

slide 1

# Aggregation in Health Resource Allocation

Weyma Lübbe  
Regensburg University/Germany

„Ethical Issues in Prioritization of Health Resources“  
April 23-24, 2009, Harvard University, Boston

1. Numbers and sizes
2. Weighing goodness and fairness?
3. Fairness trumps goodness
4. Redefining the units of goodness?

## 1. Numbers and sizes

If you want to drink much beer, there are two possibilities.

slide 2



Either you drink many glasses of beer. This is what people in my country do if they live in the city of Cologne. Or you drink big glasses of beer. This is what people do if they live in Munich. Many people, however, will say that numbers and sizes do not matter as long as we get the maximal overall amount of value for our money.

Some think that this holds for purchasing health in public health systems as well as for purchasing beer. Beer and health are certainly different goods. Still, they are both goods, and under conditions of scarcity, more of what is good seems better than less. The most obvious difference, indeed, has not to do with the goods in question but with their carriers: glasses here, persons there. The separateness of glasses is a matter of expedient transportation, the separateness of persons is not. Persons are not just locations of good. They are what goods are good for.

But what, exactly, follows from this difference? The usual answer is that “distribution” becomes a relevant issue: who gets what amount. In the debate on prioritization of health care resources, a common strategy of theoreticians is to treat this issue, under the

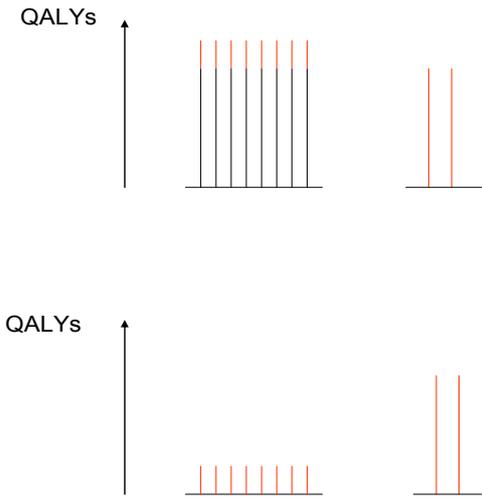
name of fairness, as a second ethical objective besides the objective of maximizing overall benefits (which some call the goodness objective). The debate concentrates on how exactly the fairness objective is to be spelled out and how the two objectives are to be combined.

I, for my part, agree that the separateness of persons is a good reason for accepting fair distribution as an ethical objective. But I believe it is also a good reason for accepting that maximizing overall benefit is not an ethical objective. I can think of no ethical basis for additive utility aggregation if such aggregation is not fair. Claims of justice, in my view, must not be weighed with other objectives anyway. They must be redefined, not weighed, if there actually are legitimate objectives that are not compatible with satisfying them.

In what comes, I will argue for this position. I am aware, of course, of the special task I have been set: to offer theoretical guidance on the question whether (or to what extent or under what circumstances) small size benefits for large numbers of patients may be preferable to large size benefits for small numbers of patients. I do not think, however, that this issue can usefully be discussed without drawing in the central topics of the other sections: priority to the worst off and fair chances.

Let me, to begin with, present some types of examples. Often, big size benefits go to patients who are worse off than patients who get small size benefits. This is because a big health benefit presupposes a big health problem in the first place. But patients who get small benefits can have big health problems, too. So, if being worse off is a relevant consideration, we should distinguish between the following two types of our decision problem, at least:

**slide 3**



The lines stand for health expectations of individuals, measured in QALYs. The red parts stand for benefits that may be added through medical treatment. Treating the many patients on the left is here and in all following cases meant to be equivalent in

costs to treating the few on the right. Treating the many or treating the few are thus alternative ways of spending a limited budget or part of a budget.

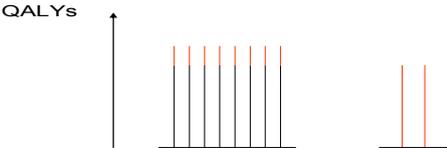
In these two examples, the overall benefit happens to be the same if small sizes go to the many than if big sizes go to the few. In the second case, all patients will die without treatment. When you asked me, however, to comment on

**slide 4**

“small benefits to a large number vs. life saving interventions to a few”

you obviously had not a case on your mind where the small benefits are life-saving too. So let us look in more detail at the first case.

