Oratie prof.dr. Han Bleichrodt

Faculteit der economische wetenschappen, Erasmus Universiteit Rotterdam

Hierop rust een embargo tot vrijdagmiddag 19 september 2008, 16 uur.

The Value of Health

Ladies and gentlemen, welcome

Usually, an oration marks the beginning of a new professor's career. Because I took some time to give this address, it also marks the end of my appointment at the DSW chair at the Erasmus Medical Centre. For those of you concerned about my well-being: my appointment at the Department of Economics has just been extended so this address still celebrates a new start.

Introduction

Today we will talk about choices, choices about life and death. We do not think about such choices everyday, but there are times where we face them. Today will be one of these.

Suppose your doctor has just told you that you have a kidney failure. There are two treatments available. The first treatment is dialysis. Patients who opt for dialysis live for another 10 years. Half of these patients react well to the dialysis and are almost symptom-free (we call this minor health problems). The other half do not react well and have major health problems. That means you are restricted in your daily activities.

The other treatment is kidney transplant. Of the patients who take a transplant 1 in 10 die within a week after surgery. Of those who survive 3 out of 10 lose the transplanted kidney within the next year and then have major health problems. Those who keep the kidney live for 12.5 years in good health. Which treatment would you choose? Is there anyone who can answer this question now? Most of us, including myself, probably don't.

Governments also face choices. Medical technology changes rapidly and new treatments become available all the time. How does the government decide which treatments to reimburse? Of course, if the new treatment is both cheaper and leads to larger health improvements than what's currently available then the government should go for it. Similarly, if the new treatment is more expensive and gives smaller health improvements than what's available then the treatment should be discarded. But what if the treatment is more expensive but yields larger improvements in health? How does the government decide that these additional benefits are sufficiently in balance with the higher costs?

We, health economists, are increasingly involved in making such choices. In many countries economic evaluations are now required before a new treatment is reimbursed. Is this justified? Are health economists up to their task? Are our methods sufficiently fine to address the important questions that patients and societies face?

My central thesis is that the answer to these question is no. No, the current tools of health economics are unreliable and lead to the wrong treatment recommendations. That is the dark side of my talk. But, it is no reason for despair. The bright side of my talk will show how we can improve our tools and come up with recommendations that better reflect the interests of patients and society. These improvements are based on new insights from theoretical economics. They have not been used in health economics yet. Health economics is an applied field of research and many health economists are not well aware of developments in theory.

But, thanks to the pioneering work of the rector and the previous deans of the departments of economics, Harry Commandeur, and health policy and management, Frans Rutten, we are currently building a health economics group in Rotterdam that is pre-eminently able to bridge the gap between economic theory and applied health economics research. It is my conviction that Rotterdam can become and maybe to some extent already is a world leader in this field.

€80,000 for a life-year

Before I come to my main topic let me sketch a bit of history. About two years ago, the Dutch Council for Public Health and Health Care, the most important advisory

body of the Dutch government on health policy, recommended that only treatments that cost less than €80,000 to let a patient live an extra year in good health should be eligible for reimbursement. The Council is not alone in its advice. In several other countries similar guidelines are common. The most notable case is the UK where such a benefit-cost ratio is required for any treatment that wants to be included in the NHS, the National Health Service.

