Priority to the Worse Off in the Allocation of Health Care

The purpose of this session of the conference is to assess the question of whether we should give priority to the worse off (or indeed the worst off) in the allocation of health resources, rather than relying solely on cost-effectiveness or equality or both.[1]

I shall be making the case for a certain kind of priority – for a view which includes a limit on the degree to which priority is to be given (this view has come to be called 'Sufficientarianism'). But let me begin by suggesting that the choice between priority and cost-effectiveness or equality, though it is a real one, is not as stark as it might seem. The cost-effectiveness principle has its roots in the utilitarian tradition. Utilitarianism is perhaps the purest form of what might be called 'welfarism' – the idea that what really matters in morality, ultimately, is how well people's lives go for them. Morality is to do with the distribution of benefits and burdens, and according to (act) utilitarianism they are to be distributed so that the level of welfare or well-being is as high as possible. As I see it, Prioritarianism is coming from roughly the same direction. Morality is indeed ultimately a question of the distribution of benefits and burdens, but according to Prioritarianism it matters exactly who receives those benefits and burdens. And because Prioritarianism includes this central idea, it is in fact closer to Egalitarianism than is utilitarianism. Indeed, I believe that many of those who have called themselves egalitarians, both real-life political activists and philosophers, have in fact held views quite close to those I shall describe as sufficientarian. What they have been concerned about is not so much the abstract relation of equality between individual people, but rather improving the lives of those who are doing especially badly.

1. Cost-effectiveness

Why do we spend money on health care? What is it, that is to say, that we are trying to purchase through buying drugs, surgery, and so on? One common view is that health budgets should be spent so as to maximize the overall health of a population. But this mistakes health for an ultimate good. In fact, the good of health is only derivative or instrumental. So the promotion of health cannot in itself justify expenditure on health care. A very large number

[1] For an excellent survey of some of the main issues involving priority, see D. Brock, 'Priority to the Worse off in Health-Care Resource Prioritization', in R. Rhodes *et al.* (ed.), *Medicine and Social Justice*. This paper is based on my 'Equality, Priority, and Compassion', *Ethics* 113 (2003), and I wish to reiterate here my thanks to those acknowledged in that paper, especially Jeff McMahan, Derek Parfit, and Larry Temkin.

of older men, for example, suffer from prostate cancer. There is treatment available for the condition. But in many cases, this cancer will have no disernible effect on the patient's well-being before they die, from other causes. So treatment would be pointless.

What, then, is the final or ultimate good to which health serves as a means? The answer is welfare, utility, or well-being – whatever it is that, in itself, makes a life worth living, or better, for the person living that life. Philosophers and others have argued for millenia about exactly what well-being consists in: pleasure or enjoyment; the satisfaction of preferences or desires; certain 'objective goods', such as knowledge, friendship, or accomplishment. For the moment, there is no need for us to go into that question, since are dealing with well-being as a concept rather than with any particular conception of it (though of course I am denying any view that makes health a good in itself). But of course it will make an important difference to your views about real-life health resource allocation which conception you accept, and I shall say a little more about that below.

It is very tempting to think that, if health care is to promote well-being, then health care budgets should be spent so as to maximize the well-being, or expected well-being, of a population. This is a version of act utilitarianism, and underpins the view that we should seek to maximize so-called 'quality-adjusted-life-years' (QALYs). This view, however, appears to require injustice. I am going to ask you to consider an example, and it might be helpful for me first to say how I understand these examples, and their heuristic role. Essentially, they are there to pose particular practical questions about various options. These options are states of affairs or outcomes, or sometimes actions seen as bringing about certain states of affairs or outcomes. So in the example below, the question is which of two drugs to purchase – A or B. The numbers represent the well-being of various individuals – those in the groups mentioned on the left-hand side. These numbers are not meant to imply that we can, in practice or even in theory, measure well-being precisely. So the numbers can be translated roughly as follows: The ten people in Group 1 are doing very badly; giving them Drug A will benefit them hugely. The two hundred people in Group 2 are doing very well; giving them Drug B will benefit them a little.

