Journal of Medical Ethics 2001;**27**:251–255

How economics could extend the scope of ethical discourse

Alan Williams *University of York*

Abstract

Ethical discourse is typically inconclusive, and with good reason. But this inconclusiveness is a distinct disadvantage when it comes to helping publicly accountable policy-makers in the health care system provide an ethical justification for their decisions. It is suggested that instead of ending with platitudinous statements such as that a balance has to be struck between the rival ethical considerations, empirical research should be undertaken to elicit the quantitative trade-offs that the affected general public would be prepared to accept when striking this balance. In the expected absence of any consensus, it is further suggested that the views of the median person be taken as the best approximation to the group view. Finally it is argued that, far from this quantitative approach lacking humanity by treating individuals as "mere statistics", it shows greater compassion than the proponents of those approaches whose fellow feeling can only be stirred by information pertaining to identified individuals.

(*Journal of Medical Ethics* 2001;**27**:251–255)

Keywords: Health economics; QALYs; equity; trade-offs; "fair innings"; "statistical compassion"

Scene-setting

Ethical discourse is typically inconclusive. There are good reasons why this should be so. The premisses on which it is based are usually contestable. There is usually more than one principle in play. No single principle "trumps" all others. The situations that are selected for analysis are complex ones where the appropriate resolution of the ethical difficulties is not self-evident. The objective of the discourse is frequently claimed to be no more than offering clarification of, rather than recommending any specific "solution" to, the issues involved. In that context the most that can be said is "if you attach more weight to principle X than to principle Y, you should choose course of action A, but if you attach more weight to principleY than to principle X you should choose course of action B". In the field of medical ethics the situations and courses of action that are presented and analysed are, typically, carefully constructed scenarios involving hypothetical people in hypothetical situations. Sometimes additional information is provided, step by step, to see how sensitive the initially recommended course of action is to variations in the initial scenario.[1] Whether and to what extent a course of action should be recommended is typically tested by pointing out not only its consequences in the stated scenario, but also where adoption of the same ethical principles would lead if applied more widely. There seems to be a desire for some broader consistency, and a general aversion to having ethical principles become too context-specific. This implies that there should be some partitioning of decisions into distinct classes, within each of which the same ethical principles should apply, but between which different sets of principles should apply. There is one further strand in ethical discourse which is worth noting, and that is the frequent appeal to intuition, or to allegedly widespread acceptability or support for the principles that have been adduced, as a further justification for according them some inherent merit (irrespective of their consequences). Since there is usually some ambivalence about some of the consequences (otherwise there would be no ethical dilemma to analyse), the bad aspects of these consequences are labelled repugnant, or unfortunate, or tragic, or some such term, and there the matter is left.[2]

Helping people to be clear in their own minds about the ethical implications of their actions is not a trivial pursuit, and in the context of publicly accountable decision making it is especially important, since the principal actors are expected to be able to provide justification for their actions, and not to behave arbitrarily or capriciously. But it is not sufficient for them merely to list the various things they claim to have taken into account. The citizenry are entitled to know what weight they gave to each, so that they can see what it was that proved decisive. Otherwise, the same bland listing of relevant principles could provide justification for almost any decision.

This is where economics can carry things forward. Economics is about what people value and how much they value different things. An efficient economic system ensures that what people have to sacrifice in order to get something they want is of less value than the thing that they want (and that this difference is as large as possible). The "thing" they value may be a commodity, or a service, or a scenic view, or an endangered species, or the upholding of the rule of law, or the observance of some ethical principle. Such "things" become of interest to economists when in order to gain them some other valuable thing has to be sacrificed, or, in plainer terms, when they are costly. "Value" can of course be interpreted in many different ways. At its crudest it becomes market value, and efficiency is then about ensuring that the market value of a good

or service exceeds the financial cost of providing it. But it could equally well refer to the subjective value, or the social value, of whatever items are involved in the transaction. An efficient system in that context is one that ensures that the (subjective or social) value gained is greater than its opportunity cost (ie what has to be sacrificed in order to obtain it), where the latter is measured in comparable units to the former.[3] Thus the "opportunity cost" of improving one person's health might be making the health of some other person worse. When there is more than one criterion in play when judging either the subjective or the social value of an item, then a slightly different "trade-off" problem presents itself, which is people's willingness to sacrifice performance on one criterion in order to improve performance on some other criterion. Technically this "trade-off" is known in economics as the marginal rate of substitution between two goods, where normally the two "goods" are commodities. But the concept is equally relevant when the two "goods" are outcomes such as the health of person A and the health of person B, or when the two "goods" are two ethical principles. The important additional matter to be noted here is the *marginal* nature of this concept. It is the rate at which, given a particular starting point, a small amount of the one good will be exchanged for a small amount of the other good, whilst leaving the overall level of wellbeing constant. Note also the important caveat about the starting point, because the marginal rate of substitution will not generally be constant across all choices, but vary according to the context. Nevertheless it can in principle be derived from a much more general "welfare function" which takes into account these different choice possibilities within a larger "decision space".[4]

