
McGILL LAW JOURNAL

REVUE DE DROIT DE MCGILL

Montréal

Volume 38 1993 No 4

McGILL LECTURE IN JURISPRUDENCE AND PUBLIC POLICY

Justice in the Distribution of Health Care

Ronald Dworkin*

In this lecture, Professor Dworkin begins by identifying two questions about justice in the distribution of health care: (1) How much, in the aggregate, should society spend on health? (2) Once established, how should this amount be distributed?

He then examines the ancient *insulation* model of health care distribution, which postulates that health care is chief among all goods and that it is to be distributed in an equal way. He concludes that this model provides no satisfactory answer to either of the two questions. It cannot answer the first, for it would require that society spend all it could on health care until the next dollar would buy no gain in health or life expectancy, something which is manifestly absurd, particularly in our age of ever-expanding medical technology. Nor, he tells us, does the insulation model provide much guidance with the second question, since its egalitarian spirit ultimately leads us to apply notions of efficiency and need which are philosophically controversial and therefore impossible to apply.

By means of a thought experiment, Professor Dworkin then develops an alternative model which he feels does provide an answer to his two questions. He asks us to imagine a society with fair equality in the distribution of resources, in which the public at large has knowledge about the cost and value of medical procedures, but in which no one has any knowledge about the antecedent probability of contracting any particular disease. Moreover, health care is not provided by the government, but, rather, each individual is free to allocate to health care (by purchasing health insurance, for example) as much or as little of his resources as he wishes. Professor Dworkin claims that whatever that society spent on health care would be just — both in the aggregate and in its distribution. Carrying the model through, he discusses its implications for our own society and analyzes possible objections to it. He concludes by stressing the importance of the question of justice in health care and by putting it in its broader political context.

Dans cette conférence, le professeur Dworkin pose d'abord deux questions sur la justice dans la distribution des soins de santé : (1) Combien, en tout, la société doit-elle dépenser pour les soins de santé ? (2) Une fois ce montant établi, comment doit-il être distribué ?

Il examine ensuite l'ancien modèle d'*isolation* des soins de santé. Selon ce modèle, la santé est le plus grand bien qui soit et la distribution des soins de santé doit se faire de façon égale. D'après le professeur Dworkin, ni l'une ni l'autre des deux questions n'est résolue de façon satisfaisante par ce modèle. Le modèle ne répond pas à la première, car il exigerait que la société consacre toutes ses ressources à la santé jusqu'à ce qu'une dépense supplémentaire ne puisse apporter aucune amélioration de la santé ou de l'espérance de vie, conclusion incontestablement absurde, surtout dans une société comme la nôtre où la technologie progresse sans cesse. Le modèle d'*isolation* ne répond pas non plus à la deuxième question, car son approche égalitaire exigerait qu'on tienne compte de facteurs comme l'efficacité et le besoin, lesquels sont très controversés et se prêtent donc mal à une application rigoureuse.

Dans le but d'élaborer un modèle qui nous permette de répondre à ses deux questions, le professeur Dworkin nous invite à envisager une situation hypothétique : une société imaginaire où il y a une distribution équitable des ressources, où le public possède des connaissances médicales sur le coût et la valeur des différents traitements possibles, mais où personne ne peut prévoir les chances de contracter telle ou telle maladie. De plus, ce n'est pas le gouvernement qui finance les soins de santé, mais plutôt l'individu, qui est libre de consacrer son argent aux soins de santé (en prenant une assurance, par exemple) comme bon lui semble. Selon le professeur Dworkin, le montant que cette société dépenserait pour les soins de santé serait juste — et dans sa totalité et dans sa distribution. Poussant plus loin son modèle, il en tire des conclusions pour notre société et examine les objections qu'il pourrait susciter. Il termine sa conférence en insistant sur l'importance de la question de la justice en matière de soins de santé et en la situant dans son contexte politique.

* Professor of Jurisprudence, Oxford University; Professor of Law, New York University School of Law. Professor Dworkin delivered this lecture on March 17, 1993 at the Faculty of Law, McGill University, as the Inaugural Lecture of the McGill Lectures in Jurisprudence and Public Policy.

© Ronald Dworkin 1993

To be cited as: (1993) 38 McGill L.J. 883

Mode de référence: (1993) 38 R.D. McGill 883

1.

You might think it very odd that, living as I do half the year in Britain and half in the United States, I've come to Canada to talk about justice in the distribution of medical care. In both those countries, the Canadian structure for the distribution of health care is taken by many people to be a model of success. Your Dean mentioned my former colleague at University College, Oxford, who is now my President. Under his administration, and under the supervision of his wife, the United States is now, as you know, engaged in a massive re-examination of health care. Almost every day an article appears in the American press about the Canadian plan, which is widely proposed as a model for Americans to follow. Nevertheless, strains on your system of health care are beginning to become evident. There's talk of rationing, and more people go south of the border to seek medical care. Doctors in negotiation with the provincial authorities claim with greater stridency that they are seriously underpaid. The system, nevertheless, is producing more Canadian doctors than economists think wise. You, too, will face the problems I'll discuss this evening. You too must worry about justice in the distribution of health care when it comes — and I'm sure it will come — to rationing health care explicitly.