I am myself assuming that the kind of well-being that matters in ethics is lifetime well-being – that is, how well individuals' lives go for them as a whole. Some might prefer to concentrate on well-being at a time – that is, how well somebody is doing at this particular time. There is a clear contrast between these two ideas. Your life might overall be a very good one; but right now, as you lie there suffering from pneumonia, you are doing very badly. Similarly, my life might be really bad, perhaps because I suffer from severe asthma;

but right now, as I enjoy my luxury cruise around the Greek islands and my asthma isn't causing me any difficulties, I'm doing fine. I should make clear that the decision to focus on lifetime well-being is itself an ethical one. I'm not claiming that it's required by, say, rationality.

The epistemic role of these and the examples that follow I see in the following way. Because of the influence of John Rawls, it has become quite common in ethical theory to seek what Rawls called a 'reflective equilibrium' between one's views about particular, concrete cases on the one hand, and ethical principles on the other. I myself see our reactions to cases as having no epistemic weight in themselves. Ethical theory is the attempt to find the correct abstract principles, and the role of reflection on particular cases is to point us in the direction of those principles. The principles are really just ways of capturing views about reasons, so what matters in these cases, when you think one outcome or action is better than another, is what reason there is for preferring that outcome or action. And whether you should accept the principle itself depends on how plausible it is taken by itself. It derives no epistemic support from the fact that it is consistent with the cases that led one to it. On the view I'm suggesting, thinking that it did would be like giving extra credence to a king's views on some matter because you find his subjects appear to hold the same views and you know that they are only repeating the king's views because he's ordered them to do so.

Anyway, to return to cost-effectiveness. Consider the following case.

	Group 1 (10 people)	*Group 2 (200 people)*
Status Quo	5	90
Drug A	50	90
Drug B	5	92.5

Total utility gains: Drug A: 10 x 45 = 450
Drug B: 200 x 2.5 = 500

The utilitarian view, then, requires us to improve the lifetime well-being of a large number of extremely well-off individuals by a very small amount, in preference to improving that of a small number of very badly off individuals by a very large amount. Utilitarians may try to avoid this objection through claiming that principles of fairness, which in practice will prevent such distributions, can themselves be justified in practice through their utility-value. But the fact remains that, if in a real-life case, Drug B were funded instead of Drug A, this,

according to the utilitarian position, would have been the objectively right policy. I think that is enough to throw utilitarianism into serious doubt. The view has a great deal of prima facie plausibility: surely the right and rational thing to do is to bring about the best state of affairs? But when we reflect upon the fact that separate individuals each have their own lives to lead, and that one person whose life is going badly can't be compensated by another's going well, we are likely to think that how well-being is shared around matters, ethically, in itself.

2. Egalitarianism

It is tempting to think that what is so objectionable about prescribing Drug B in the above example is that it exacerbates an already huge inequality between people, and hence that the value that speaks in favour of Drug A is equality. So perhaps we should, instead of cost-effectiveness, allow an egalitarian principle to govern our distributions of health care. But now consider the following case:

	Group 1 (10 people)	*Group 2 (200 people)*
Status Quo	5	90
Drug A	50	90
Drug B	5	92.5
Wastage	3	3

In *Wastage*, the government withdraws money from the health care budget and wastes it – by burning it. The pollution caused has an especially large effect on Group 2, but it also brings down the well-being of those in Group 1. According to Egalitarianism, however, there is something to be said for *Wastage*: it decreases inequality, and indeed creates perfect equality.

Now egalitarians will be quick to point out that they accept other, non-egalitarian principles, which would in fact override the equality principle in cases such as this. But I suggest that any principle according to which there is anything to be said in favour of Wastage is even more counter-intuitive than utilitarianism. At least some people benefit from the decision to fund Drug B. In Wastage, no one gains and everyone loses.

Before we reject Egalitarianism, however, we should try to understand what lies behind the kind of objection I have just outlined (the so-called levelling-down objection). Larry Temkin has argued that the force of the objection rests on:

> *The Slogan:* One situation *cannot* be worse (or better) than another *in any respect* if there is *no one* for whom it *is* worse (or better) *in any respect*.[2]

As Temkin interprets the slogan, 'if there is no one' may be taken as equivalent to 'if no one exists or will exist', and 'is worse' as 'is or will be worse'. As Temkin points out, a serious problem with such a view is that it cannot resolve Derek Parfit's *Non-identity Problem*.[3] Many of our actions will affect the identity of those who are born in the future, because they affect which sperm will fertilize which eggs. Consider, for example, the conservation of resources. If we choose not to conserve, that is going to affect our lives greatly, and consequently the identity of any children that we have – and of course the identity of any children that they have. Thus, even if we make the quality of human life in future much lower than it might have been, future generations cannot object, since no particular person can be said to have been harmed.

