

## PROSPECTS OR OUTCOMES?

Eric Rakowski

University of California, Berkeley, School of Law

Fifteen years ago, Norm Daniels suggested that the answers to four moral questions bearing on the distribution of scarce medical resources were persistently elusive, at least for those who shared widely held moral principles and who responded intuitively to specific cases involving trade-offs or priorities in what he considered, I think correctly, a common way.<sup>i</sup> A number of philosophers, not to mention a Berkeley tax professor, offered views on one or more of those issues at the time, and more have written on them since. I suppose that one implication of this conference is that, for all the light cast on these moral dilemmas over the intervening years, we as a group of informed, normatively sophisticated thinkers are not much closer to finding a consensus solution than we were then.

The realization that some of these problems bearing on a significant range of health care choices may be intractable is the prime impetus for the argument that Norm Daniels and James Sabin offered a few years back in their helpful book, *Setting Limits Fairly*.<sup>ii</sup> When agreement cannot be reached on a matter of moral substance, the argument goes, it is particularly important to find a *fair method* for making policy choices, whether at the governmental level or within a large health plan. Agreement on procedures might permit us, as a group, to get around our disagreements on substance. Among lawyers and political theorists, this a familiar point and strategy. And there is no doubting that deciding how we should, in the real world, in real time, settle matters of health resource allocation is an important and pressing question, one that implicates a

range of issues that are difficult in themselves, such as the correct account of distributive justice with respect to material resources generally and the scope and justification of state paternalism.

One point to notice for our discussion today, though, is that both the push to devise fair decisionmaking procedures and their responsible implementation do not render the difficult moral problems Daniels sketched fifteen years ago irrelevant, or for that matter less salient. Consider a parallel matter of institutional design, such as the constitution of a democratic polity. Thinking about what the contours of a just political order should be, we might conclude that it is essential to establish a set of substantive rights, such as privacy rights or rights to free expression or religious practice or even certain types of property rights, that are immune from the ordinary processes of political limitation or revision, if not completely beyond alteration through extraordinary mechanisms. We would do so not in virtue of what we consider the best account of fairness in collective choice, but because of substantive moral commitments that constrain our choice of procedural rules. Likewise, when we turn to systems for allocating health care resources, we might well believe that there are analogues to these rights – rights that cannot be abridged or that should be susceptible to modification only slowly and with difficulty – such as rights to various types of priority for the young over the old, or for those who are worse off through no fault of their own. We therefore cannot entirely escape hard moral choices in devising fair procedures. Nor can we ignore the obvious fact that no matter how fair a set of processes for democratic political decisionmaking might be, those processes can, in specific instances or over a run of decisions across time, yield unjust outcomes, which we need substantive moral principles to identify and if possible remedy.

There is a separate, overarching point that bears notice as well. No procedures for

apportioning therapeutic resources or research funding are likely to seem morally authoritative to us if they aggregate the *unreflective* preferences of citizens or members of a health plan. Insofar as we rely on some method for counting heads or intensity-weighted preferences, we will want the convictions or desires we take into account to be the outcome, in each individual's case, of engagement with the relevant moral considerations for favoring one person or group over another.<sup>iii</sup> And that means that we or anyone else whose views we think entitled to respect cannot evade the tough moral choices this conference is devoted to, even if we are fortunate enough to agree on a formula for making group decisions. As best we can, we need to answer the persistently difficult moral questions we have been discussing, both at the level of procedural design – in determining whether some types of priority should be built into the allocation system – and at the level of individual belief or preference formation, to the extent that our favored decisionmaking procedure is democratic rather than dictatorial.

### **Daniels's Statement of the Problem**

Norm Daniels stated the problem I have been asked to address in two forms – a case of what he calls “micro-rationing” and a case regarding the allocation of health care resources across larger groups.<sup>iv</sup>

His micro-rationing case involves the assignment of a scarce transplant organ. Two people are identical in every relevant respect except one: if Sarah is given the transplant, she is likely (as best we can estimate) to live another 20 years, whereas June is likely to survive only another 2 years with the transplant – a significant period of time, but a lot less than 20 years. Should Sarah be given the organ because she probably would live many more years with it than June would? Or does treating them as equals require

that we flip a coin, giving each the same chance to continue her life for a substantial span of time? Or should we perhaps find an intermediate solution, selecting the organ recipient randomly but weighting the odds in proportion to the potential recipients' relative life expectancies, so that June's chance of getting the organ becomes 1 in 10 while Sarah's would be 9 in 10?<sup>v</sup>