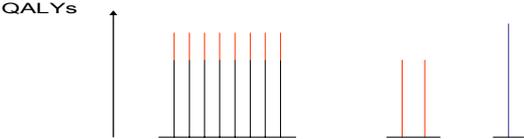
**slide 5**



Since the goodness objective leaves us free to choose whom we treat, we may concentrate on the fairness objective. It seems to tell us that we should treat the few. Not only are they worse off when left untreated; it is also the case that everybody will be equally well off if we spend the resources this way.

Now let us look at a case where goodness and fairness point in different directions.

**slide 6**



The benefits to the many are still considerably smaller than the benefits to the few. But together, they amount to more. The blue line at the right shows how long the two red lines would have to be in order to equal the aggregated benefits of the few. How should goodness and fairness be combined in such nontrivial cases?

There are several possibilities, notably: trumping, weighing, and redefining the units of goodness.

#### slide 7

1. trumping
  - goodness trumps fairness
  - fairness trumps goodness
2. weighing goodness and fairness
3. redefining the units of goodness

I will explain what exactly I mean by the third item. In the prioritization debate, trumping rules are hardly ever proposed by now. But in order to apply even this simple sort of a combination rule, we would need to spell out the fairness objective in a way that allows us to rank whatever alternatives we have. This is not a trivial matter. In commenting on my examples, I mentioned already two different criteria or aspects of fairness: priority to the worst off and equality after allocation. Equal chances is a further fairness issue.

#### slide 8

Fairness:

- Give priority to the worst off
- Realize equality after treatment
- Provide equal chances

For any two of these criteria, cases can be construed where they point in different directions. So perhaps we need already a combination rule within the fairness objective. I will come back to this, but let me comment on weighing goodness and fairness, first.

## 2. Weighing goodness and fairness?

Here is an example for a combined goodness-and-fairness measure where the combination is done by weighing:

#### slide 9

$$w_1 + w_2 - \frac{1}{2} |w_1 - w_2|$$

*J. Broome (2002), Fairness, goodness and levelling down* \*

---

\* The quoted source (in: C. Murray et al., eds., Summary measures of population health, Geneva) consistently shows “2” for the weighing factor. But this is an erratum, as the author has confirmed.

The  $w$ 's stand for the well-being of persons 1 and 2 respectively. The part at the left side of the minus sign is the goodness part, the part at the right is the fairness part. It measures equality of wellbeing. Such a combined measure allows to choose the *overall best* alternative to allocate a resource, the value of fairness included, by maximizing the value of the formula. Such formulas can of course easily be extended to many-person cases.

When presenting this measure, the author, John Broome, consciously underrepresented the complexity of the theory of fairness that he is actually advocating himself. According to Broome's own theory, fairness is not a matter of how equal everybody is off after treatment, nor is it a matter of the other criteria mentioned before. It is a matter of equal satisfaction of equal claims.

#### slide 10

Fairness:  
Satisfy equal claims equally (J. Broome)

In order to take account of the fairness objective, we would then have to ask in the first place what claims are at issue in a case. Discussing cases in the way I did above – by picturing possible distributions of benefits – would not be enough because such pictures contain no information about claims.

In the prioritization debate, authors with a primary background in welfare economics do not usually ask for such information. They speak of benefits, values, utilities, and they do speak of equal or unequal distribution. But they hardly ever speak about claims. So why did Broome – a quarter of century ago, notably, when he was still an economist writing for economists –

#### slide 11

J. Broome (1984), Uncertainty and fairness, *The Economic Journal*

why did he feel the necessity to make use of the concept of claims?

I do not ask this question for historical reasons. I think the answer teaches a point about how fairness and goodness relate to each other. In the article I refer to, Broome was not motivated by an economist's wish, fallen out of the blue, to develop a theory of fairness. He was exclusively motivated by the wish to defend John Harsanyi's proof of additive utility aggregation

#### slide 12

J. Harsanyi (1955), Cardinal welfare, individualistic ethics, and interpersonal comparisons of utility, *Journal of Political Economy*

## against a fairness-related argument by Peter Diamond

### slide 13

J. Harsanyi (1955), Cardinal welfare, individualistic ethics, and interpersonal comparisons of utility, *Journal of Political Economy*

P. Diamond (1967), Cardinal welfare, individualistic ethics, and interpersonal comparisons of utility: Comment, *Journal of Political Economy*

– an argument that would be detrimental to the proof if it was sound.