Since the slogan cannot explain why it would be wrong not to conserve resources, it should be rejected. But I myself do not believe that it is the slogan that underlies most people's dissatisfaction with the implications of Egalitarianism in levelling down cases. The slogan involves person-affectingness -- the notion that what matters morally can be only what affects those who do or will exist. But what is worrying about Egalitarianism is independent of person-affectingness in this sense. Rather, the worry arises from the idea that what matters morally could be something that was independent of the welfare of individuals. What worries people about the egalitarian position on levelling down is the claim that there is a value in making everyone worse off, and no one better off (when those who are made worse off do not *deserve* what happens to them).

Why is it that people find equality appealing? One possibility is that Egalitarianism is being confused with Prioritarianism, and I will come to that in a minute. But this is certainly not the whole story. Many find equality attractive because they believe it is *fair* and inequality is *unfair*. So, it is claimed, there is at least something to be said against both the *Drug A* and *Drug B* outcomes – that in each of them some people do much worse than others, through no fault of their own.

[2] Larry Temkin, *Inequality*, 256. The slogan is essentially a 'person-affecting claim'; see Derek Parfit, *Reasons and Persons*, 32
[3] Parfit, 357-9.

Where does this intuition about fairness come from? I want to outline an evolutionary account along the lines of the account J.S. Mill offers of the origin of justice.[4] Mill argues that human beings have a natural tendency to sympathize, and that the combination of this with our natural tendency to defend ourselves results in the notion of justice. Now consider what might happen if the natural tendency to sympathize were combined with the natural tendency towards *envy*. Envy typically involves the desire that the envied good be removed from the envied person and anger at the very possession of the good in question. Generalized through sympathy, envy becomes anger at anyone's doing better than anyone else. Note that I am *not* claiming that contemporary appeals to fairness themselves involve envy. Nor am I claiming that my account of the origin of our views about fairness is immediately debunking of any attribution of value to fairness. But I do believe that this account at least throws some doubt on such attributions.

3. Prioritarianism

So far, we have rejected both cost-effectiveness, or utilitarian, and egalitarian approaches to the distribution of health resources. In recent years, largely thanks to Derek Parfit, a good deal of attention has been paid to the idea that what matters in distribution is giving priority to the worse off. Parfit states this view as follows:

> *The Priority view:* Benefiting people matters more the worse off these people are.

Prioritarianism puts no value on equality in itself. And, since those in Group 1 will do much worse without their drug than those in Group 2 will do without theirs, the correct policyin our cases, a prioritarian may argue, is to fund Drug A.

Parfit's statement of Prioritarianism is rather vague. How might we make it more precise? On one immediately tempting position, what matters here is our being able to justify what we do, individually and person-by-person, to the worst off. Any deviation from benefiting the very worst off would be more unjustifiable to them than not benefiting those better off would be to them. This gives us:

> *The Absolute Priority view:* When benefiting others, the worst off individual (or individuals) is (or are) to be given absolute priority over the better off.

[4] *Utilitarianism*, ch. 5.

Consider the following distributions, where *WP* is the worst-off person:

	WP	*Group 1 (1000 people)*	*Group 2 (1000 people)*
Status Quo	8.9	9.1	100
Absolute Priority	9	9.1	100
Expanded Concern	8.9	100	100

The Absolute Priority view, in this case, favours moving from *Status Quo* to *Absolute Priority* rather than *Expanded Concern*. Imagine that WP is in quite serious pain, and that in *Absolute Priority* Group 1 is in pain almost as serious. All that will happen in *Absolute Priority* is that WP will be given a chocolate (her pain is bad, but not so bad that she cannot enjoy a chocolate). The Absolute Priority view favours giving her the chocolate over alleviating the serious pain of a thousand others. Because the Absolute Priority view is an 'innumerate' maximin principle, it will, like Rawls's 'difference principle', allow the smallest benefit to the smallest number of worst off to trump any benefit, however large, to any but the worst off, even the next worst off. And this, it may be thought, is almost as absurd as levelling down.