Daniels suggests that the vast majority of people believe that how long each person's life can be extended does matter in deciding whom to help, at least some of the time. If the differences in survival expectancy are very great, he says, few people would think it appropriate to flip a coin to decide who lives. Nonetheless, he adds, most people do not think that one person's likelihood of living longer than another is necessarily decisive. If the difference is slight, most are inclined to think, then we ought to accord each person a 1-in-2 chance of getting the organ. So neither extreme view seems right: we do not believe that we automatically should favor the person who is likely to live longer, but we also do not think that we always should flip a coin, because large disparities in recipients' expected futures might compel us to favor the person with the longer probable future. As to weighted-odds methods of choosing, Daniels considers them "ad hoc, adding an element of precision our intuitions lack."<sup>vi</sup> Nor will it do, he says, to locate the decisive random event in the so-called natural lottery that gave the two different survival prospects. We could hardly expect June to acquiesce in our use of a decision mechanism she never agreed to. Indeed, if the relevant event preceded her birth or the time when she could make mature, rational choices, it would be a decision procedure she could not conceivably have agreed to.

At the macro-level, Daniels invites us to consider the introduction of one of two possible medical treatments. Suppose that our health care budget allows us to pay for one but not both. And suppose that both treatments would benefit the same number of

different people who are identical in all relevant respects. The first treatment restores patients to a higher level of functioning than the second one does. Is that a persuasive moral reason to choose it in preference to the second? Or does treating patients as equals require us to flip a coin in deciding which treatment to pay for?

### **The Plan of Argument**

In my view, Daniels's fair chances/best outcomes problem is not a difficulty at all for large health insurance programs, including government-run programs such as Medicare. We should always favor the best outcome. So, let me begin by explaining what might seem a blinkered or unduly dismissive view, given the intuitive resonance of the micro- and macro-level cases that Daniels describes. My conclusion might seem all the more surprising, given my own qualified endorsement in a couple of publications of John Taurek's convictions with respect to choices between different-sized groups of people in life-or-death cases. So it will be important for me to elaborate on the significance of the special context in which these cases arise within large health insurance programs. Part of this explanation necessarily involves describing the significance of decisionmaking under uncertainty and how we ought to think about the moral implications of choices among alternatives with varying combinations of risks and rewards, both at the individual level and at the collective level. Finally, I will need to say several things about the limited reach of my rigid response to the fair chances/best outcomes problem. Even if that problem is easily solved, other issues are not, especially those relating to the priority that patients should perhaps enjoy in virtue of various morally significant characteristics, such as age or known genetic proclivities. And *combining* whatever conclusions we reach with respect to priority with the simple solution I am offering to the fair chances/best outcomes problem is a *separate* problem –

a problem that is fully as hard to solve as parts of the priority problem itself.

### **Why the Fair Chances/ Best Outcomes Problem Poses No Challenge for Large-Scale Health Insurance Decisions**

A crucial assumption underlying Daniels's formulation of the fair chances/best outcomes problem is that the two people or two groups of people whose lives or well-being is at stake are in all pertinent respects identical, except with regard to how much they might benefit from a given resource or expenditure. None of them has a moral claim to be favored because they suffered an unmerited disadvantage the others did not, or because they would be worse off if they were not helped than their rivals would be if they were not aided instead – because, for instance, one patient or set of possible beneficiaries was significantly younger than the other.<sup>vii</sup> One critical dimension of their equality is that before they landed in their predicament neither they nor their care providers or insurers knew or had reason to know that they would find themselves in the situation they later occupied, one characterized by scarcity and rival claimants. They might have recognized the possibility of that confrontation arising. But nobody knew beforehand whether, if it did arise, he or she would be the individual or a member of the group that was likely to live longer or better if treated. For, if that knowledge had been available, the people or groups would never have been on an equal footing apart from their prognoses if given treatment. One of them would have been at a disadvantage, raising a question of priority: whether justice demands some correction for disadvantages that were not the outcome of mature choice. And the fair chances/best outcomes problem as I understand it, was meant to prescind from issues of moral priority, so that the puzzle it poses might be considered on its own.

The significance of this assumption is that it places prospective beneficiaries of a particular medical treatment in a paradigmatic insurance situation: they must decide which health care options to have waiting for them, given uncertainty about the future and the value they place on having medical resources available should they need them as against other ways they might spend their wealth or the time and effort needed to acquire it. These are or should properly be seen to be (insofar as we are modeling choices that cannot actually be made one-by-one in any realistic scenario) *self-interested* choices. And rational, self-interested decisionmakers would rank treatments that yield a larger benefit or longer-lasting benefits ahead of those that produce less desirable results, for any given expenditure. It therefore follows, in Daniels's examples, that Sarah should get the lone organ rather than June, because she is likely to live eighteen years longer with it, and that it makes sense to fund a treatment program that yields larger benefits rather than one that yields smaller benefits if available funds can pay for only one.

