



Eli Hirsch

The Rational Physician

Eli Hirsch

As a philosopher I have a general interest in the concept of rationality, but my specific concern with the topic of the present paper is shared with any member of our culture who must inevitably rely at some critical moments on the judgments and advice of physicians. I was excited several years ago to pick up a book called *Clinical Decision Analysis*, authored by a group of professors of medicine and health administration, which attempts to lay out a detailed theory of how physicians can improve the rationality of their clinical decisions¹. These clinical decision analysts—as I will call the authors of this book—are explicitly promoting an agenda for altering medical education in the direction of teaching future doctors how to apply within a medical context the basic ideas of decision theory, a theory that has already achieved great prestige and influence in business and economics departments. After studying this book I ultimately came to feel that the agenda promoted in it is deeply flawed and potentially unethical. That is what I want to discuss in this paper today.

Decision theory is often presented in a highly technical and mathematical guise, but my discussion of this theory will be completely non-technical and will only occasionally appeal to an elementary level of mathematics.

The following is the sort of example that the clinical decision analysts would like physicians to learn how to deal with more rationally.

The case of vascular insufficiency in a diabetic patient. After a foot injury the patient has developed an infection and possibly gangrene. One option is to amputate immediately, with a one percent probability of death during surgery. The other option is to wait to see if the foot heals with medical care. There is a 70 percent probability that this will work, but a 30 percent probability that gangrene will spread. If the latter occurs then more extensive and disabling surgery will be required, with a ten percent probability of death during surgery.

I am not going to raise any worries about the fact that the probabilities in this example are evidently idealizations. I accept the argument of the clinical decision analysts that idealizations—estimates, approximations—are better than nothing. Let us assume, then, that we are working with these particular probabilities. The difficult problem is what to do with them. If we wait we increase the patient's chances of either dying or winding up with a more serious disability; however, we thereby give the patient a significant chance of getting away without any amputation. Should we wait or not?

¹Milton C. Weinstein, Harvey V. Fineberg, et al., *Clinical Decision Analysis* (W.B. Saunders Co., 1980).

Over the years I have put this question to hundreds of premed students. Some of them say that we must have immediate surgery in order to minimize the probability of death. Others say that we must wait in order to give the patient a chance of not needing an amputation. Most of them, after a bit of reflection, give up, and say that they would have no idea how to advise a patient in this kind of case, that it would have to be left up to the patient. But I point out to them that if you are a physician who has been treating this patient for many years, advising him about numerous medical issues, he may have every reason to feel that you should be able to help him to reach a decision in this most critical situation.

Of course the decision must be in some sense up to the patient. I don't merely mean that the patient has the legal and moral right to decide, but that the patient's distinctive preferences and concerns will play some role in determining what decision he makes. A professional athlete may reasonably care a lot more about losing a leg than someone in another profession does. There may even be a range of feelings about how bad it is to die, and these diverse feelings may strike us as perfectly rational. Now decision theory purports to give us a way to factor in the distinctive cares and feelings a patient has about disability and death, and then, on that basis, to be able to advise the patient what it is most reasonable for him to do. The physician can therefore respect each patient's individuality while not simply abandoning the patient to make the decision on his own. This does sound like a salutary goal. But can it really be achieved?

The first step in applying decision theory is to assign numerical measures to the patient's attitudes about disability and death. In our example we will assume that the best outcome from the patient's point of view is recovery without amputation, and the worst outcome is death. We may conventionally assign the measures 1 and 0, respectively, to these outcomes. To simplify slightly, we can assume that the only other relevant outcomes are two: surviving with the less extensive amputation or surviving with the more extensive amputation. Each of these intermediary outcomes is to be assigned a number between 1 and 0. We need to find out how close to 1 or 0 these outcomes are in the patient's perspective. Merely to ask the patient to assign numbers to these outcomes would seem futile, for who would know how to assign such numbers? The trick is to get the patient to implicitly reveal enough about his attitudes for us to assign the proper numbers. The method used is to ask the patient to rank each intermediary outcome relative to a gamble in which there is a probability of getting the best outcome and a probability of getting the worst outcome. According to the clinical decision analysts the physician (or some member of his staff) should attempt to interview the patient and put to him, perhaps in some more elaborate form, the following series of questions: "Which would you choose if you could either get the certainty of surviving with the less extensive amputation (an intermediary outcome) or get a lottery—a gamble—in which there is a 60% chance of recovery without amputation (the best outcome) and a 40% chance of death (the worst outcome)? How