What is required, then, is a principle that allows us to give priority to the worse off, but in giving priority to take into account the size of benefits at stake and the numbers of people who will benefit. So understood, the Priority view is essentially a non-lexical weighting principle:

> *The Weighted Priority view:* Benefiting people matters more the worse off those people are, the more of those people there are, and the greater the benefits in question.

The Weighted Priority view, then, will permit us to benefit those who are better off if the benefit to them is significantly greater than to the worse off, or if they are greater in number. Exactly how the factors of absolute position, size of benefit, and number of beneficiaries are to be weighted is of course an important question, but I now want to suggest that whatever weights are attached to these factors the Weighted Priority view allows too much weight. Consider the following proposal:

Improvement in level	**Weight**	**Overall value**
1-->2	100	100
2-->3	99	99
3-->4	98	98
.....		
98-->99	3	3
99-->100	2	2
100-->101	1	1

This table represents the moral weights attached to improvements in the position of any individual. If, for example, I can offer one unit of the good to be distributed to P, who already possesses one unit, two units to Q, who has two units, or three units to R, who has 98, the overall moral value of my so doing will be:

P (has 1, gets 1): 100
Q (has 2, gets 2): 99 + 98 = 197
R (has 98, gets 3): 3 + 2 + 1 = 6

Attention to the size of benefits, then, requires me to benefit Q in this case. But if we imagine that there is *another* person, P', in P's position to whom I may also give one unit when benefiting P, and that everything else remains the same, the overall value of benefiting P and P' will be:

P + P' (have 1 each, get one each): 100 + 100 = 200

In this case, the number of beneficiaries outweighs the importance of the amount of benefit available to any particular individual.

The Weighted Priority view will judge *Expanded Concern* clearly superior to *Absolute Priority*. But because it allows for straightforward aggregation across persons the following problem now arises, regardless of the weighting.[5] Consider the following situations, involving ten people doing pretty badly and fifteen thousand pretty well:

[5] Compare here Temkin's 'Repellent Conclusion' (Temkin, p. 218).

	Poor (10 people)	*Rich (15,000 people)*
Status Quo	1	98
Pain-relief	51	98
Chocolates	1	99

In *Pain-relief*, each poor person gains *fifty* units, and there are ten such people. The value of increasing the level of the poor in this outcome is thus:

(100 + 99 + 98 + 97 … + 53 + 52 + 51 = 3775) x 10 = 37,750

The value of giving a chocolate (a really *good* chocolate!) to the rich is:

3 x 15,000 = 45,000

In other words, the Weighted Priority view, though it may avoid requiring us to give the smallest benefits to the smallest number of worst off at the largest costs to the largest number of those only slightly better off, does require us to give tiny benefits to those who are very well off at huge costs to the worst off. Its readiness to aggregate straightforwardly 'all the way up' leads it to fail to attach the appropriate moral significance to size of benefits and numbers of recipients. This seems, if anything, an even less palatable position than the Absolute Priority view, since that view at least always skews distributions in favour of the worst off.

It will not solve the problem to allow the weighting to operate only at lower levels. For as long as the number of the rich is large enough, priority may be given to benefiting them to a small degree rather than benefiting the worse off to a large degree. The problem is arising from straightforward aggregation, so one possible solution here would be to decrease the weight attached to numbers of individuals. Consider:

> *The Number-Weighted Priority view:* Benefiting people matters more the worse off those people are, the more of those people there are, and the larger the benefits in question. But the number of beneficiaries matters less the better off they are.

It may be claimed, for example, that the importance of numbers asymptotically approaches zero as they become large. Thus, a weighting could easily be devised which

ensured that, in the *Pain-relief/Chocolates* case, when aggregating, one weighted the second rich person's contribution to the sum at somewhat less than three, the third at even less, and so on, in such a way that the total was lower than 37,750, thus respecting our intuition that we should relieve the pain.