Let me now move through the argument slowly, to address concerns that are apt to arise along the way.

If we step back and ask how much health care there ought to be and how it should be distributed, we will find what might seem a surprising degree of consensus among non-consequentialist theories of distributive justice as to how these questions should be answered, bearing in mind our key assumption, for the moment anyway, that everyone who might be insured is competent to make decisions about his or her future and none of them is known to face above-average risks of physical or mental infirmity. In fact, my belief is that the most attractive consequentialist approaches would be in broad agreement too, though I will not try to demonstrate that today.

At one end of the spectrum lie libertarian theories, some soft – for instance, in

suggesting that all natural resources and perhaps some collective human creations are commonly owned – and some more austere. Libertarians contend that losses should remain where they fall, except when relieved by charity. Thus, a life marred by ill health or a fatal disease grounds no claim of justice against those who are more fortunate. Health insurance is therefore essentially a private matter, although one can imagine the government providing at least some of the coverage, even within a libertarian framework, whether as the most efficient mechanism for amalgamating individual charitable impulses into a health care safety net or as a low-cost alternative to insurance from a panoply of private companies. The overall health care budget, as well as the allocation of expenses within any insurance plan, would in this world reflect the sum of individual preferences for insurance, given the risks people believe they face, their incomes or wealth, the disutility of labor, the efficacy of known treatments, and the costs and possible pay-offs to further research. Individual choices under circumstances of inevitable uncertainty form the touchstone for both the overall amount and the assignment of medical resources, as private actors and the state seek to satisfy individual preferences most efficiently.

At the other end of the spectrum – and there are of course many points in between – are theories of distributive justice that seem to demand far more. I defended a view of this sort some years ago, which I called equality of fortune.<sup>viii</sup> These views start from the notion that unmerited inequalities ought to be erased, to the extent that can be done consistent with honoring competing values and insofar as our technological capacities and psychological willingness allow. Ideally, the opportunities and resources people have should reflect their choices against the backdrop of unavoidable scarcity and the diversity of other people’s preferences alone, not bad luck they are powerless to guard against. So, as an initial matter, it appears that people stricken by leukemia or bipolar disorder are owed compensation by their more fortunate peers sufficient to put them as close as money and medical science can on the same plane of opportunity they occupied before their

health deteriorated. But appearances are misleading. This approach would amount to a mandatory insurance scheme that provides exceedingly generous protection at an exceptionally high cost. And there is no reason to believe that this is what people would want. At least it is clear that a great many people would not want to pay for so comprehensive and costly a shield, given other ways they might spend their time or earnings. If one believes, as I do, that a just society generally should respect people's informed choices if it can so without cost to others, then compulsory insurance at this high level cannot be justified. If personal autonomy matters, people who wish to provide less assistance to others if they find themselves in need, in exchange for accepting that they will in turn get less help if they are afflicted instead, should be free to do so. By accepting less, they would in no way harm anyone who wished to remain in the expensive, high-level insurance pool, for each simply bears the expected cost of her coverage and those who wanted more protection would be able to get what they paid for.

Hence, even theories that seem to prescribe considerable redistribution to narrow differences between people that do not flow from their informed choices come to embrace a model for health insurance that one may loosely call market-based, in that how much money gets spent, and what it gets spent on, should be responsive to people's preferences, taking into account ever-present trade-offs between costs and likely benefits, pairings that shift as technology and scientific understanding develop.<sup>ix</sup> This is the conclusion that Ronald Dworkin arrives at as well in developing his hypothetical insurance model of distributive justice.<sup>x</sup> It seems to me the outcome that any rights-based theory of justice will endorse so long as it gives wide scope to the value of personal autonomy, because deferring to individual choices means that any set of initial entitlements suggested by a background theory of justice may be set aside if people would prefer to exchange them for something they value more.