about if the lottery gave you a 65% chance of recovery? A 70% chance? ... How about a 99.9% chance?" Suppose that this sort of interview eventually reveals that for this patient the outcome of surviving with the less extensive amputation seems preferable to the gamble only when the latter provides less than a 99% chance of recovery without amputation. We then assign the measure .99 to that outcome. This represents a measure of the value or utility of this outcome for this particular patient. By repeating this process we assign a utility measure to the other outcome of survival with more extensive amputation.

Armed with these utility measures the clinical decision analysts believe we can advise the patient whether he ought to have immediate surgery or wait. The most fundamental commandment of decision theory is to maximize "expected utility." Expected utility is something like an "average utility," which takes into account both the utility measure and the probability of each outcome. We calculate an expected utility of a course of action by multiplying the probability of each outcome by its utility and then summing all of these probabilities-times-utilities together. Knowing, as we now supposedly do, the utility for this patient of each of the four possible outcomes of either having immediate surgery or waiting, we can easily calculate which course of action has the higher expected utility. That is the action that the patient is rationally obliged to do.

Here again I am not going to worry about the fact that we are evidently dealing with idealizations. Of course it will typically be difficult—and sometimes positively harmful—to conduct the kind of interview with a patient that I have just sketched. Nevertheless, at least a partial interview of this general sort may often be quite feasible, and may suffice to establish an approximate utility measure for the patient. Approximate probabilities and utilities may be better than nothing. The clinical decision analysts seem, then, to offer an elegant solution—at least ideally and in principle—to what is unquestionably a major dilemma for both physicians and patients. The question is whether the solution really makes sense.

Let me introduce a pair of examples that are somewhat less complicated than the example of vascular insufficiency that I have been discussing. These new examples will allow me to illustrate more concretely how decision theory works while at the same time exposing a major problem for the theory.²

Imagine first a case in which a patient has a disease of the brain that will certainly be fatal unless it is treated. There are two treatment options available which we will call *A1* and *B1*. We assume that these options are alike in discomfort, expense, and every other relevant factor, except for their differing

²My example is an adaptation of a monetary example given in Amos Tversky and Daniel Kahneman, "The Framing of Decisions and the Psychology of Choice," *Science*, Vol. 211, January 1981, 453-211, at 455. The example is a variation of "Allais's paradox." The paradox is discussed throughout the decision theoretic literature, but seems to be ignored by the clinical decision analysts.

probabilistic consequences. If option *A1* is chosen the patient will certainly survive but with a fairly serious speech impairment. On option *B1* there is a 20 percent chance of death, but if the patient survives it will be without any impairment.

We can summarize this example, which I will call Case I, in the following table.

Case I

A1 Certainty of survival with speech impairment.

B1 0.2 probability of death; 0.8 probability of survival without impairment.

As in the previous example, it seems clear that the patient's distinctive attitudes about death and the impairment will be relevant in determining what decision the patient is likely to make. Let us now compare Case I to a second case. Here again the patient has a disease of the brain that will certainly be fatal unless it is treated. The prognosis in this case, however, is more grim, for the only two treatment options available do not promise a high probability of saving the patient. The relevant facts about these options are as follows.

Case II

A2 0.75 probability of death; 0.25 probability of survival with speech impairment.