Let us then explore the ways in which this economic mode of discourse might complement and extend the usual mode of ethical discourse. I will first set up a hypothetical scenario, then work my way through a (slightly caricatured) typical ethical discourse related to it, and then continue beyond that into the territory of economics, to see where such an excursion might lead. I shall conclude that it can greatly improve the clarity and relevance of the ethical discourse.

The scenario

Suppose there exists a publicly accountable agency charged with the two tasks of improving the health of the community and reducing inequalities in health within that community. Suppose further that there is general agreement within that community as to what health actually is (for simplicity let us assume temporarily that it is years of life lived, or life span). Suppose further still that the resource constraints facing the agency when it is pursuing these objectives are quite clear, and beyond its control. The fundamental ethical issue which the agency has then to resolve is whether a year of life *is to be regarded as* of equal value no matter who gets or loses it, and, if not, how the differential value of a year of life is to be determined.[5] Note the phrase "is to be regarded as" does not mean "is", because we are assuming that the agency itself has the task of working out what a year of life to different individuals is to be valued at *from the agency's own perspective*, and this may be very different (say) from what each individual would be willing to pay (or accept by way of financial compensation) to have his or her own life lengthened or shortened when viewed from the individual's own perspective.

A possible ethical discourse

It is tempting to start by arguing that, on some egalitarian principle or other, every year of life enjoyed by every member of the community should be accorded the same value. But this simple-minded approach will not do in the present context, because to adopt it would imply the total rejection of one of the two objectives which the agency is supposed to be pursuing, namely to reduce inequalities in health. It is extremely unlikely that the single-minded pursuit of the health-maximisation objective will, quite coincidentally, also minimise inequalities in health. It is far more likely that a balance will have to be struck between them, which means that extra years of life to the more deprived (in terms of health) will have to be regarded as of greater social value than extra years of life to the more favoured (in terms of health).

The next stage of the discourse is therefore likely to concern itself with rival ethical bases for choosing how much weight to give to the reduction of certain kinds of inequality versus other kinds of inequality, and whose health should be sacrificed in order that any particular inequality should be reduced. Here the discourse is likely to focus on whether any particular inequality is beyond anyone's control (either as regards its cause or its alleviation), and, if it is controllable, whether it is the product of the voluntary (and well-informed) behaviour of the individuals affected, whether the inequality is beneficial or damaging in some other way to society, and how large it is. Generally it will be possible to deploy ethical arguments both in favour of and against any particular stance on any of these issues, and, a fortiori, on the weight to be attached to each when looking at the big picture.

The manner in which the protagonists then attempt to convert others to their particular standpoint is, typically, to point out the distasteful consequences of the wider adoption of their opponents' views. Thus those who think that we should not regard smokers as blameworthy for their shorter life span point to other behaviour which adversely affects health, and ask whether those people should not be treated in the same way as smokers, and, if they were, at the end of the day will anyone be left blameless? On the other side, the question posed is why should those who have behaved responsibly with respect to their health be penalised by having resources diverted away from them to help the others, thereby reducing the very life span they had striven to extend. Surely this is to

deny people their just reward for virtuous behaviour, and that sends out the wrong signals to the members of the society, thereby encouraging antisocial behaviour of all kinds.