A Harsanyi-style proof of additive utility aggregation was and is the basis of Broome's theory for the goodness objective. It rests on premises of seemingly trivial normative content – notably on the Pareto Principle, the idea that an option is to be preferred to another option when it is better for at least one person while not being worse for anybody. The proof secures that every unit of benefit that is produced with limited resources counts morally, and that more units count more, at whatever location they show up. Broome saw, however, that there was no adequate answer to Diamond's objection within the conceptual framework of this proof, the framework of welfare economics (more specifically: of that part of welfare economics that allows for interpersonal comparison of utilities). It was in order to save this framework that Broome introduced claims as a second sort of reasons for acting, besides utilities, and postulated that these reasons have to be dealt with separately – which makes goodness and fairness separate ethical objectives, to be weighed with each other. As I said before, I believe that goodness and fairness cannot be dealt with separately. Looking at Diamond's argument helps to understand why.

Here is the decision problem Diamond was referring to:

### slide 14

	state 1 (heads)	state 2 (tails)
A	David 100, Peter 0	David 100, Peter 0
B	David 100, Peter 0	David 0, Peter 100

A and B are alternative ways of allocating a life-saving medical resource. There is only enough of it to save either David or Peter. A, the first option, consists in saving David right away; Peter is left to die. B, the second option, consists in making a random choice with equal chances for both patients (heads: David will be saved; tails: Peter will be saved). Diamond argued, plausibly enough, that it would be better (or, as economists usually put it, “socially preferable”) to choose option B because, giving equal chances, it would be fairer while not producing less utilities. But Harsanyi's proof of additive utility aggregation rests on an axiom that is incompatible with saying that option B is preferable because it gives equal chances. The relevant axiom is called the Sure Thing Principle. It says that the comparative value of two prospects (like A and B) can be

assessed by comparing separately what happens if heads fall and what happens if tails fall.

If we do this comparison (state-by-state comparison it is called), we will see its incompatibility with Diamond's contention that B is the better option. First, look at state 1. David will be saved in option A, as well as in option B. The outcomes are identical and thus, of course, of equal value. Now look at state 2. David will be saved in option A, Peter will be saved in option B. Since we neither realize more goodness nor more fairness by saving Peter or David, these two outcomes are of equal value as well. So, if state-by-state comparison is a sound procedure, we must conclude that option A and B are equally valuable. The Sure Thing Principle says that state-by-state comparison is a sound procedure. But the value of providing equal chances cannot be seen under state-by-state-comparison. The fact that David and Peter get equal chances is a feature of option B that spreads over both states, heads and tails. The Sure Thing Principle does not allow that such a feature contributes to the value of an option.

Other welfare economists reacted to Diamond's argument by denying that option B is preferable. Broome, however, thought that Diamond had a point. In certain cases, providing equal chances might indeed be fair. So Broome tried to find a way to explain how B can be the better choice without violating the axiomatic basis of additive utility aggregation. It had to be an explanation that makes the higher fairness value of option B visible within the single states, not across states. After arguing in some detail (plausibly, in my view) why the fairness value of option B cannot be made visible by introducing additional individual utilities, Broome presented the following solution:

**slide 15**

	state 1 (heads)	state 2 (tails)
A	David 100, Peter 0, the selection not having been fair	David 100, Peter 0, the selection not having been fair
B	David 100, Peter 0, the selection having been fair	David 0, Peter 100, the selection having been fair

He simply added the information that option B includes fair selection, while option A does not, within the descriptions of the single outcomes.

As you will probably feel yourself, this is not an impressive solution. On the face of it, it saves the Sure Thing Principle all right because the different fairness value of the two options is now visible within the single states. But if equal chances still are what makes us say that the selection has been fair in option B, this would simply be a redescription trick. Axiomatic frameworks are not to be deceived by such tricks. When they ask for separable outcomes – outcomes the value whereof is independent of what happens in other states – we must fill in outcomes that really are separable. So the question is whether the fairness information that is written within the single states can be explained without reference to what happens in both states, yes or no.