B2 0.8 probability of death; 0.2 probability of survival without impairment.

Everyone will agree that in either case both options might be rational, depending on the precise nature of the impairment and how the patient feels about it. It seems obvious, however, and is well documented in the literature, that many people will be more attracted to the *A*-option in Case I than in Case II. The rationale for this attitude might be expressed as follows: "In Case I the *A*-option gives me the certainty of survival, whereas the *B*-option gives me a significant possibility of dying. I'm not going to risk my life; I'll take the *A*-option and deal with the impairment. In Case II, on the other hand, I'm probably going to die no matter what I do, and the probability of dying is only a little more in *B2* than in *A2*. I may as well go for it, risking my life a little more in the hope that I'll come out without any impairment." In Case I the idea is not to risk one's life unnecessarily; in Case II, where a serious risk to one's life can't be avoided, the temptation is to "go for it." Certainly this will not immediately strike anyone as intuitively irrational.

The surprising fact, however, is that decision theory rules out choosing the *A*-option in Case I and the *B*-option in Case II. This can easily be seen if one considers that someone's reaction to the choice in Case I between option *A1* and option *B1* is precisely the sort of reaction we were trying to elicit in the interview sketched earlier designed to define the utility measures for different outcomes. If someone chooses *A1* over *B1* this implies that for this person the

utility of surviving with the impairment is more than 0.8. A trivial calculation then shows that the expected utility of A_2 is greater than the expected utility of B_2 . The more general result is that the expected utility of A_1 can be greater than the expected utility of B_1 if and only if the expected utility of A_2 is greater than the expected utility of B_2 . Since decision theory demands that we always maximize expected utility, decision theory forbids combining the A -choice in one case with the B -choice in the other case.

Suppose that a patient confronting Case II is uncertain what to choose. If the clinical decision analysts had their way the physician would ask this patient what choice he would make if he had been confronting instead Case I. Case I is in a way simpler than Case II (for it involves fewer probabilities), so we can well imagine that the patient feels quite definite about making the A -choice in that case. The physician will then inform the patient that he ought to make the A -choice in Case II, the case he is actually confronting. If the patient eventually comes to feel that he would actually prefer the B -choice in Case II, even though he remains attracted to the A -choice in the hypothetical Case I, the physician would try to persuade him that this is irrational.

I would hope that at this stage of the discussion you will want to say, "So much the worse for decision theory." Where, you may well ask, is the commandment written that one has to maximize expected utility? Since it seems perfectly okay to make the A -choice in Case I together with the B -choice in Case II, what we see, apparently, is that there is no rational demand in general to maximize expected utility.

But things are not so simple. There is in fact an extremely straightforward and powerful argument in favor of the decision theorist's point of view. I want now to present the essentials of that argument, and then to explain why I reject it. This argument has been formulated in many different ways, but at its center there is always a certain principle that does the main work. I will call this the "master principle" and I will formulate it in a somewhat rough manner that I hope will convey its essential intuitive content.

The master principle. If you prefer X to Y then you ought to prefer a probability of getting X to the same probability of getting Y (everything else being equal).³

Let us first consider a completely trivial application of the master principle. Suppose we are going out to dinner and you tell me that you prefer Chinese

³Well known versions of the master principle include the "sure thing principle" in Leonard Savage, *The Foundations of Statistics*, second edition (Dover, New York:1972), pp. 21-22; the "substitution principle" in Howard Raiffa, *Decision Analysis* (Addison-Wesley, 1970), p. 62; the "best prize principle" in Michael Resnick, *Choices* (University of Minnesota Press, 1987), pp. 91-92; and the "independence axiom" in Edward F. McLennan, "Sure-thing Doubts" in P. Gardenfors and Nils-Eric Sahlin, ed., *Decision, Probability, and Utility* (Cambridge University Press, 1988), p. 167, and in Mark J. Machina, "Dynamic Consistency and Non-expected Utility" in M. O. L. Bacharach and S. L. Hurley, ed. *Foundations of Decision Theory* (Basil Blackwell, 1991), p. 47.