Some people, particularly in America now, say there is really no need to ration health care. They agree that medical expense already constitutes an alarming proportion of the American economy, and that, even though America spends that much in the aggregate, forty million Americans are wholly uninsured or without any adequate health care, which is intolerable. But they deny that correcting these deficiencies will require some form of rationing: they say that there is so much waste and inefficiency in the health care system that, if these were eliminated, we could save enough money to insure that everyone had all the medical treatment he needed. We know what they mean. The administrative inefficiency of United States medical insurance companies and carriers is legendary. American doctors' salaries are large and, according to many people, inflated. The average medical salary two years ago in the United States was over \$160,000 — the *average*. Nevertheless, a series of recent studies suggests that even if administrative efficiency were greatly improved, and even if doctors' salaries were capped at some reasonable level, rationing of health care would still be inevitable, because by far the biggest cause of the explosion in health care costs (not only in the United States, but, I believe, in Canada as well) is a massive supply of new technology. It isn't that we're paying all that much more for what we formerly bought cheaper; it is that we now have so much more to buy.

Many politicians and some doctors say that much of the new, expensive technology is "unnecessary" or "wasteful" or "inappropriate." But if you look to see what they mean, you find they have in mind techniques that are (as it's often put) "low yield," which is not the same as "no yield." They point, for example, to massive mammography screening of women under the age of fifty, or to the heavy use, in some medical facilities, of magnetic resonance imaging. A society that spends a great deal of money on routine screening or expensive diagnostic equipment may not save *many* more lives than a society that does

not. But it will presumably save *some* more lives, and that means that we cannot appeal just to efficiency as an abstract value to justify saving the cost. We cannot recommend eliminating “inappropriate” medical care without deciding what medical care is appropriate, and why, and that, in turn, depends upon how we answer the question: “how much medicine *should* a society provide?”

That question can usefully be divided into two more specific ones. The first is the question of a community’s *aggregate* health care budget. Money spent on health care (I include not just acute care but also preventative medicine, care of the chronically sick or disabled, and so forth) is money that might be spent on education, or on economic infrastructure that will produce more jobs. How much of the overall budget should be devoted to health care instead of other plainly valuable projects, like these? The second question, though it’s really part of the first, is the question of distribution. Once it’s established what a society should spend overall on health care, then it must also be decided who should have that care, and on what basis it should be allocated. Of course, nations struggling with health care costs must resolve many issues beyond these twin questions of justice. There are economic questions — what are reliable predictors of how much a particular health care plan or structure will cost? There are administrative issues: what is the most efficient organization for administering any particular plan? There are medical questions: what is the likely impact of a particular program on morbidity and mortality? Above all, there is the political problem: what plans will a particular democracy in fact be willing to accept and pay for? I don’t mean to denigrate the importance of these various problems, or to deny their evident connections with the problems of justice. But I shall concentrate on the latter. To repeat: In all justice, how much should a decent society spend on medical care, broadly described? In all justice, how should that society distribute what it does spend — who should get what? Behind these two questions lies a more explicitly philosophical one. What is the right standard to use in answering these questions? What should we take as our ideal of justice in medical care?

2.

I begin by describing an ancient and attractive ideal that many people instinctively accept, which I shall call the ideal of *insulation*. It has three features. The first argues that health care is, as René Descartes put it, chief among all goods: that the most important thing is life and health and everything else is of minor importance beside it. The second component of the insulation ideal is equality. The ideal supposes that even in a society which is otherwise very unequalitarian — indeed even in a society in which equality is despised as a general political goal — medical care should nevertheless be distributed in an equalitarian way so that no one is denied care he needs simply because of an inability to pay. The third component (it really flows from the other two) is the old principle of rescue, which holds that it is intolerable when people die, though their lives could have been saved, because the necessary resources were withheld on grounds of economy.