But to count people's evaluative contribution to populations in this non-individualistic way strikes me as against the spirit of welfarism and the thought that from the moral point of view each person counts as much as any other. Further, I now want to suggest that all prioritarian strategies which allow prioritarian concern for all, however well off, are mistaken. So now consider what I shall call the *Beverly Hills case*, in which you can offer fine wine to different groups of well-off individuals:

	10 Rich	*10,000 Super-Rich*
Status Quo	80	90
Lafite 1982	82	90
Latour 1982	80	92

If we use the weighting method above, the value of giving Lafite 1982 to each of the ten Rich is:

(20 + 19) x 10 = 390

Now let me assume that a number-weighting has been devised, such that the value of giving the Latour 1982 to the ten thousand Super-rich comes out as less than 390 (how the figures might be calculated is not important for the purposes of this example). Once again, I suggest, a modification of the Priority view has taken us from one extreme to another -- from allowing that numbers count straightforwardly to denying them appropriate relevance. It seems somewhat absurd to think that the Rich should be given priority over the Super-rich. Indeed, what the Beverly Hills case brings out is that, once recipients are at a certain level, any prioritarian concern for them disappears entirely. This implies that any version of the Priority view must fail: when people reach a certain level, even if they are worse off than others, benefiting them does *not*, in itself, matter more. And this seems to me to be true even if, in a Beverly Hills case, the utilities are equal. That is, even if the benefits to each of the Rich and the Super-rich are identical, and their numbers are the same, there still seems to me

nothing to be said for giving priority to the 'worse off'. At this level, only utilities matter, so there would be nothing to choose between the two distributions.

4. Sufficientarianism

Cost-effectiveness or utilitarianism cannot avoid implications of injustice. Egalitarianism runs into doubts over the value of fairness when no one is benefited by it. And unrestricted Prioritarianism attributes importance to giving priority even when those concerned are extremely well off. So we must look elsewhere for an account of why we should select Drug A in my original example.

One useful heuristic for constructing distributive principles is via the idea of an 'impartial spectator. As Adam Smith puts it:

> We can never survey our own sentiments and motives, we can never form any judgment concerning them; unless we remove ourselves, as it were, from our own natural station, and endeavour to view them as at at certain distance from us ... We endeavour to examine our own conduct as we imagine any other fair and impartial spectator would examine it.[6]

One common way to use the idea of the impartial spectator is as an argument for utilitarianism. The spectator is assumed to be impartial in the sense of caring for each 'unit' of well-being without concern for its location (that is, for who 'owns' it), and the virtue ascribed to her is that of benevolence. But as we have seen such a spectator is blind to the significance of the fact that each person is, to use John Findlay's word, 'separate'. Well-being isn't the only thing that matters; it is also of great significance in whose life that well-being is found.

In the case of those who are very badly off, which virtues come into play in response? One obvious one is, of course, benevolence. But benevolence, construed impartially in the standard way, leads us to utilitarianism. Another significant virtue here is *compassion*. Compassion seems ordinarily to be a response to the *suffering* of another person. And, at some point, it tends to run out. Imagine that in the Beverly Hills case the wine goes to the Super-rich. Does one feel any *compassion* for the rich, who were already worse off and have now lost out further? I think not.

[6] *Theory of the Moral Sentiments* (ed. Raphael and Macfie), 110.

If we now ascribe the virtue of compassion (as well as benevolence, perhaps) to the impartial spectator, we will provide the foundation for a theory of justice which gives special concern to those who are badly off, but not to those above a certain threshold:

> *The Compassion Principle:* Absolute priority is to be given to benefits to those below the threshold at which compassion enters. Below the threshold, benefiting people matters more the worse off those people are, the more of those people there are, and the greater the size of the benefit in question. Above the threshold, or in cases concerning only trivial benefits below the threshold, no priority is to be given.

According to the Compassion Principle, justice is especially concerned with ensuring that individuals have *enough*. Enough what? According to the view I outlined above, enough well-being in their lives as a whole. That is why I prefer not to see *Sufficientarianism,* as it has come to be called, as a *needs-based* principle. Justice is not a matter of responding to needs, but rather to those whose lifetime well-being will otherwise be below a threshold. I think there are anyway two problems with needs as a foundation for theories of justice. The first is that the value of the satisfaction of needs is merely instrumental. You may need some drug to remain healthy. But satisfying that need will benefit you only if being healthy is itself going to enable you to advance your own well-being in some way – for, as I said above, health itself is not valuable in itself. The second problem is that – as we can see from reflection upon the Compassion Principle – needs give out before justice is satisfied. Imagine a society which includes, among a large number of very wealthy and flourishing individuals, a group which is very poor, but whose basic, and indeed non-basic, needs are met. Compassionate concern for the badly off speaks in favour of at least some transfers from the rich to the poor, even if the poor use any resources gained to purchase goods which they could not be said to need.