food to Italian food. There are, I inform you, two restaurants of these sorts in the area, alike in all relevant respects except their fare, but, I add, given how late in the evening it is there is about a 20 percent chance that the restaurants are already closed. "Oh," say you, "then let's head for the Italian restaurant." "But I thought you said you prefer Chinese food." "Yes," you reply, "I prefer the certainty of Chinese food to the certainty of Italian food, but I prefer an 80% chance of getting Italian food to an 80% chance of getting Chinese food." That attitude would seem completely absurd. In this kind of example the master principle surely seems right: If you prefer Chinese food to Italian food then you ought to prefer an 80% chance of getting Chinese food to an 80% chance of getting Italian food.

The remarkable fact, now, is that if we apply the master principle to Cases I and II we get the result that if you prefer the A-option in Case I then you ought to prefer the A-option in Case II. To see this one only has to realize that a 25% chance of getting an 80% chance of getting something is equivalent to a 20% chance of getting that thing. If I tell you that you have a 25% chance of being chosen as one of the candidates for a job, and if you are chosen as a candidate you then have an 80% chance of getting the job, I have told you that you have a 20% chance of getting the job. Now if you look at the relationship between B2 and B1 you see that the 20% chance of surviving with no impairment provided by B2 is 25% of the 80% chance of surviving with no impairment provided by B1. And if you look at the relationship between A2 and A1 it's immediately clear that the 25% chance of surviving with the impairment provided by A2 is trivially 25% of the 100% chance of surviving with the impairment provided by A1. But the master principle implies that if you prefer a 100% chance of surviving with the impairment to a 80% chance of surviving with no impairment, then you must prefer a 25% chance of surviving with the impairment to a 20% (equals 25% of 80%) chance of surviving with no impairment (everything else being equal, in the sense that the only alternative to surviving with or without the impairment is death). In other words, the master principle implies that if you prefer A1 to B1 you must prefer A2 to B2.

The argument that I have just gone through shows that, in cases I and II, if you do not maximize expected utility—which means that you choose the A-option in one case and the B-option in the other case—you will violate the master principle. To generalize from these cases to all other cases would require a bit of technical work that I will not go into. The thrust of decision theory is, however, this central idea: You are rationally required to maximize expected utility because you will otherwise violate the master principle. This is unquestionably a brilliant and powerful argument, and the seminal decision theorists who first thought it up deserved all the accolades they received. But have they really given us the right advice about how to make decisions? I don't think so. Let me try to explain why.

When we focus on cases I and II our strong intuitive feeling is that there is nothing irrational about making the A-choice in case I and the B-choice in case II. (Maybe the opposite choices could also be rational, but that seems

harder to explain, so I will put it aside.) On the other hand, when we look at the master principle our immediate intuitive reaction is that it seems right. What we have here is a dilemma that is typical in philosophy: our intuitions about a particular case run counter to our intuitions about a general principle. There are, roughly speaking, two methodological stances that might be taken towards this kind of dilemma, that of the "particularist" and that of the "generalist." Wittgenstein made the remark: "Don't think, but look!"⁴ He meant: "Don't rely on general principles, but look at particular cases." I think that Wittgenstein's particularist methodology is the right one, at least in most cases. When general principles clash with particular cases, typically the principles ought to yield. That is what is meant by giving a "counterexample" to a generality. In the face of a counterexample one typically needs to modify the generality, not hold on to the generality and reject the example.