This ideal of insulation has exerted great power throughout history. Hospitals have always been paradigm examples of appropriate charities, and religion has, from ancient days, always been associated with them. Contemporary political philosophers — I have Michael Walzer in mind, for example — say that the provision of medicine constitutes a separate sphere of justice, and that in that sphere decency, community, solidarity, and equality must reign. The power of the insulation ideal is so great that people often think that though the administrative, medical, economic and political problems I described are intellectually daunting, the questions of justice are not: that it is clear what the ideal of justice demands in health care, and that our only problem is that we are unwilling to live up to that ideal. That is, I believe, a serious mistake. The crisis in health care includes a crisis in our conception of what a just health care system would be — what answers we should give to the questions of justice I set out. We face that intellectual crisis because it has become clear that the insulation ideal, for all its ancient popularity, is now irrelevant. Consider the first question I posed: the problem of the aggregate expense a decent society will commit to health care, as against competing needs and values. What advice does the ideal of insulation give? It says a society should spend all it can on health care until the next dollar it spends would buy no gain in health or life expectancy at all. Of course no society ever did organize its affairs in that way, any more than any sane individual organizes a plan of life with the goal of making that life as long and as healthy as possible. In past centuries, however, there was not so significant a gap between the rhetoric of the insulation ideal and what it was medically possible for a community to do. It was possible to give lip service to the ideal, and charge social failure to live up to it to collective moral shortfall. But now — when technology continues to produce more and more ways to spend great sums on medical care — it is self-evidently preposterous that a community should treat health as lexicographically prior to all other values. Any community that really tried to do so would secure for its citizens marginally longer lives, perhaps, but these would be lives barely worth living. Once, however, this suggestion of the ancient ideal is rejected as incredible, the ideal has nothing more to say. It has, as it were, no second best or fall-back level of advice. It simply falls silent.

In fact, as a result, philosophers, theorists and medical specialists who nominally subscribe to the ideal of insulation all despair of attacking the first question. After some discussion, they announce that the size of the overall medical budget will be “decided in politics,” which is an academic way of saying that abstract considerations of justice have nothing much to contribute to this part of the health-care discussion. I believe that that is a mistake; if I am right, then the dominance of the insulation ideal has been a hindrance, and not just not a positive contribution, to achieving justice in health care.

Now look at the second question, the question of distribution. When the theorists finish saying that politics will set the overall health-care budget, they quickly add that justice will require that that budget, whatever it is, be spent in a fair way. But how does the insulation ideal help us to define a fair distribution? It tells us something negative and undoubtedly important: that how someone is medically treated should not depend, in our society, simply on abil-

ity to pay. It tells us that if rationing is necessary, the principle of rationing should not be, as it now largely is in the United States, the pocket book. But we need more positive advice. What should the principle of rationing be, if it is *not* to be money? Once again, the ancient ideal has very little to say. The egalitarian impulse of the ideal seems to recommend that medical care be distributed according to some principle of efficacy and need. And so people committed to the ideal speak about rationing according to cost-effectiveness or according to some principle that requires money to be spent where it will do the most good.

As many of you know, the state of Oregon established a commission some time ago to try to give structure to that idea, to try to describe what rationing health care in accordance with effectiveness would mean. The difficulty, of course, as that commission discovered, is that the concept of doing the "most good" (or, in more academic terms, of maximizing welfare, or utility, or well-being, or happiness, or capability) is systematically and multiply ambiguous. These various terms, when properly used, do not name psychological concepts. Or medical or, in my view, economic ones. They name contested ethical concepts: the proposal that health care money should be spent to do the most good means that it should be distributed in whatever way will make the lives of citizens better lives to have lived, and that goal cannot be restated, without controversy, as the goal of making lives more pleasant, or economically more productive, or socially more beneficial. Whenever you attempt to describe in more detail what making the lives of citizens better actually means, you enter the kind of controversy that it was the promise and hope of the insulated ideal to avoid, and it would be sheer disaster to try to reduce that ideal to something mechanical enough to be measured by a computer. The Oregon commission discovered this. It developed mechanical measures of the cost-effectiveness of various sets of treatment matched to various kinds of disease, typed these metrics into a computer, typed in a great deal of further information, and watched the computer produce a ranking of cost-effectiveness that ranked capping a tooth higher in social priority than appendectomy. It's perfectly true, the computer said, that you will die if you have appendicitis and don't have your appendix removed. But it costs four or five thousand dollars to do that and dentists can cap a great many teeth and prevent a massive amount of toothache if you spend that five thousand dollars on dentistry instead. Well, of course, as soon as that result appeared, the commission saw that its algorithms were hopeless, re-rigged its operational definitions, and produced something at least less implausible than that. But the story indicates the character of the problem I have in mind.

So the old ideal of insulation fails to answer our second question as well as our first. Its proposal, that health care should be distributed according to need, or so as to do the most good, or so as to improve overall welfare, is fatally ambiguous, and becomes evidently unattractive when the ambiguity is resolved by defining success in terms of some utilitarian reading. We have not, after all, settled the question of what justice in health care means, and that philosophical problem stands beside the economic, medical and administrative problems we know we face, and it may be at least equally daunting.

3.

This evening I shall try to construct, at least in very broad outline, an alternate approach to justice in health care, which is based not on the insulation of health care as a separate sphere of justice or activity, but, on the contrary, on the *integration* of health care into competition with other goods. I shall describe an approach that, I believe, is more instructive about the two great issues of justice I named. I can state the central idea in advance: we should aim to make collective, social decisions about the quantity and distribution of health care so as to match, as closely as possible, the decisions that people in the community would make for themselves, one by one, in the appropriate circumstances, if they were looking from youth down the course of their lives and trying to decide what risks were worth running in return for not running other kinds of risks.