A similar set of conflicting views can be paraded with respect to intergenerational equity. The maximisation of life span is likely to concentrate emergency treatment upon the young, because by so doing the gains in life expectancy are much greater. To counteract this tendency, the right of all citizens to equal access to life-saving treatments is adduced as a relevant principle to prevent any discrimination on grounds of age.[2] Against this it might then be argued that the older members of the community will already have had a "fair innings" in terms of life span already enjoyed, so greater weight should be given to helping younger people whose lives are under threat than to helping older people whose lives are under threat. Indeed the "fair innings" argument might justify discriminating against the old to an even greater extent than would be justified by the health-maximisation objective.[6] The persuasive powers of the various protagonists are then typically deployed on a line of argument where each probing question starts "how would you like it if....". How would you like it if, when you are old, having paid all your life for health care to be available when you need it, and never having needed it until now, you are told that you are now too old to be a high priority? Conversely, how would you like it if it were your child that was denied life-saving treatment because this old person is getting higher priority because he or she has paid more taxes into the system over his or her lifetime than you have? And so it goes to the end of its inconclusive path.

One possible way out of this dilemma is to put it to a vote amongst the members of the community whose welfare is at stake in one way or another. My impression is that the public opinion poll solution is distrusted by most ethicists, who usually cite capital punishment as an example of where public opinion (which is reckoned to be in favour of capital punishment) is obviously wrong. Part of this distrust concerns how issues are framed in such polls, part of it is about how representative the sampled population is, whether they mean what they say, and whether they really understand the problem and what is at stake. The favoured solution[7] tends to be to ensure that the voters are thoroughly briefed (by neutral fair-minded experts rather than by advocates of particular positions) before expressing their opinions, and that they are then required to produce an acceptable ethical justification for their choices (this acceptability being determined by ..... whom?) Thus, in the capital punishment case, I suppose that a person voting for capital punishment because they believed it would deter potential murderers would have their vote disqualified, because (say) there is no evidence that capital punishment does have a deterrent effect on murderers. But anyone who voted for capital punishment because of a strong belief in retribution (an eye for an eye, a tooth for a tooth ... a life for a life) would have their vote accepted. A vote cannot be disregarded just because you don't agree with the ethical principle on which it is based. If the predominant ethic in the community in which you live is not acceptable to you, you can try to convince people that you are right and they are wrong, but until you have succeeded in that task you either put up with it or move. You could try to generate a more general discussion of the rival objectives of the criminal justice system (retribution, deterrence, rehabilitation, protection, etc) and ask people also to consider the relative importance of ensuring that no innocent person is ever unjustly punished, as opposed to ensuring that no guilty person ever goes unpunished, in the hope that this moves people in your direction, but it might well not do so. Those hell-bent on retribution might start dismantling the rest of the apparatus to which you have so foolishly drawn their attention!

A possible excursion into economics

It is precisely when matters get into an impasse that cannot be resolved by logical argument that it is useful to look behind the usual words that are used in ethical discourse when the problem is handed over to the reader. Characteristic phrases are "a compromise has to be found", "a balance has to be struck", "a judgment has to be made" or "it will depend on how much weight is to be attached to ....". All of these point to the presence of what I earlier termed "a trade-off problem", and the first task for any economic analyst would be to formulate that problem in such a way that it becomes susceptible to empirical analysis. But the empirical analysis here goes beyond head counts (though the "weight of opinion" may well include a head count of some kind). It starts from the observation that it is not simply that some people are 100% for retribution and to hell with everything else, while others are 100% for rehabilitation or some other selected objective of the criminal justice system. More often we would each give *some* weight to more than one of these objectives, and perhaps to all of them. The question is, how much weight, and how context-specific is it?

One way into this is to present people with binary choices and see which alternative they opt for, and then, step by step, make their preferred option less and less attractive until they switch. The stage at which they switch can then be used to estimate the quantitative weight they attach to the two rival objectives. Let us go back to the earlier scenario of an agency charged with the twin objectives of improving the overall health of the population as much as possible, and of reducing inequalities in health within that population (where "health" is still being interpreted as "life span"). Let us assume that the context is that there are two equally-sized groups in the community, of which the better-off group enjoys a life expectancy at birth of 78 years, and the worst-off enjoys a life expectancy at birth of 73 years. We could then offer for consideration the following alternative policies (which cost the same to implement):

Option A: Increases the life expectancy of everybody by two years.
Option B: Increases the life expectancy of the worst-off group by four years and leaves the better-off group unaffected.