Generalists don't see it that way. This difference of approach goes back to the dawn of philosophy. Zeno's paradoxes, propounded before the time of Plato, raised a number of problems about how motion is possible. One of the simplest paradoxes goes like this: To move from one side of a stadium to the other side, you have to first pass through one half; then you have to pass through half of what's left; then you have to pass through half of what's left; then you have to pass through half of what's left; and so on without end. So before you get to the end of the stadium you always have to first do something else. So you can never get to the end of the stadium. From this kind of paradox Zeno drew the conclusion that time and motion are delusions. The paradox depended on a number of general principles, for instance, the principle that you can't do an infinite number of tasks in a finite time. Certainly that principle must have seemed highly intuitive to people in Zeno's pre-calculus days; it may even seem pretty intuitive today. But I think that reasonable particularists always assumed that the principle cannot be quite right, since we obviously do move from one end of the stadium to the other. When my two daughters were each about 10 years old I presented them with the paradox of the stadium, which they were well able to understand but of course could not answer. To my relief they both reacted as strong particularists, who continued to confidently believe in the reality of motion—and who have in fact treated all of my subsequent philosophical questions with utter disdain.

Generalists have always been popular among philosophers. A number of years ago I attended a post-doctoral seminar conducted at Brown University by the eminent philosopher Roderick Chisholm. He would present arguments in the following vein. If the object *C* is made up of the two objects *A* and *B* then *C* cannot get bigger unless either *A* or *B* gets bigger. Sounds like a good principle, right? Well, then it follows that if you have a two-story house you cannot make it bigger by adding a third story. Chisholm, who was a generalist metaphysician to the core, accepted this conclusion, but as a particularist I

⁴Ludwig Wittgenstein, *Philosophical Investigations* (Macmillan, New York, 1953), p. 66.

consider it to be preposterous. I would say that the principle on which Chisholm's argument was based evidently has to be revised in some way.

I think it's fairly obvious that particularism is closely related to a standpoint that is sometimes called "common sense philosophy." It seems virtually true by definition that what is more general, more abstract, more theoretical, is less commonsensical. Generalists are naturally attracted to highfalutin intellectual exercises, whereas particularists tend to be drawn more to pedestrian down-to-earth judgments of the sort made by ordinary people rather than professional intellectuals. There is a connection here, also, to Pascal's famous remark that "the heart has its reasons that reason knoweth not." Admittedly, there may be little connection between "the heart" in any ordinary connotation and the judgment that one's house can be made bigger by adding a story, but "the heart" does seem to have very much to do with one's judgment that it makes sense to save the certainty of one's life with the *A*-choice in case I and to go for broke with the *B*-choice in case II.⁵

So part of my reaction to the decision theorist's argument is to say, "Don't think, look." There must be something wrong with the master principle, because if you look at particular cases, such as cases I and II, you find that the principle gives the wrong answer. I would say this even if I had no way to explain what is wrong with the principle (just as my daughters continued to believe in motion even though they couldn't say what was wrong with Zeno's principles). But I think that I can explain what is wrong with the master principle.

To give this explanation I need to distinguish between two kinds of choices. Sometimes a choice involves what I will call a "conflict of risk." All of the medical examples that we have so far considered involve conflicts of risk. In all of these examples one option is attractive in that it is more likely to achieve at least a certain relatively low level of benefit, but the other option is attractive in that it holds out the greater promise of achieving a relatively higher level of benefit. Consider the initial example of the patient with an infected foot. Immediate surgery yields the higher likelihood of at least surviving without extensive amputation; that speaks in favor of immediate surgery. But waiting holds out the possibility of surviving with no amputation at all. As another example of a conflict of risk, look at the alternatives in case II. *A2* is more likely to achieve at least the benefit of survival, but *B2* holds out the possibility of the higher benefit of no impairment.