At some point (as those of you who have read any political philosophy written after the middle ages know) an imaginary story gets told. My story has the virtue of being less imaginary than some others, but it will nevertheless require you to exercise your imagination. Suppose that your community were to develop and change in the following three ways.

First, *per impossibile*, suppose it developed into a society in which the economic system provided "fair equality" in the distribution of resources. I mean that government recognized its inevitable responsibility to choose amongst economic and tax structures, and chose a structure that treated all members of the society with equal concern. I have my own idea about what that means in practice, and I've tried to spell this out in a series of articles.¹ I said (this is a very crude summary) that an economic structure treats all members of the community with equal concern when it divides resources equally, measured by the opportunity costs of each person owning a particular resource, and then leaves each member free in principle to spend those resources designing a life that each believes appropriate. That conception of equality will not make people equal in the amount of money or goods each has at any particular time; still less will it mean that everyone will lead the same kind of life. Some people will have invested and some people will have consumed. Some will have spent early and some will have saved for late. The result will nevertheless be egalitarian, because the choices people will have made will answer to their own conceptions of what life is right for them.

These are my views about what a just economic system would be like, but I offer it only by way of illustration. You may — you probably do — have a different conception of what economic structures genuinely treat all people with equal concern; if so, your view of how that community would have changed, in order to meet my first condition, will be different from mine. That does not matter for the present exercise: I merely ask you to assume that it has changed, in whatever way you think justice and equal concern require.

Second, imagine that your community is also different in that all the information that might be called, roughly, state-of-the-art knowledge about the value

¹The central article, for purposes of this lecture, is *What is Equality? Part II*, which appeared in *Philosophy and Public Affairs*, Fall 1981.

and cost and side-effects of particular medical procedures — everything, in other words, that very good doctors know — is known generally by the public at large as well.

Third, imagine that no one in your community — including insurance companies — has any information available about the *antecedent* probability of any particular person contracting any particular disease or infirmity that he or she does not evidently already have. No one would be in a position to say, of himself or anyone else, that that person is more or less likely to contract sickle-cell anemia, or diabetes, or to be the victim of violence in the street, than any other person. So no information exists about how likely it is that young blacks, as distinguished from people generally, will die in violent fights, for example.²

The changes I am asking you to imagine in your community are heroic. But they are not, I think, beyond the reach of imagining, and I am not inviting you to imagine other changes. Indeed, I am asking you not to: I want you to assume that your preferences and ethical convictions, and those of other members of your community, have remained constant in spite of these changes. Very well. Suppose that your community is indeed changed in those three ways, and then also suppose that health care is simply left to individual market decisions — in as free a market as we can imagine. Medical treatment is not provided by the government for anyone, as it is for everyone here and for some people in America. Nor are there any government subsidies for health care — in particular, the premiums people pay for health care insurance are not, as they are now in the United States, tax-deductible. If people choose to purchase such insurance, they do so as they buy anything else: out of post-tax funds.

What kind of health care arrangements would develop in such a community? How much of its aggregate resources would end up devoted to health care? How would medical treatment in fact be distributed among its members? Well, of course, it is hard to say; indeed it is impossible to say with any precision, though I shall offer you some speculations in a moment. But I'm anxious to make two claims in advance of any such speculation, to show you why the question of what such a society would do is important. The first is that *whatever* the society I've just described spends as its total health-care budget, which means simply the aggregate of what individuals spend, would be the just and appropriate expense for that society. The second is that *however* health care is distributed in that society would be a just distribution of health care for that society. I must qualify those two dramatic claims to some degree, but the qualifications I need are not major, and I'll relegate them to a footnote.³ So I shall

²I am ignoring an important issue that I will have to consider in a subsequent full presentation of this material. Is it right, in the hypothetical exercise I am constructing, to exclude information relating risk of disease to voluntarily chosen behaviour? Should insurance companies be in a position to charge cigarette smokers or mountain climbers higher premiums, for example? If so, then what counts as voluntary behaviour? Should sexual behaviour of a particular kind be treated as voluntary for this purpose? Should insurance companies be able to charge active male homosexuals higher premiums because they are more likely to contract Aids?

³Some paternalistic interference with individual decisions about health care insurance, particularly those people make early in their lives, might be necessary out of fairness to people who might

proceed on the flat assumptions just stated: nothing that society does, by way of health care arrangements, is open to objection on grounds of justice, though of course what it does might be questionable or objectionable in many other ways.