Broome tried to come to a yes by developing a theory of fairness that does not refer to utilities at all, but only to claims – these being allegedly a separate sort of reasons for

action. Fairness, according to Broome's theory, means (as we have heard) to satisfy equal claims equally. So what claims are at issue in our case?

Whenever Broome has written about the fairness of random selection, he supposed that the claims at issue are equal claims to be saved. The fairness of option B would then have to be explained in the following way:

**slide 16**

	state 1 (heads)	state 2 (tails)
A	David 100, Peter 0, equal claims to be saved not satisfied equally	David 100, Peter 0, equal claims to be saved not satisfied equally
B	David 100, Peter 0, equal claims to be saved satisfied equally	David 0, Peter 100, equal claims to be saved satisfied equally

B would be fairer because within each of its outcomes the two patients' equal claims to be saved would be satisfied equally. But they are not, of course. One candidate dies and the other survives. This invariably happens in all outcomes, in option B just as well as in option A. If fairness means that we satisfy equal claims equally, and if option B is to be fairer than option A, there is only one way to represent the case:

**slide 17**

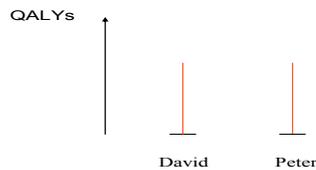
	state 1 (heads)	state 2 (tails)
A	David 100, Peter 0, equal claims to <i>chances</i> to be saved not satisfied equally	David 100, Peter 0, equal claims to <i>chances</i> to be saved not satisfied equally
B	David 100, Peter 0, equal claims to <i>chances</i> to be saved satisfied equally	David 0, Peter 100, equal claims to <i>chances</i> to be saved satisfied equally

The only claims that option B satisfies equally while option A does not are equal claims to chances to be saved. This redescription, however, includes a reference to chances, and in order to verify this reference, one must look at both states, heads and tails – which we cannot do under state-by-state comparison. This is the reason why claims to chances to be saved are never to be found in Broome's writings.

I conclude that the answer to the central question is no: The fairness value of option B cannot be explained without violating the Sure Thing Principle. The axiomatic basis of additive utility aggregation and the acceptance of random choice as a fair procedure are incompatible.

If this is true (we should probably discuss whether it is), what would be the consequences? Diamond's case is a very simple case. Pictured in our earlier notation, it looks like this:

slide 18



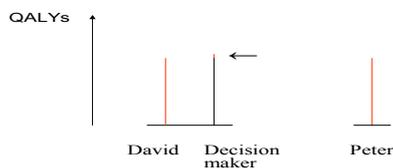
Individuals do not differ in need, nor do they differ in gain, nor do the alternatives differ in numbers of people saved. So this aspect of fairness seems to have a rather limited scope of application. Does it tell us anything at all about the more complicated sort of cases we started with?

I think it does. It does because it teaches some general points about claims, and about their relation to benefits. Let me explain this in my next section.

### 3. Fairness trumps goodness

How would the case of Peter and David have to be decided if information about benefits and their more or less equal distribution was all the decision maker needed to make a proper decision? Well – the decision maker would be free to choose. She would not produce more benefits in choosing either way, nor would she realize a more equal distribution. Why not, then, choose the patient that seems more agreeable to her personally? David, for instance. If the decision maker wrote David's agreeableness to herself down as an extra bit of benefit, she could even justify her choice in pointing out that it makes a difference under the goodness objective.

slide 19



Justifications like these constitute an embarrassment for the theories that support them. Broome thinks so, too. The intuitive judgement is that a choice between two patient's lives should not depend on such bits of additional benefit, be it to the decision maker, to the patients themselves, to their inconsolable surviving dogs, or to whomsoever. A theory that recognizes claims can account for this intuition – *if* claims are understood properly. Notably, conflicting claims are not to be handled in the way the utilitarian tradition handles conflicting benefits: by weighing. The conflicting claims in our case

are David's claim to survival and Peter's claim to survival – equal claims, by assumption. If it was appropriate to solve the conflict by weighing, the weights would be equal – just like the weights of the patients' benefits. We would again be free to choose and any additional bit of a benefit would tip the balance.