Before proceeding, it probably would be wise to note that this highly idealized picture needs to be modified in a variety of ways before it can serve as a practicable guide. Any theory of justice except perhaps the most uncompromising form of libertarianism would allow scope for paternalism on both ends. It would permit the state to require people to purchase some minimum amount of health coverage – not only to effect the outcome that they themselves would want if they knew more and were thinking clearly and to help people overcome a potentially quite harmful instance of weakness of will, but also to prevent them from imposing on others, who in the absence of their contribution would have a duty to help those in need if their care was reasonably cheap. Moreover, it would allow the state, acting as the efficient agent of its citizens' reasonable desires for protection, to regulate the practice of medicine, to evaluate the safety and efficacy of drugs, and to safeguard the soundness of insurance plans. In addition, the utopian picture of individual people considering every health risk they might run, evaluating the usefulness and cost of treatments in relation to it, then paying for the package they assemble, is almost silly. Made-to-measure health insurance is unrealistic on a variety of levels. No consumer – indeed, no single doctor, for that matter – could possibly assemble all the information needed to compose complete insurance packages in this way, and administering millions of individualized plans would make the practice of medicine even more of a bureaucratic nightmare than it already is. That said, the important point to carry away is that the aggregation of individual choices, in light of the likely costs and benefits of treatment options, is the proper benchmark for what care should be provided. To be sure, there remains the question of what care to provide to people who are unable to choose for themselves, due to youth or mental incompetence. Without going into this question further, suffice it to say that, for both theoretical and practical reasons, it seems to me that their coverage ought to track that which competent, rational agents settle on.

Let us go back to the argument's thread. The upshot of our hasty review of the gamut of nonconsequentialist theories of distributive justice on the assumption – and I apologize for belaboring this point, but it is very important – that everyone is a responsible adult with seemingly the same set of chances of suffering medical adversities through no fault of their own is that the treatments to be made available for various health problems with varying prognoses, along with the investments to be made in research, ought to be dictated by people's preferences as to the risks they run and the amount of their time and savings they wish to devote to this purpose. These are self-interested choices, decisions about how much time and cash and caution to devote to health insurance and possible technological advances *after* a just distribution of resources and opportunities has been established. To return to Norm Daniels's macro-allocation example, this means that people would have to decide, individually in each case and together collectively, whether neither or both or just one of the treatments he describes is worth paying for. Similarly, they would have to decide whether organ transplants are worth the money in a variety of circumstances, given their likely outcomes. They would also have to determine what to do when they would be willing to pay for treatments but not everyone who qualifies can receive it at a given time. The supply of transplant organs might be chronically inadequate, for example, or an emergency might cause a spike in demand that prevents everyone who ordinarily would qualify for help from getting it. In these cases, it seems self-evident to me, people would choose to have help provided in whatever way would yield the best result from a personal vantage point – the greatest improvement over the longest period. This is the standard justification for funneling health care dollars to those projects or treatments that seem likely to yield the most QALYs. It would be perverse to decide in any other manner. And so, in Daniels's example, doctors should not hesitate to give the transplant organ to the patient who likely will live 20 rather than 2 years if she receives the transplant. This is precisely what we, and June and Sarah themselves, would have wanted, even if none of us in fact ever

considered this precise scenario and recorded our wishes before we found our lives in the balance.

## Objections and Replies

Let me turn to some worries this account might engender.

First of all, somebody might question whether it is appropriate to assume, in the absence of actual insurance choices by people in a fairly priced and transparent market, that those whose hypothetical insurance choices this approach relies upon would rank treatments, including treatments in tie-breaking situations, according to the magnitude and quality of the benefits they predictably would yield. Frances Kamm raises this doubt in connection with a parallel argument, based on rational choice in ignorance of who one turns out to be, for saving the larger of two groups of people straight away if both cannot be saved:

The procedure . . . suggests that an individual consider what *his* probabilities (i.e., the probabilities of one individual) are of falling into any given position and into the more numerous group. But others have suggested that an individual reasoning behind the veil must consider what it would be like to be each of the actual people who will actually occupy the various real-life positions. This is a different procedure; the first may result in favoring maximal average utility as a policy, the second may result in maximin or, perhaps, equal chances for each in our Taurek conflict cases.<sup>xi</sup>

Kamm does not seem to endorse any of the views she says have been suggested, and in fact those philosophers she cites as raising this possibility – Tim Scanlon and Tom Nagel – do not endorse them either. So perhaps this is a straw position. It might be worth brief attention regardless, because the challenge makes plainer the nature of the hypothetical choice standard I – certainly not uniquely – have been proposing.