If so, then it is indeed important to consider what health care arrangements our society would make if it were changed in the ways I described, because, as I shall argue, what they would do through independent decisions can serve as a guide to what we should do, in whatever way we can, to improve justice in our own circumstances. So speculation seems worthwhile. It seems likely that even though the members of the imagined community — our community transformed — would perhaps begin by making individual insurance decisions, they would soon develop, through these individual decisions, collective institutions and arrangements; it also seems likely that progressively more and more people would join those collective arrangements. They would develop very large cooperative insurance plans, or very large health maintenance organizations which provide stipulated categories of medical care for a stipulated advance contract price, or both, for example. As such plans became larger, and more efficient, it would become progressively more and more expensive, relatively, for people to make wholly personalized, individual medical arrangements for themselves, and progressively fewer and fewer people would do so. (Remember that in this society wealth is much more equally distributed than it is in our society now, and though some people are relatively rich there, they are mainly people who have decided to concentrate on saving.) So the number of people who could and would turn their back on the economies of scale and administration of the collective provisions will be few, and, as the process continues, fewer still. The result of the process might very well be something functionally very close to the single, comprehensive health care provision scheme that you have reached here in Canada. Large insurance cooperatives or health maintenance organizations might negotiate a basic scheme of provision that would be much the same for everyone. If so, however, the community would probably also develop a secondary insurance market: people would be free to negotiate specialized insurance in *addition* to that basic insurance package. What form that secondary market would take, and how large a market it would be, would, of course, depend on factors we cannot sensibly predict. But even in a much more egalitarian society, some people would be able and willing to make provision for queue-jumping, or elective cosmetic surgery, or other benefits that the basic provision made available through general collective schemes would not provide. (In a more egalitarian society, the cost of some of these special benefits might well be lower than it is now — since doctors' salaries, for example, would presumably be lower, specialized services might be available at lower cost.)

We need not dwell on the character of that secondary market: it is more important to consider the basic, standardized coverage packages which I'm assuming that the large cooperative institutions would provide. What would be the character of those packages? Well, of course, that would depend upon the

make imprudent insurance decisions when young. And some constraints and requirements might be necessary in the interests of justice toward later generations.

mix of preferences and convictions. But we can speculate with some confidence about what would *not* be covered in such a plan. Some private insurance decisions would be plainly irrational in the imagined community: they would be what the economists might call *dominant* mistakes, by which I mean they would be mistakes, even in retrospect, no matter what happened in the future, including the worst. I'll give you one or two examples: they are extreme, but of course they would be, given the claim I've just made about them.

Almost no one would purchase insurance that would provide life-sustaining equipment once he had fallen into a persistent vegetative state, for example. That would be a dominant mistake: the substantial sum spent year-by-year in insurance premiums to provide that coverage would be at the expense of training or experience or culture or investment or jobs that would have enhanced real life. Even someone who lived only a few months after purchasing the insurance before he fell into a vegetative state would have made, in retrospect, a mistake, giving up resources that could have made his short remaining conscious life better to buy a longer unconsciousness. My second suggestion might seem more controversial. I suggest that almost no one would purchase insurance providing for expensive medical intervention, even of a life-saving character, after he entered the late stages of Alzheimer's disease or other forms of irreversible dementia. Almost everyone would regard that decision, too, as a dominant mistake, because the money spent on premiums for such insurance would have been better spent, no matter what happens, making life before dementia — life in earnest — more worth while.

Now I come to a further suggestion, more controversial still. In most developed countries, a major fraction of medical expense — in the United States it approaches forty percent of the health care budget — is spent on people in the last six months of their lives. Of course, doctors don't always know whether a particular patient will die within a few months no matter how much is spent on his care. But in many cases, sadly, they can say, with considerable confidence, that he will. I believe that if people reflected on the value of buying insurance that would keep them alive, by heroic medical intervention, four or five more months, in the condition in which most such patients undergoing that intervention live, compared with the value the premiums necessary to purchase that insurance could add to their earlier lives if spent in other ways, they would decide that buying that kind of insurance was not a wise investment. That is not to say, of course, that most people would not want those additional months, no matter in what state or condition they spent them. Many people want to remain alive as long as possible, provided they remain conscious and alert, and provided the pain is not too great. My point is rather that they would not want those additional months at the cost of the sacrifices in their earlier, vigorous life that would be necessary if they had to make that choice. They would think the money better spent, earlier, on job-training or education or investment or on something else that would benefit their lives as a whole more than just taking on a few months of very limited life at the end. I cannot quite make the claim here that I made about persistent vegetative state or advanced-stage Alzheimer's disease: that purchasing insurance for costly procedures extending life a few months would be a dominant mistake. We can imagine circumstances — some-

one falls fatally ill the day after buying a policy providing for such care — in which, in retrospect, the decision to buy it turned out to be a good one. But most people would agree, I think, that in the circumstances we are imagining — in which, remember, no one knows he is more likely than anyone else to contract a disease not already evident — that decision would be an *antecedent* mistake.