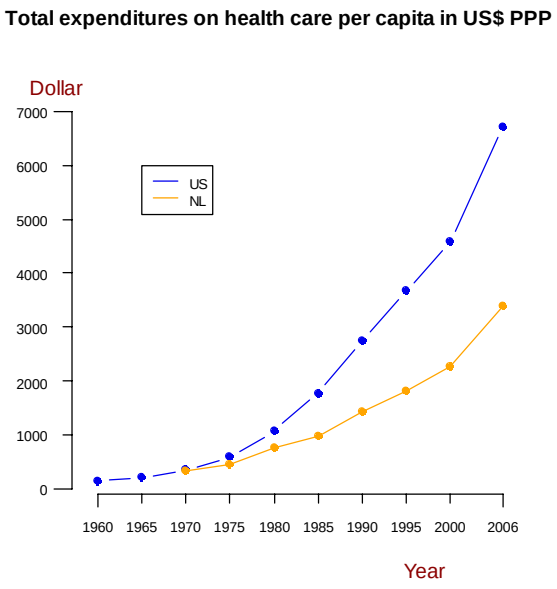

An underlying reason for this recommendation is displayed in the above graph. It depicts the evolution of health care expenditures per head and adjusted for inflation in the Netherlands, the orange line, and the US, the blue line, from the early 1960s until now. In the Netherlands, health care expenditures per head rose from about $300 per head in 1960 to more than $3000 per head now. A tenfold increase. In the US the situation is even more dramatic with a more than twentyfold increase to nearly $7000 per head now. What is more, the increase in health care expenditures seems to accelerate. The health budget will come under increasing pressure in the future due to the rapid development of medical innovations, increasing expectations of what medical care can achieve, and the ageing of society.

It is against this background that the Council came to its advice. The health care budget is limited and we cannot fund everything that is medically possible. Choices have to be made. In making these choices, the Council recommends to perform those activities that maximize the benefits of health care. If we have €200,000 to spend then we better reimburse a treatment that costs €50,000 to gain a year in good health than one that costs €100,000 to gain a year in good health.

Misguided criticism

This may sound obvious, but the Council's recommendation was nevertheless received with much criticism. Most of this made little sense.

The first criticism, voiced in nearly all Dutch newspapers, reflected a misunderstanding. Commentators thought that expensive treatments should no longer be reimbursed. The advice was perceived as a device to cut costs. This is clearly not true. The Council said that medical programs that cost less than €80,000 to gain a year in good health should be reimbursed. The recommendation involves both a cost and a benefit side. A heart transplant, for example, is a very expensive treatment. But, it generates sufficiently many extra years in good health to pass the criterion set by the Council.

A second group of critics came from ethics. They took issue with attaching a monetary value to a year in good health. They argued that human life is priceless and that we can never put a price tag on it. Judging from his recent public statements, this view is shared by the current Dutch secretary of health, Ab Klink. That's too bad. Saying that human life is priceless sounds perhaps appealing at first hearing, but it does not stand scrutiny. It does not provide any guidance when we have to decide which treatments to fund.

Moreover, it conflicts with reality. In our daily lives we often make trade-offs between risks of death and other considerations. Many workers take up risky jobs in exchange for extra pay. If they considered their lives to be infinitely valuable they would never do so. Similarly, most of you probably arrived here by car. If you consider your life infinitely valuable then this was a very unwise decision. I strongly

urge you to return by public transport because this markedly reduces your chance of being involved in a fatal accident. Those of you who came by bicycle probably did not wear a helmet. Irresponsible if your life is infinitely valuable. Just in case you have second thoughts I have one here which you can buy for €100,000. A bargain if your life is infinitely valuable.

The government also assigns a finite value to human life as it does not undertake every road or rail safety program available. Some programs are simply considered too expensive. In the UK and the US the value of life that is used is made explicit. For example, the EPA, the US environmental protection agency, uses a value of life of $6.9 million. A curious detail is that 5 years ago this value was still $7.9 million. Apparently living in America has become less valuable under George W. Bush.

To the best of my knowledge no explicit value of life is used by the Dutch government. This is unfortunate because using such a value makes government policy much more transparent.

A third criticism, coming from political scientists, was that reimbursement decisions should not be taken by economists but are better left to politicians. Apparently the argument is that the political process will ensure that everyone can have their say and that this will lead to better decisions. I don't think the Council had in mind to leave health policy to economists. I agree that such supremacy would be undesirable. But, economists can give refreshing insights about the way public money is spent. Let me give an example.