Our goal is not to make the overall level of health care as well as its distribution mirror the choices that imagined persons would make behind a veil of ignorance of the sort that John Rawls believes is helpful when trying to arrive at bedrock principles of distributive justice. Our goal in solving the fair chances/best outcomes problem is only to approximate the choices that actual people would make about how to spend their money or their equal claim to a fixed health care budget in the belief – this is part of the problem’s structure – that they face the same chance of anyone else of benefitting from a transplant organ or a course of treatment in circumstances in which other potential beneficiaries are in all relevant ways identical except for their prognoses. Ideally, as I said, each person would choose the details of her medical insurance plan for herself, given her awareness of the possible benefits of various treatments and their expected burdens and given her other aspirations. To be sure, people’s purchase of medical insurance might conceivably be influenced by their concerns about other people, such as spouses or children or other relatives who depend on their continued income or direct assistance. But this influence is rarely more than negligible. Health insurance is not akin to life insurance, which it makes no sense to buy (even with the tax preference for whole-life policies in the US) unless one cares about helping specific people if one dies. All that we are, or should be, trying to do is sum the decisions that potential patients would have made in the face of pervasive uncertainty about how the world will be. It is hard to imagine many people choosing for themselves being indifferent as to whether a treatment would provide a larger or a smaller benefit, or whether it would give them an added 20

years rather than 2.

I think that a better question that this response prompts is whether a covert assumption of Daniels's fair chances problem is warranted. In formulating his challenge, Daniels assumes that some unspecified "we" – are *we* the health care providers? the population of prospective patients? just the rational, competent ones? – have to pick between random choice and favoring the individual or the group of patients who will live longer or with better functioning. Those are hardly the only plausible alternatives, though. We could readily conceive of an allocation system, for example, that gave priority to those willing to buy it in advance, if not across the board then in commonly recurring situations in which the supply of life-saving or functionality-improving resources is insufficient to meet demand. We let people, quite properly, work longer or harder to earn more money to acquire goods or forms of security that others are unable to buy, given their own values and choices. The question, then, is why transplant organs, for instance, should be treated differently, if our health care system ought to mirror the summation of individual purchasing decisions. There are many ways in which a market for transplant organs might be structured, though for various reasons it seems to me that the equilibrium point might be a system that permitted insurers to sell priority points for a surcharge.<sup>xii</sup> Or perhaps the market would drive out that alternative and settle on a different tie-breaker, such as best prognosis. Or maybe the cost of the priority ticket should be made payable in some medium other than cash, such as a non-cancellable contractual obligation to donate one's own organs at death.

There are many schemes one might consider, using decision rules other than the options Daniels mentioned, namely, straight- or weighted-odds random selection or allocation by prognosis. I put the question of their attractiveness aside for a different occasion, though it probably bears noting that this family of alternatives raises a variety of difficult questions. For example, if priority is determined in part by market choices or

the express willingness of patients to contribute, say, their own organs or marrow later, what should one do about those patients who lack competence to decide on account of age or mental condition? If the background distribution of wealth and income is far from just, should that rule out market-based approaches? Indeed, even in a just order – think now of a society governed by Rawls’s difference principle, or one with tax and spending programs consistent with Dworkin’s notion of equality of resources<sup>xiii</sup> – some people almost certainly will be able to earn, after tax, more money for their labor than others, not only because of their wise or lucky choices in training themselves and choosing a career but because of genetic, educational, and familial advantages accruing to them through no merit of their own. Should that be enough to make money irrelevant to what Daniels calls micro-allocations, even though money-backed desires would generally have to be allowed to shape macro-allocations? These are issues for another day.

A second objection is that a method of choosing patients to receive a limited medical resource or for deciding which treatment regimen to fund under conditions of temporary scarcity might end up allowing some tiny difference in the expected outcome to determine who lives and who dies, or who experiences improvement and who does not. After all, a rational individual choosing between two alternatives would opt for whichever one was better, even if it was only very slightly better. And this might seem intuitively unsatisfactory. It seems to permit what Frances Kamm calls “irrelevant utilities” to play a role when instinctively we know they should not.<sup>xiv</sup> Kamm’s example of an irrelevant utility is a cure for a sore throat, a benefit which we are asked to imagine would flow to a third person if we were to save the first of two people from death rather than the second. No benefit so trivial by comparison with the value of life itself should convince us to depart from selecting the lucky party randomly. And yet that is what relying on the rational choices of self-interested decisionmakers implies. Ergo, according to this objection, the approach must be flawed.

It is hard to resist noting the ad hominem point that the suggested alternative to allowing some small difference in expected outcome to determine who survives and who does not is scarcely more appealing by its own standard: a flip of a coin is itself an irrelevant utility, if anything is. Making some small benefit on one side but not the other the relevant randomizing event seems clearly superior to using a method that contributes no added benefit at all. It also is worth remembering that differences in the expected payoffs to competing patients or groups of patients are not likely to be as minuscule in the real world as they are in Kamm's fanciful example. A new heart is not going to cure a cough in one patient but not another. Differences in patients' prognoses that are outside a sensible margin of error will surely be much greater, which blunts the force of Kamm's point considerably. Even if, conceptually, there are no irrelevant utilities on the approach I am recommending, in practice no life-saving choice is going to turn on an additional source of utility that minor. In addition, given the context – self-interested decisionmaking under uncertainty – it does not seem to me at all unsettling that a small perceived difference in the benefits of two different courses of action may produce a large difference in the fortunes of two different people. The same is true in a wide array of other situations in which we must decide where to invest our savings or our affections or whom to hire for a particular job. The candidates might be quite similar but the ramifying effects of choosing one over others can be huge.