How, then, are we to place our threshold? Remember first that it will be a threshold of *lifetime well-being*. So we could not, for example, on the basis of a 'fair innings' argument, immediately set a threshold in terms of age. For some older people may have had much less valuable lives than some much younger people. Probabilities also need to be factored in, using standard cost-benefit analysis functions. Let's begin with a case which might plausibly be said to be a clear one:

	Patient 1	*Patient 2*
Status Quo	5	90
Drug A	10	90
Drug B	5	95

What do the numbers represent here? Imagine that patient 1 is young, with a low (but positive) quality of life. Drug A will, relatively speaking, improve her life massively, though the absolute increase is only 5 units. Patient 2 is much older – let's say 95 years old – and has had a pretty wonderful life so far. Drug B would give her the same absolute benefit as Patient 1.

It is not implausible to claim that, given real-world conditions of scarcity, Patient 2 is clearly above the threshold, whereas Patient 1 is well below. One obvious objection to this suggestion is its apparent parochialism. Imagine that human lifetimes were both very happy and usually much longer – averaging, say, 500 years. Then it is quite unlikely that we would consider as a fair innings a life of 95 years of life even at a high quality.

Now, though it is true that we probably would take this view, that is not to say that it is correct. One option here would indeed be to make the truth about thresholds context-sensitive in certain ways. So in a happy population with average lifetimes of 500 years, someone with 95 years of happy life would be below the threshold, whereas in our world they are above it. This seems to me unsatisfactory. If we are thinking about compassion as experienced in our world, we do have to see it as limited in various respects. There is a very great deal of very serious suffering, and this is a call on our attention which should drown out a lot of the claims of less serious suffering which might otherwise require our concern. But we should not imagine our impartial spectator in this way. She is God-like, able to apportion exactly the right degree of concern to each being, regardless of context. We are not gods, of course, so all we can do is make judgements about the threshold to the best of our abilities, in the light of the facts and of reflection upon possible worlds in which things are very different from in our world. There is a general hermeneutic problem here, since anyone's judgements about the threshold will be shaped to some extent by their upbringing and the world in which they live. But human reason has the ability to transcend its origin and circumstances to some degree, and that is what we must do here as in other areas of ethics.

But what should happen below the threshold? One important factor here is the correct account of well-being. It might be thought that, in general, the younger you are when you die, the worse it is for you. But on some views of well-being, it may be worse for you to die when

you are a late adolescent or in early adulthood than it is to die when you are very young. For by the time you die you will have begun certain projects, and the failure to carry them through is more harmful to you than never having the opportunity even to begin them. This view, however, strikes me as implausible. In general, the longer your life, and the more you have packed into it, the better for you. But there are differences between positions that allow that more life is, in general, better. On some 'ideal' or 'perfectionist' views, for example, your well-being is really going to start picking up when you accomplish something with your life, or acquire some reasonably high level of understanding of the world. So a year for an adult, say, may be worth more to that adult than a year to a child. I myself do not accept idealism, but, like any view of well-being, it can be plugged into Sufficientarianism to offer guidance on the justice of a distribution.

On the Sufficiency view, then, priority should be given in the distribution of both lifesaving and therapeutic health care to those below the threshold over those who are above, and – below the threshold – to those who are worse off. But numbers of beneficiaries also matter. The relevant factors are, once costs are known: well-being without treatment; well-being with treatment; the numbers that can be treated. Here is a clear case.

	Patient 1	*Patient 2*
Status Quo	55	60
Drug A	65	70

Here, if a choice has to be made, it seems clear that Patient 1 should be provided with the drug.

Here is a harder case:

	Patient 1	*Patient 2*
Status Quo	40	45
Drug A	50	65

Here the greater size of the benefit to Patient 2 has to be weighed against concern for Patient 1's lower well-being. The 'efficiency' component of Sufficientarianism is in tension with its 'individualist' component.