In the 1990s researchers from Harvard University compared 500 interventions to save lives. They divided these interventions into three categories: medical interventions, injury reduction, and environmental measures. If public money were effectively spent then these different interventions should gain life-years at basically similar costs.

Table 1: Median Costs per Life-Year Saved for Three Types of Intervention (in 2007 dollars)

Type of Intervention	Costs/Life-Year
Medicine	$27,000

Injury Reduction	$68,000
Environmental	$4,000,000

This was clearly not the case. Overall, medical interventions offer the best value for money: they cost $27,000 per life-year saved. Some measures save life-years at very low costs or even save costs. For example, methadone provision for heroin addicts saves more than it costs, mainly through prevented crime. Injury reduction is somewhat less cost-effective, saving life-years at an overall cost of $68,000. The price varies wildly, however. Seat belt laws cost only $98 for a life-year saved. Strengthening buildings in earthquake-prone areas on the other hand costs as much as $18 million dollars.

By far the most expensive are environmental measures. They save life-years at a cost of 4 million dollar. There are of course exceptions. The reduction of lead in gasoline saved more than it cost. However, in general environmental measures are expensive. For the rector, it is perhaps interesting to know that what it costs to save a life-year through removing asbestos in roof coatings: more than $7 million. Very cost-ineffective. So whether the reconstruction of the aula of this university was really worth the effort….

You may be inclined to believe that these data are specific for the US, which is in many ways special indeed. Your suspicions may be raised further when I tell you that one of the authors of this study later went on to become a chief advisor of the current president of the USA, who is indeed not known for his environmental credentials. Nevertheless, studies in Sweden and the UK showed similar patterns. These data suggest that public money is not always effectively spent: by shifting money, for example from pollution control to medical interventions, we can save extra lives at no extra cost. We have to keep in mind that many interventions have benefits other than survival. What these data show, however, is that politics can benefit from the insights of science. Politicians who are deaf to science waste public money.

Is health economics up to its task?

What I have said thus far hopefully justifies that the advice of the Council is sensible. Economic evaluations can help to make health policy more transparent. The recommendation can only work, however, if we, health economists, are up to our task. A critical assumption is that we have the tools to compare the benefits and costs of health care programs in a reliable manner. In particular, it assumes that we can reliably measure the value of health. As I already mentioned, I do not believe that we have these tools. Our current tools are flawed and unreliable.

Current practice

To make my point, let us return to the treatment decision I started my talk with, the choice between dialysis and kidney transplant. It's a difficult decision because we must balance different life durations and complex and unfamiliar health states against each other and give due weight to the risks involved. To analyze the decision we must have a procedure, or model as economists call it, to combine these different aspects of the decision. The model that is commonly assumed is called expected utility.

To use expected utility we must know what health is worth to you. We need to know the value of health. How do we determine this? We could of course look at life duration only. Then dialysis gives 10 extra years for sure and kidney transplant gives either 12.5 extra years, or 10 extra years or nothing at all. But looking at life duration alone does not tell the full story. Not all life-years are spent in the same health quality. Life-years in good health count for more than life-years with major health problems.

To capture this, we adjust life-years for the quality of life in which they are spent. The resulting number is called a *quality-adjusted life-year* or QALY for short. QALYs are easy to interpret. One QALY corresponds with one year in good health. There are many ways in which we can gain one QALY. It can mean, for example, one more year in good health, or two more years in health of only 50%, or 4 more years in 25%.

A cause of concern is to what extent QALYs reflect how people perceive health. When I use QALYs to value the health states does this indeed capture how you feel about health? Can your feelings be captured by such a simple formula or are they

more complex? Many health economists believed that QALYs were too simple and new measures of health have been proposed, most notably the so-called healthy-years equivalents. In the end these new measures did not stand scrutiny. The new developments in theoretical economics that I will discuss shortly have also led to a reconsideration of QALYs. I believe that QALYs are fairly accurate as a reflection of how people feel about health. They could be improved somewhat but for all practical purposes they perform well enough.