Like my response to the first objection, however, this reply raises a more challenging question. I have been talking as though the additional benefit that would be experienced – the small if not irrelevant utility – is a benefit that would accrue to one of the two people in peril if she were saved in preference to the other. But Kamm herself spoke of the bonus as going to a *third* person, who stands to profit, if only from having her cough end, when one person is saved instead of the other. And this forces the

question of whether possible gains or losses to people other than those whose lives or well-being hang in the balance should affect the choice between them when medical resources have to be assigned. If June and Sarah would both live the same additional number of years with a new organ but June has a husband and children and parents and close friends whose lives would predictably be impoverished if she were not chosen whereas June does not, can this third party benefit determine who prevails? What if Sarah appears likely, if saved, to do far more good for other people, given her occupation and ambitions?

My hunch is that most of us would be extremely uneasy at best about giving preference to one patient based on her friendships, her family, or her job description, at least outside of special organizations such as the military or exceptionally special people, such as our President. One of the advantages of the model I have sketched is that this limiting intuition seems implicit in the framing of our inquiry. We ask what each person, taken as a moral equal, would buy *for herself* in the form of health insurance. And in fact this is the only way that a health insurance plan based on actual or hypothetical choices that serves a diverse universe of patients could work. One could hardly expect those who would be disadvantaged by a rule that took third-party benefits into account agreeing to that rule to their own disadvantage. To my mind, this way of thinking through health care priorities, along with the outcomes it favors in head-to-head rationing cases, is noticeably superior to straightforward consequentialist approaches that seem directly to require that third-party benefits – to parents, children, friends – be given weight in deciding whom to help.

The view I have been defending might now seem incredible, in suggesting that third-party benefits can be ignored in deciding whether to help one person or group rather than another. Suppose that a third person's leg could be saved if we kept alive one patient

rather than another. Could saving one person's leg really be irrelevant to the choice between patients' lives? It would be a strange health insurance plan that disregarded any benefit to a third person just so long as it was less than the benefit that is to be provided to one of two people in direct conflict. The collective expression of such a rule, across the broad run of cases, would be counterintuitive and deeply unsatisfactory. People choosing prospectively would rather have life-saving and leg-saving benefits as a part of their health package for a fixed price rather than just the first.

This last point is true, certainly, but intended as criticism this line of thought misses the mark. There are two separate points to keep in mind. First, no treatment of a third person or group, small or big, comes free. That must mean that in the choice between two people or groups who could be provided with the *same* benefit, one of them can be helped more cheaply than the rival claimant. Otherwise, there would be no money left over to pay for the third-party benefit. And that removes the apparent difficulty. The third-party benefit breaks the tie not because it adds an extra dollop of objective value to one of the two options on the table, but because it entails that one of the two equal benefits can be had less expensively than the other, leaving cash in hand to pay for an additional health care benefit too. And rational consumers would choose whichever coverage promises a given benefit at lowest cost. This way of thinking about the relevance of third-party benefits is especially appealing, because it relieves whatever pressure might be felt to allow non-medical benefits to third parties to influence primary treatment decisions.

Second, suppose that somehow a third-party medical benefit could be produced *costlessly* if one of two identical patients was treated but not if the other one was. Should that be allowed to break the tie? In my view, that depends on the nature of the third-party benefit. If the benefit is purely medical and does not flow from the identity of the person

saved – I cannot imagine a realistic example that fits this description – then it pretty plainly ought to serve as a tie-breaker, because it falls within the scope of a health insurance policy and a purchaser would prefer, as we already said, to have some probability of a transplant *plus* some other health benefit to just the transplant. If, however, the benefit to the third person would obtain as a result of some personal tie to one of the patients – maybe the loss of a loved one would trigger some devastating reaction in a third person – then it should have no bearing on the decision, for reasons we already have mentioned. Our guiding question is what people deciding for *themselves* would choose, not ghostly decisionmakers who do not know what they want in the world and who have no responsibility for their personal associations.<sup>xv</sup>

