unambiguously, we still have the problem that it is quite implausible that life-saving is the only thing that matters morally. And every time we trade off one life-saving against a large number of broken bone settings, as we must do, we put an implicit, finite value on the life that is not saved.

A similar duality emerges when we examine how to evaluate whether a life has been or is being lived well. In many ways, the first-person perspective on this must be the primary perspective, but we cannot deny that the third-person perspective is relevant in some instances.

Perhaps the only thing we can do is to hold onto this duality. The value of a life is decided by the person whose life it is, and all lives are equal in this respect, but their value is not infinite, and it can, in some circumstances, be traded off against other equally important values.

Notes

1. Herodotus, with an English translation by A.D. Godley. Cambridge: Harvard University Press; 1920, at 1.32.7
2. Harris J. *The Value of Life*. London: Routledge and Kegan Paul; 1985.
3. See note 2, Harris 1985, at 89.
4. See note 2, Harris 1985, at 101.
5. See note 2, Harris 1985, at 102.
6. Harris J. Does justice require that we be ageist? *Bioethics* 1994;8(1):74–83; Harris J. The age-indifference principle and equality. *Cambridge Quarterly of Healthcare Ethics* 2005;14(1):93–99.
7. Parfit D. Innumerate ethics. *Philosophy & Public Affairs* 1978;7(4):285–301.
8. See note 6, Harris 1994; Quigley M, Harris J, Roberts J. Personal or public health? In *International Public Health Policy and Ethics*. Cham: Springer International Publishing, 2024:31–46.
9. Harris J. *The survival lottery*. *Philosophy* 1975;50(191):81–87.
10. See note 9, Quigley et al. 2024.
11. Harris J. QALYfying the value of life. *Journal of Medical Ethics* 1987;13(3):117–123.
12. See note 11, Harris 1987, at 117.
13. See note 11, Harris 1987, at 121.
14. See note 11, Harris 1987, at 122.
15. See note 11, Harris 1987, at 122.
16. Holm S. The ethics of death policies. In *The Routledge Handbook of Ethics and Public Policy*. London: Routledge; 2018:381–392.
17. Berthelsen PG. Manual positive pressure ventilation and the Copenhagen poliomyelitis epidemic 1952. *Acta Anaesthesiologica Scandinavica* 2014;58(5):503–507.