Problems

If QALYs are basically fine and we want to work with them, one question remains. To use QALYs we must measure quality of life. How do do that? Suppose I want to measure how many QALYs one year with, say, major health problems is worth to you. The recommended method is to give you two options. Option 1 is the certainty of staying with major health problems for the rest of your life. This means that you have difficulty walking, are unable to participate in most leisure activities, and have moderate pain.

Option 2 is to undergo a treatment that can either succeed, in which case you live the rest of your life in good health, or fail, in which case you die within a week. I then ask you which risk of death you are willing to take. 0%, 5%, 10%, 20% or even 50%? Please give it some serious thought. Just for 20 seconds. Which is the maximum risk of death that you are willing to take to undergo the treatment?

In real data this maximum risk of death is about 10%. If the risk of death is less than 10%, we prefer the treatment, if it exceeds 10% we prefer to live with major health problems. Under expected utility, this means that our value of major health problems is 90% of the value of good health. So we think major health problems are not very serious. Their value is almost equal to good health. If we have them, we're quite all right and a health policy maker should not care too much about us.

This procedure is known as the *standard gamble*, an unfortunate name because you probably link the term gamble more with a casino than with a medical treatment. In this context, the term gamble just means that the treatment is risky and offers several

possible outcomes. Because the term is so common I will continue to use it even though it is not well-chosen.

Is the standard gamble right? Does it measure quality of life correctly? In the above example, is our value of major health problems indeed 90% of the value of good health? Before I answer that question let me again sketch a bit of history.

Expected utility was introduced in the 18th century by a Swiss mathematician Daniel Bernoulli. The theory has a little Dutch flavour to it as Bernoulli was born in Groningen. Bernoulli proposed expected utility as an explanation for a behavioural pattern that he observed. He did not give any reasons why people should behave according to expected utility. It took more than two centuries for this to be established by a mathematician John von Neumann and an economist Oskar Morgenstern.

Von Neumann was a brilliant man, what Einstein was for physics he was for mathematics. He has also affected your life as he is one of the developers of the computer. Von Neumann was a very colourful personality who was known for the wild parties he gave at least twice a week. He was a very bad driver, mainly because he used to read books while driving. And he was an insatiate womanizer. In Los Alamos where he worked on the atomic bomb during World War II, the secretaries covered the open space under their desks with cardboard to stop him staring at their legs. No comparable stories circulate about Morgenstern, but then he was an economist.

von Neumann and Morgenstern clarified why it makes sense for people to behave according to expected utility. But, that it makes sense for people to behave according to expected utility does not mean that they really do. Let us look again at the standard gamble.

We found that we are willing to take a 10% risk of death to avoid major health problems. I told you that this means that our value of major health problems is 90% of the value of good health, so almost perfectly healthy. A policy maker deciding how to spend the money should not give much priority to us. Treating us only leads to a small gain in health.

But is this really true? Can we deduce from our answer that the value of major health problems is really as much as 90% of the value of good health? I don't think so. We believe major health problems are serious and treating these will benefit us a lot. However, we are not willing to run a large risk of death either. What is wrong is the conclusion that a small risk of death translates into a large value of major health problems. In other words, the standard gamble is wrong. It overestimates the value of major health problems. Consequently, benefit-cost analyses based on the standard gamble underestimate the burden imposed by major health problems. Consequently, the number of QALYs that can be gained by treating these health problems will be underestimated and too little will be spent on treating major health problems.