This is why claims are not to be weighed. Broome agrees. He opts for weighing benefits within the goodness objective, and for weighing goodness with fairness in the end, but not for weighing claims with claims within the fairness objective. Weighing conflicting items against each other aims at maximizing. Broome's fairness definition, however, is not "maximize the satisfaction of claims". It is "satisfy equal claims equally".

When commenting on Diamond's case in the last section, we did not solve the conflict by weighing Peter's and David's claims to survival. Going beyond Broome's analysis, we solved it by redefining the claims such that they do not conflict anymore. We said that a fair procedure in such a case must acknowledge equal claims to a chance to be saved, not equal claims to be saved. The decision maker is then not left with a tie that sets her free to act as the whim takes her. She is left with a clear duty how to proceed. And when the random choice is made, claims and duties are clear again: The winner has a claim to be saved, and there is no conflicting consideration left – no conflicting claim, *and no benefits that would have to be counted*. The foregone benefit of the other patient's survival is not to be counted separately, of course. It has already been adequately considered in acknowledging each patient's claim to an equal chance.

Generally speaking, this amounts to a consequence that Broome would never draw because he never thought of leaving the utilitarian tradition, only of supplementing it:

**slide 20**

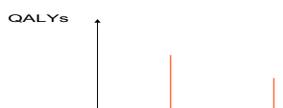
Reasons for acting:

- not: Benefits that we can produce
- but: Benefits that we can produce without violating claims

*Benefits that we can produce* by allocating scarce resources do not as such constitute reasons for acting. Only *benefits that we can produce without violating claims* constitute reasons for acting.

If this is a general point to be learned from Diamond's example (again, we should discuss it if you disagree), what can we say about the following sort of decision problem?

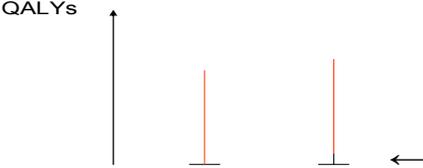
**slide 21**



The application of what I said so far seems to be this: Watch out for claims; do not count benefits separately; do not handle conflicting claims by weighing. Okay – so what claims are at issue here? When possible gains differ?

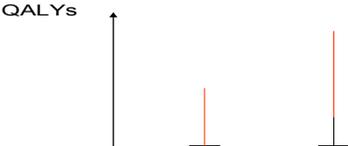
Or here:

slide 22



When one patient is worse off before treatment– but only very slightly. Or here:

slide 23



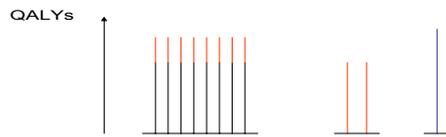
When needs *and* gains differ. Or here:

slide 24



When only numbers differ. Or here

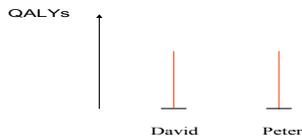
slide 25



(one of the cases we started with): Needs *and* gains *and* numbers differ. What, in all these cases, are the claims?

The answer is: We do not know. As I said earlier, pictures that show distributions of possible benefits do not contain information about claims. For Diamond's example,

slide 26



this was of course true as well. When discussing it, Broome *assumed* for the sake of the argument that the two patients have equal claims to survival. He did not derive this from the utility informations presented by Diamond. Indeed, Broome mentions explicitly that the same picture would correctly describe the situation after the random choice has been made. As long as nobody is yet treated, possible gains are still the same. But claims are not. Obviously, claims are not a function of individual's possible gains.

This is another point that may be learned from Diamond's example: In watching out for claims, we may have to gather information about the history of a decision problem. History, even if completely contingent (as is the outcome of a lottery), can be a legitimate source in defining claims. This makes further analysis somewhat fuzzy. What I just termed "history" can include complex institutional environments: insurance contracts, established social security programs, constitutional citizen's rights, charters against the discrimination of the disabled, and so on.

Perhaps health economists have been wise in never even attempting a theory of claims. Still, more and more of them have realized that they will not have much of a say in the politics of prioritization if they simply continue to advocate QALY maximization. So

they look for ways of taking account of the fairness issue that are more congenial to their traditional conceptual frameworks. One of these ways is to propose redefinitions on the goodness side.