The third objection I will mention builds on cognitive dissonance. To those of us who share John Taurek’s orientation towards life-saving in cases involving different numbers of people<sup>xvi</sup> – a small minority, perhaps – the adoption of a health care policy that routinely favors patients likely to achieve the larger benefit, such as a longer life expectancy or a higher level of functioning, may seem disconcerting. If Taurek and I were on our private yacht off the coast of a small island as a long dormant and seemingly inactive volcano erupts, with time enough to rescue the villagers at either the northern end or the southern end but not both, we would flip a coin before setting course, even though the two groups differed in size.<sup>xvii</sup> None of the islanders was to blame for her plight and none chose to live in one place rather than another in consciousness of the risk of the volcano’s erupting. Because we care about them equally as individuals, we would give the same 1-in-2 chance of survival to each. When somebody loses his life, he loses what typically is most precious to him, and the loss he sustains does not increase if others perish alongside him. Crucially, from this perspective, what matters morally is not the vanishing of a valuable object, for which morality makes us stewards. What matters is the losses that specific individuals suffer, people with whom we can identify one by one,

and nobody (apart from the universe poetically conceived) suffers the sum of those losses. If one subscribes to that view, though, is it not jarring, if not outright inconsistent, to champion a life-year or functionality maximizing policy for dispensing medical goods or services?

I do not think it is. Crucially, in the medical decision cases the health care resources belong not to the person deciding but rather to patients themselves through their insurance purchases, or perhaps to the community generally in the case of naturally scarce resources such as cadaver organs. Moreover, in those cases we can make reasonable assumptions about the rules they would have approved for assigning the resources they have brought into being for a specific purpose. If the islanders had paid for a coast guard vessel and given it a general directive to save as many people as it could if a situation did not allow everybody to be saved, then it likewise would be bound to steer for the more populous end of the island when lava began to flow. Ownership and the circumstances of presumed consent to a method of allocation make all the moral difference.

The final objection I will take up disputes the overriding importance my approach gives to the decisions people would have made in advance, in the face of uncertainty as to whether anyone would need help and how many people might benefit from a course of treatment at some future time. It begins by noting that people often have a poor grasp of probabilities, fail to appreciate how they will feel if they turn out to be unlucky without adequate provision for their loss, and in particular are unlikely to understand how they will come to regard a choice between two people, or two groups of people, in circumstances where scarcity denies help to both. And it suggests that the right response to these cognitive and dispositional failings is to consider the choice between June and Sarah, or between the two potential treatment groups, on its own, without regard to broader health policies. People's status as moral equals, it might be asserted, demands no

less in situations in which some will be helped and some not. This narrowed focus may have symbolic value as well in spotlighting the moral importance a just community attaches to the life and welfare of each of its members. To be sure, it may not be clear exactly how we ought to decide who lives or which group gets treated once we have adopted this vantage point. That is why Daniels's original question continues to frustrate us. But it is clear that these trying decisions should not be made mechanically by recourse to a simple-minded maximizing principle that is oblivious to the direct confrontational nature of the choice.

These claims seem to me entirely unpersuasive. No doubt people often do have a difficult time weighing small probabilities and accurately visualizing the consequences of various events. The cure for these shortcomings, however, is to help them reason better, in the light of more complete information about the benefits of varying types of insurance. It seems the reverse of respectful of people's autonomy as agents to insist that they incorporate into their insurance policies rules of choice they almost certainly would reject because those rules would not conduce optimally to their expected well-being. There are sound reasons, not all paternalistic, for effectively forcing people to buy a minimum amount of insurance against medical adversity. But there are none for foisting on them a decision rule in case of scarcity they would not have wanted beforehand, knowing what we do about the circumstances and phenomenology of choice. The symbolic value of yoking people to a decision rule they almost all find repellant is hardly positive.

### **Remaining Problems**

Even if the fair chances/best outcomes question admits of an easy answer, there remain difficult issues in deciding how health care dollars should be spent, over and

above the problem of determining which procedures for resolving disagreements are best to which I alluded at the start.