How much further can we go down this road? How much more insurance can we be reasonably confident people would not buy in the circumstances we are imagining? I'm not sure, and anyway have no time to explore other examples now. But I do want to raise, at least, one further issue which, as I suggested to you earlier, is already of crucial importance and will become even more critical in the next decades. How far would people in the imagined community make provision for access to the ultra-expensive medical equipment now in use, or which is being developed, or is still over the horizon?

I came here from attending a meeting at the Harvard Medical School in which new advances in technology were being described. You ain't seen nothin' yet. I've already mentioned advances in diagnostic radiology: expensive magnetic resonance imaging, for example. Much of the talk at this meeting was about molecular biology: about, for example, promising research into treating cancer by creating monoclonal antibodies specific for each patient, from the patient's own genetic material, at stupendous cost, and new, very expensive, blood tests that marginally — very marginally — improve the accuracy of a diagnosis of heart disease. Each of these examples illustrates, though in different ways, how technology might come to be regarded as "low yield" relative to its large cost. Both would undoubtedly save some lives. But at a cost, in development and production, that might seem very high when we consider how a community might use the funds in other ways that would enhance the economy and provide more jobs and a higher standard of living for more people.

Would people in the imagined society, ultimately deciding for themselves how to allocate their resources, provide for expensive and/or speculative technology? People informed and reflective might make distinctions along the following lines. They might pay to provide life-saving techniques for diseases that tend to occur relatively early in life, particularly when these techniques have a high probability of success. But they might not spend to insure for technology that is very speculative, even though it will save some lives, or for technology whose main results benefit people in relatively old age. Paying all our lives to secure the latter kind of technology, if we need it, might seem a poor decision when it means that we run a higher risk than we need to run of unemployment or an otherwise less satisfactory life. I won't pursue these speculations further. I hope I've given you some idea of the kind of choices that people in the conditions we're imagining would have to make, and of how they might be tempted to make them.

4.

Do you resist my claim that whatever such a society spent, through collective institutions governed by individual decisions of this character, would be just, and that the distribution of health care such a society achieved would also

be just? You will not, if you accept a conception of social justice that assigns individuals responsibility for making the ethical choices for their own lives against a background of competent information and a fair initial distribution of resources. If you accept that vision of a just society, then you will accept my claim — though, as I said, you may well have a different conception of what a fair initial distribution of resources would be like, and how unjustified inequalities should be remedied, than I do — in which case your understanding of the conditions I described will be correspondingly different from mine.

So I will assume that you do agree with my main claim: that whatever our imagined society achieves, by way of health care arrangements, cannot be faulted on grounds of justice. I suggested, earlier, that we might therefore make practical, political use, for our own communities, of at least our less speculative conclusions about what people in the imaginary community would provide for themselves. There is a natural way in which we might be tempted to do this. Almost all government-sponsored or supervised health schemes now in existence, and almost all of those that have been proposed as vehicles of reform in the United States, define a basic health-care package of benefits that must be made available, at responsible cost, to everyone, and supplied without charge to those who cannot pay that responsible cost themselves. We might use our speculations about the imaginary society to help us to define what should be in that basic package, and what that responsible cost should be.

In one way, at least, the imaginary story might be helpful for countries, like the United States, who have not settled on a particular structure for health care reform. As I said, many people in America believe we should follow your example in constructing a single-payer arrangement in which government, not private insurance firms or health care providers, decides what medicine to offer and at what price. But others think the United States should adopt what is called a scheme of “managed competition,” in which private insurers compete to offer a basic package stipulated by government, and government supervises their performances and premium structures. As of this evening, at least, most commentators predict that a managed competition scheme will be adopted, primarily, they say, because it is better suited to the political culture of the United States than a single-payer scheme would be. But our imaginary story might be helpful in guiding the choice between the two forms of scheme, in the following way. The decision might turn, among other things including suitability to the political culture, on the degree of confidence we have in our speculations about what people would choose in the imaginary community. If we were reasonably confident that we knew roughly what such people would buy — what the dominant collective arrangements they would reach would provide — then that would argue for trying to set in place a single-payer system like yours or like the National Health Service in Britain. Government can more effectively guarantee people what it is persuaded justice demands that they have if it is free to provide it itself, in some such way. To the degree we are uncertain about what people in the imaginary world would decide, however, that argues for a scheme of managed competition with enough flexibility to allow different people to choose different packages all meeting a common stipulated standard. The choices that

actual people make among such schemes would provide a self-regulating mechanism that would bring us closer to the just distribution of the imagined world.