"Solutions"

What now? How to solve this problem? We cannot use the standard gamble because it gives the wrong values of health. But what then? One line of rescue comes from psychology. It says that most people do not have well-defined preferences for health. When you answered the standard gamble, you were asked to make a comparison to which you probably have not given much, if any, thought before. You did not have a ready-made answer. Instead you were probably searching for a way to come up with a plausible answer within 20 seconds. Your response was the sum of two terms: your genuine preferences, which we want to know, and a term that reflects your response strategy. To arrive at your true value of major health problems we should try to eliminate the second term. This is not easy, however, and requires much effort. In economic evaluations, we generally do not have the time and resources to do so. In theory this psychological approach is the best solution to obtain reliable values of health. In the practice of economic evaluation, however, other ways must be found to solve our problem.

A second possibility is to discard the standard gamble altogether and to use other methods. This is a common strategy in health economics. For the initiated, I am thinking here about methods such as the time trade-off and the visual analogue scale. Unfortunately these methods have their own problems. They are at best ad-hoc solutions to the measurement problems we face, but, in my view, they cannot be taken as the final answer.

That leaves a third option, which I favour. That is to accept the answers as they are and to make the best of them. We must try to find a way to derive more plausible values for health from the standard gamble. To do so we must abandon expected utility. We need a theory that better explains how people answered the standard gamble. It is here that we apply the new insights from economics and psychology.

The solution is prospect theory. Prospect theory was developed by two psychologists Daniel Kahneman and Amos Tversky. It is the most important theory of decision making in economics today. In 2002 Daniel Kahneman received the Nobel prize in economics for the introduction of prospect theory and I am happy to let you know that next year he will accept an honorary doctorate from this university. Unfortunately Amos Tversky died in 1996 and he could not share this honor as the Nobel prize is not awarded posthumously. Tversky was not only a brilliant scientist, but also a decorated war hero, a rare honour amongst scientists. His fame came for saving the life of a fellow soldier. The soldier froze in panic after placing an explosive charge, literally lying on top of the explosive. Tversky, then 19 years old, knew the explosion would occur within a few seconds. Nevertheless, in a split second he took the gamble and risked his life. He ran to the soldier, picked him up and threw him to safety, only to be wounded himself. For this display of courage, Tversky earned Israel's highest military decoration.

To understand how prospect theory modifies expected utility, let us look one final time at the standard gamble question. Under expected utility, the value of major health problems is 90% of the value of good health. Prospect theory disputes this. It argues that expected utility has two important problems. First, suppose I ask you how much are you willing to pay to reduce the risk of a cancer, from 1% to zero. Is that the same amount that you are willing to pay to reduce the risk from 5% to 4%? Expected utility says yes, but it appears that the two reductions are not the same. It's a different choice. Most people pay more for the former than for the latter. We find completely eliminating risk very comforting. We want to feel safe.

Prospect theory has another key insight. We want a gain but we certainly do not want to lose. So if we can win €100 when the head of a coin comes up but lose €100 when the tail comes up most of us are not willing to play. We don't want to lose.

These two insights imply that the standard gamble overestimates the value of health. Consider the standard gamble question. We don't want to lose so we don't want to run a large risk of death. Under our current tools we conclude that our health problems are not serious. Prospect theory gives us a better picture of reality. It says that our health problems are much more serious and that their value is less than 90% of the value of good health.

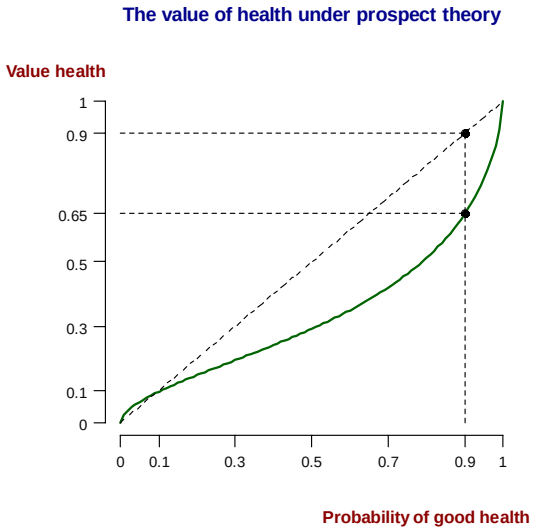

How wrong is expected utility? What's the difference between the value of health under expected utility and under prospect theory? The above figure shows this. It shows for each answer to the standard gamble the implied value under prospect theory. The diagonal shows the value under expected utility. The green curve, showing the values under prospect theory, lies almost everywhere under the diagonal. In other words, the standard gamble leads to an overestimation of the value of health. Moreover, the overestimation is substantial. In our standard gamble question we concluded that the value of major health problems was 90% of the value of good