#### 4. Redefining the units of goodness

As I mentioned earlier, the traditional economic way of speaking about the goodness side is not “goodness” but “preference”. Not goods are aggregated but utilities – these being numerical representations of individual preference relations. If the public debate suggests that citizens, as far as their health system is concerned, have preferences not only for as long and as healthy lives as possible but also for being treated fairly and having their fellow-citizens treated fairly, utilities should be redefined accordingly. The true economic framework, after all, is liberal. Why should it not fit for people who value fairness?

For instance,

slide 27



when people think (this has to be empirically tested, of course) that it would not be fair to treat the many in such a case, although it would realize more QALYs, people do obviously not value a unit of health benefit like any other unit of health benefit. The value of a QALY at the upper end of the QALY scale might, in the eyes of the citizens, be smaller than the value of a QALY at the lower, life-saving end of the scale. If we hypothesize this, citizens’ preferences can be seen to be perfectly rational even if they do not maximize QALYs. In opting for the treatment of the few, people opt for the outcome that maximizes what they actually value in resource allocation: not units of health, but the “societal value” of such units – as this author calls it:

slide 28

“[A] health care insurance scheme, be it public or private, should try to be as valuable as possible to its members. [...] This is the same as saying that it should give priority to activities that have a favorable ratio between benefits and costs. In that way it will maximize the membership value, hereafter called the *societal value*, of health care.”

*E. Nord (1999), Cost-value analysis in health care, p xix*

So the real measure of goodness is the societal value of health benefits. And as it happens, it varies along the QALY scale. And what is best, nobody has to bother about a theory of claims under this approach. This part of ethical theory can be replaced by empirical evidence:

**slide 29**

“I define a *fair* resource allocation in health care as one that accords with societal feelings about the strength of claims of different patient groups [...]. A resource allocation that violates such feelings is defined as *unfair*.”

*E. Nord (1999), Cost-value analysis in health care, p 23*

Very well. So what would the author hypothesize about societal feelings in a case that resembles our David-and-Peter case?

**slide 30**



If, as Nord assumes, citizens watch out for outcomes that maximize societal value, why should they dream of giving equal chances to groups that have equally strong claims, in their eyes? Equal chances are not outcomes, and providing them does not maximize anything.

One might object that this sort of decision problem is improbable to show up at the level of program evaluation, which is obviously the level the author has in mind when he compares patient “groups”. On this level, we might always manage to stay within our budget by drawing a line that separates groups with claims of different strength. But imagine that the patients in group A have a slightly less costly disease, such that we could satisfy the claims of a third group in addition

**slide 31**



– weak claims, if we keep in mind the value compression at the upper end of the QALY scale. Still, societal value maximizers must go for groups A & C. The value of satisfying the claims of group B has been outweighed.

Such things happen if the fairness objective is integrated into the welfare economic framework by redefining the units of goodness, while leaving the modus of aggregation unchanged. The essential point with such theories is that they do not part with the old utilitarian idea that there actually *are* units of goodness. As long as you stick to this idea, you will naturally imagine that the task of resource allocation is a maximizing task. Nord, like other economists who have written against QALY maximization, stays within a framework that takes at least *this* piece of the utilitarian heritage to be safe from criticism:

**slide 32**

“After all, we do want to do as much good as possible with the limited resources available for health care, don’t we?”

*Erik Nord (1999), Cost-Value Analysis in Health Care, p. xv*

What if we don’t? What if citizens are non-consequentialists?

I, for one, have argued that resource allocation should satisfy claims, not maximize value. The definition of claims, true, is a difficult task. I do not believe that it can be carried to its very end by however an astute philosophical analysis. Concrete claims – claims that are defined specifically enough to correlate with identifiable duties in identifiable situations of choice – are usually defined by political institutions, not by philosophers. Philosophers, however, can give political advice. One advice I like to give to health care decision makers when I am asked (not the only one, of course, but an important one, in my view) is this: Whenever experts begin to talk to you about better outcomes, ask back: better for whom, please? And if the answer is “better for society”, or even (which can happen if the expert is a philosopher) “better for the Universe”, ask back again: Who exactly is that, please?