But of course whether the United States ultimately chooses a single-payer scheme like yours or, as seems more likely, a scheme that includes private competition, is more likely to depend on considerations other than justice. Nor, I think, is justice decisive of that issue one way or the other. Both types of scheme include the idea of a basic package (or set of such packages) of insurance made available to all, and the main issues of justice consolidate in the question of what should be in that basic package or set of packages. That is the question, as I suggested, that is most directly responsive to the exercise I've been imagining. I offered you reasons for thinking that certain kinds of insurance or health organization contractual provision would be rare in the community we imagined, and that, I now submit, is a good reason why that kind of provision should *not* be part of the basic package that will be the heart of any reform in the United States and any readjustments here in Canada. Since those are expensive provisions, this is an important result. But it is a negative one, and the exercise must be conducted on the other side as well. I have little doubt that people in the imagined community would insist on provision for standard prenatal care, for example, and on the kinds of primary medical care, including relatively inexpensive routine examinations and inoculations that poor people in the United States so conspicuously lack. It follows, from the argument I have been making, that these are essential elements in the basic package that any responsible health care reform would establish.

5.

I have been exploring ways in which practical health care administration and reform could be guided by the exercise I hold out: trying to imagine what health care people in the imaginary circumstances I described would provide, out of their own pockets, for themselves. It is past time, however, for me to consider the drawbacks and pitfalls of my overall argument. One danger is evident: my suggestions about how people would behave in the imaginary society are speculative, and even though some of these speculations seem very plausible, we cannot test them by asking how everyone actually behaves in communities as they are now constituted. Resources are unjustly distributed among us: Canada is not as bad in this respect as is the United States, but even Canada is very far from ideal justice in economic distribution. We obviously don't have a society in which people enjoy state-of-the-art information about medicine. On the contrary, people's medical ignorance is often cited as one reason why medical expenses continue to rise. And, of course, our insurance companies do know that risks are higher among certain groups within the community than others, and the curse of experiential rather than community rating for premiums has dogged attempts to make commercial medical insurance fair.

But it doesn't follow that our speculation about what people would want under very different, and fairer, circumstances must remain just speculation. The choices Americans of average income make about their employee insurance package in wage negotiations, for example, can offer some guidance. And

research and publicity can provide better guidance. Not only government but private organizations — large medical schools, for example — could help design a few sample paradigm insurance protocols representing different insurance strategies. Some of these would provide for catastrophic care or transplant surgery in circumstances in which others denied it, for example. The protocols could be accompanied by the medical information of the kind that is crucially missing from public awareness now: by some realistic expert opinion of the expected consequences for mortality and morbidity from a public commitment to each protocol, together with, for each, some estimate of its total cost and consequent macro-economic effect. If information of that sort were put into the public domain, and challenged and debated there, the resulting discussion would be at least minimally informative about how much people value what kind of care, and might be very informative. When we think of the kind of opinions that pollsters examine now, and that feature on television discussion shows and radio phone-ins, we might welcome a shift to the kind of discussion I'm now imagining.

A second difficulty is potentially much graver, however, at least conceptually. I've imagined a utopian (in some respects) society and I've then suggested that we set out to copy one feature of that utopian society: the provision it would make for medical care. An economist will remind me that, when the first best is impossible, the second best is not always achieved by mimicking the first best *partially*. That may, indeed, make matters worse than the *status quo*, and it is not difficult to see this possibility as a threat to my argument. Suppose, for example, that we decide that if our community were just, and different in the other ways I imagined, the standard medical package nearly everyone would purchase through collective insurance arrangements would include a particular set of benefits. If we decide, therefore, that that set of benefits should make up the basic package that must be made available to everyone in our own community, some relatively low-income people may end up paying a higher share of their actual income for medical care (for themselves and, through taxes, for others) than they would have chosen to pay in a just society. Or, to put the matter the other way around, they may have less left over for other expenses than they would have chosen in those circumstances. That may not seem, particularly to them, an improvement in justice.

I do not want to minimize the problem this hypothetical example illustrates. But the possibility that the test of justice I propose *might* produce unjust results is not, in itself, a sufficient argument against accepting that test; someone who objects must show a strong likelihood that the result would in *fact* be worse, from the point of view of justice, than using some other defensible standard for designing the basic package of protection. This is not a question of who has the burden of proof. If it is true that if our economic structure were just, everyone would be able to and nearly everyone would purchase a particular medical provision, that supplies a very strong even if not decisive argument that our structure would be closer to a just one if we made sure that everyone had that provision now. We should act on that strong argument unless we have some positive reason, not just the bare possibility, that it is mistaken.

It is true, however, as the example I just gave demonstrates, that the new model of health care provision and distribution I am proposing will work more dependably as the community's tax system grows more just. If relatively low-paid workers pay much more than their fair share of taxes, because the tax structure is insufficiently progressive to be fair, then any governmental program that relies on the redistribution of tax proceeds to improve justice for those with scarcely any income at all will be compromised for that reason. It will involve an unjustified transfer to the worst-off group from the almost-worst-off-group. That reflection provides a strong reason why tax reform must be at the centre of any general campaign to improve social justice. It would be ironic and disappointing, however, if the point were stood on its head, and if those who resisted redistribution to the very poor were able to point to imperfections in the tax structure as justification for doing nothing, and retaining their own privileges under the *status quo*.