Those who select A are obviously unconcerned about inequality reduction. Those who select B are concerned, but the interesting question is *how concerned* are they? Suppose we came back to the latter group and presented them with a second choice, in which Option A is the same as before, but Option B is rather less attractive, offering only three extra years to the worst-off, and again nothing to the better-off. Some people will now switch to Option A, but not all, so we can go on and reduce the gains to the worst-off under Option B to two years. Those who still vote for B are telling us that they care so much about inequalities in health that they would be prepared to deny gains of two years to the better-off (which could be obtained at no extra cost) in order to reduce the level of inequality from five years to three years. And there may be some who would go still further.[8]

What should we make of this diversity of opinion if we were the policy-makers within the democratically accountable public agency which has to choose between these options? There is a lot to be said in favour of taking the median person as the one to whom we should regard ourselves as accountable, because in any majority voting system that person will always be in the majority. But from one decision to another it will not always be the same person who is the median voter, and we cannot submit every decision we have to make to some kind of referendum, so we need to do something rather more complex than the simple experiment sketched out above in order to get some systematic guidance as to where the (changing) median voter is likely to stand in the complex territory in which we are operating.

We could start by filling in some detail about the characteristics of the groups that were designated "better-off" and "worst-off" above. Suppose we had said initially that they were (respectively) babies born into households where the chief breadwinner was a professional person, and babies born into households where the chief breadwinner was an unskilled manual worker. We now change that, and in a subsequent experiment say that the better-off group are females, and the worst-off group are males (the initial differences in life expectancy at birth being equally true in this case). Does this change the original ratings in any way? We could then inquire why the ratings had or had not changed. And we could go on, in the same way, to explore ethnic differences, geographical differences, intergenerational differences, and so on, and in each case get some quantitative indicator of how important these different inequalities in life expectancy at birth actually are to the median voter. We could also find out whether they are associated with the particular socio-demographic or other characteristics of respondents (for example the extent to which they might be self-interested). Thus instead of limply concluding that a balance has to be struck between potentially conflicting objectives, we could go on to say where the general public would like the balance to be struck. It seems to me that evidence of this kind has more to commend it than vague appeals to "intuition" or *alleged* "wide acceptability".

But there are other complications that also deserve consideration. To regard "health" as simply a matter of life span is clearly inadequate. The (health-related) quality of people's lives also matters. We would not regard someone whose life had been filled with pain and disability as having been as healthy as someone who lived just as long but without those disadvantages. So instead of measuring people's healthiness in terms of number of years lived, we need to measure it in terms of quality adjusted years of life, where the quality adjustment in a public policy context also needs to be that warranted by the views of the median person. This too can be elicited by trade-off experiments of the general kind outlined earlier. Such an approach will enable us to judge how much extra weight we should give to a particular health gain because it will go to a long term disabled person compared with the weight it would attract if it went to someone whose past health had been good. The implication of this richer concept of health is that instead of being concerned with inequalities in people's expected life span, the focus has shifted to their expected *quality adjusted* life span.

This might well have some bearing on the two experiments sketched out above, between social classes on the one hand, and between males and females on the other.[4] The differences in life expectancy are very similar in each case, but the differences in *quality adjusted* life expectancy are not. In the case of social class differences, the group with the lowest life expectancy also has the worst experience of pain and disability during their lives, so that the differences in *quality adjusted* life expectancy are much larger than the differences in life expectancy itself. But in the case of males and females, the group with the longest life expectancy (the females) has a slightly worse experience of pain and disability during their lifetime, so that the differences in *quality adjusted* life expectancy are slightly less than the differences in life expectancy itself (though a marked difference in favour of females still remains). So among the experiments that need to be conducted are those designed to test the sensitivity of these ratings to the manner in which the inequality is measured and to its magnitude (since it seems likely that for those who are averse to a particular kind of inequality, they will be more averse the greater the inequality is).

Aversion to quantification

However, assessing aversion to inequality by trying to find out what sacrifice in the overall level of health people would accept in order to reduce it by some specified amount, seems to elicit in some people an aversion to quantification. One understandable reason for this is that some people do not think in that way at all. At best they find it hard to

grasp, and at worst they are totally alienated by it. It is always a good strategy to maintain a sceptical attitude to things you do not fully understand, and to try to formulate questions to the cognoscenti designed to allay or confirm your doubts or fears. But aversion to quantification often goes further than this, with outright rejection on the grounds that it lacks humanity, and reflects a rigidly mechanistic view of life, lacking sensitivity to the infinite variety of human experience.