The first is what Norm Daniels calls the priorities question. Even in a just society, some people will be worse off than others through no fault of their own, as a result of events they could not possibly have insured against. They might have had poor parents, less than nurturing surroundings growing up, bad health or accidents during their minority, congenital handicaps, or genetic predispositions that set them at a disadvantage in buying insurance or in pursuing any number of goals that matter to them. A just society would try to mitigate disparities in opportunities that arise in these ways, but it could not conceivably eliminate them altogether. In the realm of health insurance, giving each person a legal right (and giving that right teeth) to buy policies on the same terms as everyone else, or providing them with government-sponsored coverage on the same terms as others, would erase most of their disadvantage. But it would not remove the problem completely in rationing situations. We would still need to determine whether a disadvantage for which someone is not responsible and which society cannot redress completely qualifies him morally for some form of favoritism in the allocation of future benefits. It may be worth noting that this problem is likely to become more acute over time, as we learn more about the genetic origins of or links to later life-shortening diseases.<sup>xviii</sup> If one person would likely live less long with a scarce organ than another because of a genetic condition that she has been known to have since birth, we cannot imagine that person agreeing to a simple life-year maximizing rule if her actual preferences are consulted. We would then need to decide whether to rely on an even more depersonalized conception of hypothetical choice, as Ronald Dworkin does, or to adapt our answer to the fair chances/best outcomes problem to accommodate this known liability. This is one instance of the broader problem of combining whatever answers we think best to the priorities question with our life-year or functionality-maximizing

conclusion to the fair chances question to yield a single set of decisions rules in scarcity cases. We might be fortunate, as moral philosophers if not as patients, that the benefits that come from donor-donee organ matching shrink the scope of the scarcity rationing cases significantly.

Second, we face a different sort of priorities question when we consider offering medical help to people in substantially poorer countries. Some are apt to be *much* worse off than others, and blamelessly so. To what extent a single earmarked form of assistance, such as health care, should be directed primarily to offsetting much larger resource inequities rather than devoted to producing the largest aggregate health benefit is a difficult issue about which I have nothing now to say.

---

1. Norman Daniels, "Four Unsolved Rationing Problems: A Challenge," 24 *Hastings Center Report* 27 (July-August 1994).

ii. Norman Daniels & James E. Sabin, *Setting Limits Fairly: Can We Learn to Share Medical Resources* (2002).

iii. See Ronald Dworkin, "Liberty and Moralism," in *Taking Rights Seriously* 240, 248-55 (1977).

iv. Daniels, *supra* note 1, at 27-28.

v. Daniels attributes this view to Dan Brock and Frances Kamm. See Dan Brock, "Ethical Issues in Recipient Selection for Organ Transplantation," in *Organ Substitution Technology: Ethical, Legal, and Public Policy Issues* 86-99 (Deborah Mathieu ed. 1988); Frances Kamm, *Mortality, Mortality Vol. 1: Death and Whom to Save From It* (1993). In the nearly analogous case of choices between the lives of different numbers of people, a number of authors favor assigning chances of survival in a manner proportional to the sizes of the competing groups, at least in many contexts. See Kamm, *id.* at 123-43; John Broome, "Selecting People Randomly," 95

---

*Ethics* 38 (1984); John Lockwood, "Quality of Life and Resource Allocation," in *Philosophy and Medical Welfare* 54 (J.M. Bell & Susan Mendus eds. 1988).

vi. *Id.* at 28.

vii. My view is that age is in many instances relevant to choices between people and that younger patients often are entitled to priority in cases of medical scarcity and research funding. See Eric Rakowski, "Should Health Care be Rationed by Age?" in *Controversial Issues in Aging* 103 (Andrew E. Scharlach & Kenard W. Kaye eds. 1997)

viii. Eric Rakowski, *Equal Justice* 19-106 (1991).

ix. For another expression of this view, see David M. Eddy, "Rationing by Patient Choice," 264 *Journal of the American Medical Association* 105 (Jan. 2, 1991).

x. Ronald Dworkin, *Sovereign Virtue* 307-50 (2000).

xi. Kamm, *supra* note 5, at 120.

xii. I discuss organ allocation and possible market-based schemes for determining priority in *Equal Justice*, *supra* note 8, at 313-32.

xiii. Dworkin, *supra* note 10, at 65-119.

xiv. Kamm, *supra* note 5, at 144-64; Norman Daniels, "Rationing Fairly: Programmatic Considerations," 7 *Bioethics* 224, 229 (1993).

xv. I admit that there is a difficulty here as concerns young children who are dependent on one of the patients in need of a scarce live-saving resource. Children do choose their parents and are not competent insurance buyers in their own right. My sense is that, although this point counts in favor of allowing benefits to them to sway medical choices, that point is outweighed by the importance, real and symbolic, of treating patients in need equally, based on their own needs and likely prospective commitments, and not on their popularity or their independent decisions to have children.

xvi. See John Taurek, "Should the Numbers Count?" 6 *Philosophy & Public Affairs* 293 (1977).

xvii. Eric Rakowski, "Taking and Saving Lives," 93 *Columbia Law Review* 1063, 1155 (1993).

xviii. I address a number of puzzles and difficulties in compensating for genetic bad luck in Eric Rakowski, "Who Should Pay for Bad Genes?" 90 *California Law Review* 1345 (2002).