health. Under prospect theory this value is only 65% of the value of good health. A huge difference and much more realistic.

You may retort “but is what you tell us true? Does prospect theory indeed lead to better values for health?” A legitimate question, indeed. To answer it I will show you some data that will give substance to my claim that the current tools of economic evaluation are wrong and can be improved by prospect theory.

In a study with Jose Luis Pinto and Peter Wakker we collected data for three methods that should give the same values for health when expected utility is true.

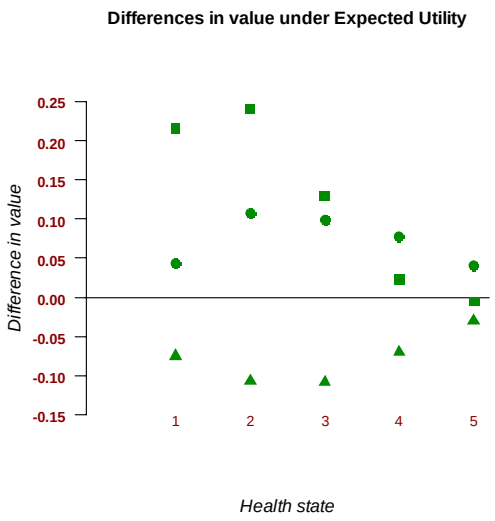

The figure shows the differences between these methods under expected utility. The circles indicate the differences in health state values between the first and the second method, the squares the differences between the first and the third method and the triangles the differences between the second and the third method. If expected utility were true there should be no differences. All points should be located on the horizontal axis. This is clearly not the case. The points are all over the place and expected utility is clearly wrong.

Now let us see what happens when we analyze these data under prospect theory. The blue circles, squares and triangles show the differences under prospect theory. Under prospect theory, all differences go to zero. All points are close to the horizontal axis. Prospect theory indeed performs much better than expected utility.

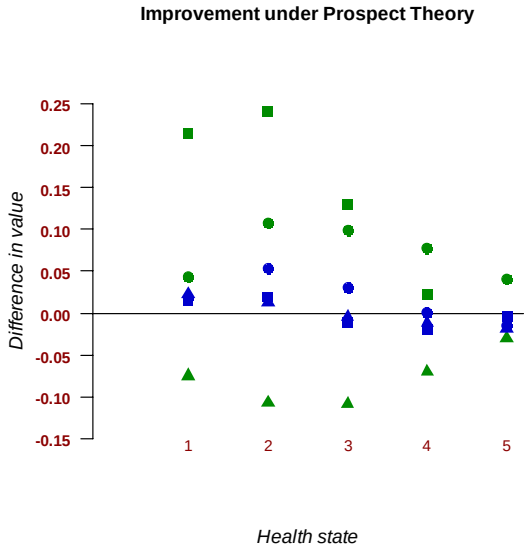

Conclusion

Let me conclude. Economic evaluations of health care can help to make better medical decisions. Decisions about life and death. Our current methods are wrong. Reality is different from what we did believe. Policy decisions based on our current tools are not at all in the best interests of patients. We, scientists, have the responsibility to bridge the gap between theory and reality. Prospect theory respects reality. Now, we have all the opportunity to reliably value health. And, these values can even be obtained free of charge: we do not have to collect additional data. We can apply prospect theory immediately. What remains is to spread the word. I hope I have convinced you.

Thank you.