However they are reached, all decisions about priorities, made on behalf of groups, are bound to "lack sensitivity to the infinite variety of human experience" if by that is meant that they will not reflect the choice that each and every individual would have made had that individual had no one but himself to think about. So that is not a consequence of quantification, but a consequence of having to make a single decision on behalf of a heterogeneous group. However, since this is inevitable, it has to be accepted as inevitable, and the issue becomes how best to minimise its adverse consequences. Ensuring the fullest possible public accountability of the collective decision makers is the process that has mostly been relied upon for this purpose, allied to a majority voting system through which those affected can make their views count. But that is precisely how the "trade-off" ratios were generated, and on a more sophisticated basis than anything that ordinary voting systems can generate. These trade-offs can be made to generate "equity weights" which can be promulgated by the public agency as reflecting its policy, so that we know precisely where it stands on each issue (and which it can use internally as a test of its own consistency, which is important for purposes of horizontal equity).[8]

But then we run into the objection about a "rigidly mechanistic view of life" which this seems to entail. It is, of course, no more "rigidly mechanistic" than the rule of law, which enshrines sentencing principles which are also quantitative, and which are based on some weighing of the rival purposes of the criminal justice system that were outlined earlier. But the weight given to each of these principles changes through time as "hard cases" come to light, and/or as public opinion shifts, and/or as new circumstances arise to which the existing rubric does not seem to apply at all comfortably. The courts are publicly accountable agencies establishing priorities in the light of public policy, and then assessing allegedly factual evidence to decide which side of a line a particular case falls. Presumably those who think that quantification in setting health care priorities is rigidly mechanistic think the same about sentencing policy by the courts. But we live with it, because it protects us from a greater source of injustice.

Finally we come to quantification as lacking humanity. The more specific accusation here appears to concern a lack of compassion generated by treating people as "mere statistics". I must confess that it continues to bother me that people can show enormous amount of concern for a single identified sufferer, but fail to show much concern when there are 100 anonymous people reported in the same condition. This seems to me to reflect a lack of humanity on the part of the unimaginative people who cannot see that behind the "mere statistics" there are (a lot of) real people, who do have names and addresses and worried loved ones, even though we don't know them personally. It is these unimaginative people, not the statisticians, who are insensitive to the widespread experience of struggling humanity. The statistical approach is an important mechanism for trying to keep everyone's interests in balance. I think a major educational effort is required to inculcate a sense of "statistical compassion" in people, so as to enlarge and extend our concept of individual compassion.

Conclusion

Resort to numbers is not an abdication of moral responsibility, nor an attempt to conceal the moral issues involved. On the contrary, it is a means of avoiding the impasse which much moral argument reaches, when platitudinous generalisations about striking a balance, or finding an acceptable compromise, or weighing conflicting views one against another, block further thought and progress. What we need to do at that point is to seek some context-specific guidance about how much weight should be given to one principle versus another. In a publicly accountable body making decisions on behalf of groups of people within a community, there is much to be gained from finding out what the members of that community think those weights should be, since staying close to them would appear to be a good justification for the policies pursued by such an agency. It is here that the techniques of economic analysis which have been developed to elicit the quantitative trade-offs that people are prepared to make when called upon to sacrifice a little of one good thing in order to get a little more of another good thing can play a useful role in extending the scope of ethical discourse. I commend it as part of the curriculum for all medical ethicists.

Alan Williams is quasi-retired Professor of Economics in the Centre for Health Economics, University of York, Heslington, York.

References

1 Ubel PA. *Pricing life: the controversial role of cost-effectiveness in health care*. Cambridge MA: MIT Press, 2000.
2 Harris J. *The value of life*. London: Routledge and Kegan Paul, 1989.
3 Tsuchiya A, Williams A. Welfare economics and economic evaluation. In: Drummond MF, McGuire A, eds. *Economic evaluation in theory and practice*. Oxford: Oxford University Press, 2001.
4 Williams A. Intergenerational equity: an exploration of the 'fair innings' argument. *Health Economics* 1997;**6**:117-32.
5 Williams A. QALYs and ethics: a health economist's perspective. *Social Science and Medicine* 1996;**43**:1795-804.
6 Lockwood M. Quality of life and resource allocation. In: Bell M, Mendus S, eds. *Philosophy and medical welfare*. Cambridge: Cambridge University Press, 1988.
7 Daniels N. *Just health care*. Cambridge: Cambridge University Press, 1985.
8 Dolan P, Shaw R, Tsuchiya A. To reduce inequalities in health or to maximise health? Paper presented to the Health Economists' Study Group, Nottingham, July 2000 (an expanded and updated version is under review by the *Journal of Health Economics*).