And here is another harder case:

	Patient 1	*Patients 2-1*
Status Quo	40	50
Drug A	55	60

Here, the size of the benefit to Patient 1 is greater than that to any of Patients 2-11. But the overall gain in well-being is five times as great if Drug A is administered to Patients 2-11. Again, we have a conflict between efficiency and concern for individuals, to be resolved through careful reflection and judgement.

Should we see the Sufficiency principle in health care as merely an application of a more general principle of application, or are there what Michael Walzer called 'separate spheres', so that distributions can be assessed sphere by sphere and in the light of principles tailored to suit each sphere?

As I said above, my inclination is to see ethics as ultimately concerned with abstract, general principles. So though we cannot rule out separate spheres, we should not introduce them without good reason. Further, it strikes me that there is at least one good reason for not introducing them. We have seen that the compassion principle is concerned with how well individuals' lives go for them. Other things being equal, it does not matter to any individual whether her life goes worse for her because of poor health, poor income, lack of citizenship rights, or whatever. The upshot for the distribution of health resources is that priority in health care may be given to certain individuals whose lifetime well-being is lower than certain other individuals for non-health-related reasons.

Because of Sufficientarianism's generality, it also covers our relationship with non-humans. So one implication of it may be that we have to give priority to perfectly contented non-humans over possibly quite poorly off or sick humans, because the lifetime well-being of the non-humans is so low. I myself am prepared to bite this bullet. But it should be noted that the implication can be blocked if a good argument can be provided for giving priority to humans over non-humans.

Introducing thresholds is often philosophically dangerous. It may lead to your having to accept, for example, that the addition of one tiny grain to a large number of grains brings a heap into existence. In the case of Sufficientarianism, the problem could be seen as follows: Consider the following scenario, where *WP* is again the worst off person:

	WP	*Group 1*
Status Quo	22	26
Below	24	26
Above	22	100

If we assume that the threshold is 25, we can see that the view will prefer the smallest benefit to any number of individuals below the threshold to any benefit, no matter how large, to any number of individuals above the threshold. That is, it will view *Below* as superior to *Above*. This implication of the view, however, may not be as implausible as it seems, once we give proper recognition to the fact that the threshold is the point at which compassion no longer applies. There really is something special to be said for benefiting the worst off individual which cannot be said for benefiting those above the threshold. It may seem that, at some point, sheer quantity of well-being or number of beneficiaries must always trump benefits to sub-threshold individuals. This strikes me, however, as doubtful. Let me adapt a case of Temkin's. Imagine that you have a choice between alleviating someone from the most severe agony imaginable, which will otherwise go on for another fifty years, or giving billions of individuals a very small pleasure (such as a lick of a lollipop). It is not implausible to suggest that *no* number of lollipop licks could provide a reason strong enough to outweigh the reason to alleviate the suffering.

This point may be further illustrated by considering how a Sufficiency principle might emerge from a contractualist set-up such as that of Rawls. In effect, such forms of contractualism extend principles of intra-personal rationality to groups of individuals. The interpersonal distributive principles that emerge depend on the conception of intra-personal rationality used in the original set-up. So Rawls's maximin principle, according to which rationality requires one to make one's worst possible outcome as good as possible, delivers at the interpersonal level the difference principle, according to which institutions should aim to make the position of the worst off as good as possible, at the interpersonal level. It seems to me that many people would accept an intrapersonal threshold, so that an option which involves the possibility of fifty years of agony could not be made preferable to another option with no chance of such agony *however* large the gains in well-being that might be made through choosing the first option.

It is not my brief to apply the Sufficiency view to our three real-life cases. But let me conclude with four general practical principles for the distribution of health resources that might, on certain plausible assumptions, emerge from Sufficientarianism:

• Give absolute priority to those under the threshold at which compassion (impartially understood) gives out to those above.

• Choose treatments which are as effective as possible at boosting well-being.

• Target the worst off in particular, especially in lifesaving treatments for young patients who are likely to go on to lead a full and healthy life.

• Minimize costs.

Often, there will not be a strong case for expensive, lifesaving treatment for the elderly; whereas there will be a strong case for cheap, beneficial interventions aimed especially at younger people – such as e.g. health promotion education, or certain types of treatment for mental illness.

Roger Crisp April 2009

St Anne's College
Oxford Uehiro Centre for Practical Ethics
University of Oxford