We must next consider a very different kind of issue, which I must not evade, though my views on that issue, I fear, will disappoint many of you. Suppose that everything I've been describing as possible came to pass. Suppose that, after the right kinds of collective consultation, after meetings and discussions and polls and electronic politics and all the rest, we settled on a particular basic program of medical care that we collectively thought government should, in one way or another, make available to everyone. That basic package, as my earlier argument suggests, will not include some treatment that rich people are now in a position to buy for themselves. I said earlier, for example, that the test I proposed would very likely rule out ultra-expensive marginal diagnostics or extraordinarily costly treatments that have some but very little prospects for success. Some people in Canada and America now have the money to buy health care that would be excluded from the basic package. They have the money to buy a liver transplant when the odds are very small but nevertheless real that the procedure would save their lives. In England, people are standardly denied even renal dialysis on the National Health Service when they are sixty-five. So people of that age die if they cannot afford to pay for dialysis themselves.

If we adopted the kind of scheme that I'm describing, in our admittedly imperfect society, and took no steps to forbid people buying more expensive care than the basic package provides, some people would have better medical care — some people would live longer and healthier lives — only because they had more money. In most cases, since the basic economic structure would continue to be unjust, because they unjustly had more money. Should we therefore take steps to prohibit or constrain the private market in medicine? Should that be part of any respectable campaign to improve justice in health care? Of course, we couldn't actually abolish the private market in health care altogether: we would end by producing back street dialysis. But should we do what we can, aiming to prevent anyone from buying better medical care than the basic package provides, so far as that is possible?

The insulation model of medical justice I began by describing, if taken seriously, would insist that we should, and I believe many people here this eve-

ning would agree. Solidarity is compromised, they think, when some people can live while others die only because the former have more money. That seems to me the wrong answer, however. The spirit of the argument I have been making suggests that no one can complain *on grounds of justice* that he has less of something that someone else does, so long as he has all he would have if society were overall just. And, of course, in the circumstances we are now considering, people whose basic provision does not include liver transplants, and cannot afford to buy such an operation for themselves, are by hypothesis not denied what they would have if economic justice were perfect.

Some of you will hate that argument, as I said: you will think it intolerable. May I remind you, however, that the hypothetical inequality in medical care I'm now considering is, in one important respect, relatively benign compared to other inequalities in our society. If health care were rationed in the way we are contemplating, then everyone would have at least the medical care he would have in a just society, and that would not be true in most other departments of resource allocation. In education, employment, culture, recreation, travel, experience and a host of other goods and opportunities that for most of us make up the value of being alive, the poor would continue to have much less than they would if we had reformed not just health care but our economic and social life more generally. If we somehow manage to succeed in providing the poor with the medical care that justice requires, it would be perverse, given that a rich man can spend on more comfortable housing or better education for his children, not to allow him to spend on more expensive health care. We would do better to put an excise tax on special health care, and use the proceeds of that excise tax to improve public education, or the economic infrastructure, or to reduce public debt that blights employment prospects, or in some other ways that would make the community distinctly more egalitarian.

6.

I will offer you no final summary of my somewhat discursive remarks; I shall try, instead, to broaden the argument in closing it. I began by criticizing the insulation model, as I called it, and you may think I've been undermining that criticism in the last part of the lecture. I've been arguing how we might make our communities better in just one respect, and that goal seems to assume, with the insulation model, that health care is special, "chief" among goods. But my special interest in medical care is largely practical. Medicine is now a problem for people so high up in the economic scale — well up into the middle, fat part of the economic diamond where the votes are. People generally, not just the poor, agree that government should take a larger role in structuring, controlling and financing the provision of health care. We can seize on this opportunity to make the distribution of health care more just as well as more efficient.

But if America does make new progress in that direction, as Canada already has, then the lesson might be of more general political importance. For one thing, it might teach us that the bad press the ideal of equality has had for some time is unjustified. There is a rap against equality: that accepting equality as an ideal, even one among others, means levelling down and requiring every-

one to live the same kind of life. But the conception of equality I've been relying on has quite the opposite character: it is dynamic and sensitive to people's differing convictions about how to live.

I end with this further observation: the question of health-care reform in America, including politically acceptable and fair health-care rationing, is ideologically leveraged. If we find, after all the fuss, that politically we can't do much to make the distribution of medical care more just, in spite of the apparent present opportunities to do so, then a pessimistic conclusion may be irresistible: we may abandon hope for any more widespread or general democratic concern for social justice. But if we do now make substantial and recognizable political progress in this one urgent matter, we may learn more, from the experience, about what justice itself is like, and we might find it to our taste, so that we can steadily, bit by bit, incrementally, fight the same battle in other areas. So the war against injustice in medicine that you have been fighting so well here, and that we are about to take on in America, is indeed a crucial one. Health might not be more important than anything else — but the fight for justice in health might well be.