Many choices do not involve a conflict of risk. If I am asked to choose

⁵I should add that "particularism" is itself a generality that, as a particularist, I will want to examine on a case by case basis. There may be cases (especially in Ethics, perhaps) where the intuitive force of a general principle seems powerful enough to overcome conflicting intuitions at the level of particular examples. The distinction between "particularism" and "generalism" can no doubt be further refined, but I think it captures a deep division of philosophical tendencies. Compare with the comments on particularism and generalism in Ned Markosian, "Brutal Composition," *Philosophical Studies*, Vol. 92, 1998, 211-49, at 236-7.

between going to a Chinese restaurant and going to an Italian restaurant I am not generally faced with a situation in which one alternative is more likely to secure at least a low level of benefit but the other alternative promises the possibility of a greater benefit. Here I am simply making a choice based on which alternative seems more likely to benefit me more. This choice may involve a conflict of some sort but not what I am calling a conflict of risk. I may like the food more in the Chinese restaurant but I may know that I could pay less in the other restaurant. Then I have to try to do a kind of cost-benefit analysis, judging whether Chinese food at price x is worth more or less to me than Italian food at price y . This kind of cost-benefit conflict, however, has nothing to do with *risk* as such. It is not a conflict in which I am weighing the probability of one level of benefit against another probability of another level of benefit.

Now the basic suggestion that I want to make is that the master principle applies only where there is no conflict of risk. If you prefer Chinese food at price x to Italian food at price y then certainly you ought to prefer an 80% chance of getting Chinese food at price x to an 80% chance of getting Italian food at price y . In this kind of case, where you are simply making a judgment as to which alternative is probably more advantageous all things considered, the master principle holds. The error of decision theory is to try to apply the master principle to situations of risk-conflict, situations such as cases I and II.

To get a better grasp of the distinction I'm making let's look at another medical example. In this example there are several alternative treatment options involving different probabilities of the outcomes that figure in cases I and II.

Case III

- A3. 0.4 probability of death; 0.5 probability of survival with impairment; 0.1 probability of survival with no impairment.
- B3. 0.1 probability of death; 0.4 probability of survival With impairment; 0.5 probability of survival with no impairment.
- C3. 0.2 probability of death; 0.1 probability of survival with impairment; 0.7 probability of survival with no impairment.

In the choice between A3 and B3 there is no conflict of risk. B3 is clearly the option that is more likely to be advantageous at every level of benefit. B3 improves the likelihood of evading death and also improves the likelihood of survival with no impairment. One may be momentarily confused by the fact that A3 has the higher probability of the intermediary outcome "survival with impairment." But a little reflection shows that this is no advantage in favor of A3. What we need to look at are *levels* of benefit. In this example there are two levels: the lower level is evading death, that is, surviving either with or without the impairment; the higher level is surviving with no impairment. B3 offers the higher probability for both levels of benefit. So we can simply say about B3: "This is the course of action that is more likely to be advantageous." Whenever

we can say that it means we do not have a conflict of risk. The point here is not that we have no risk; whenever we are dealing with probabilities there are risks. *B3* may in fact kill the patient. But there is no "conflict of risk" because, at each level of benefit, *B3* is more likely to yield that level.

Contrast, now, the choice between *A3* and *B3* with the choice between *B3* and *C3*. In the latter choice we do have a conflict of risk. *B3* offers a 90% chance of at least achieving the benefit of evading death, whereas *C3* offers only an 80% chance of achieving that level of benefit. On the other hand, *C3* offers a 70% chance of achieving the maximum benefit of no impairment, whereas *B3* offers only a 50% chance of that higher benefit. Here we cannot simply say of either option, "This is the course of action that is more likely to be advantageous," since that depends on what level of benefit one has in mind.

In the standard literature there is, to my knowledge, no general name for the cases that I am calling "risk conflict," but there is a name for the other cases: "dominance".⁶ One alternative is said to (probabilistically) *dominate* a second when the first is more likely than the second to achieve a certain level of benefit, and the second is not more likely than the first to achieve any level of benefit. *B3* dominates *A3* because the former is more likely than the latter to achieve either level of benefit. If two treatment options offered the same likelihood of evading death but one offered a greater likelihood of no impairment, then again the second would dominate the first. But neither *B3* nor *C3* dominates the other, since one offers a greater likelihood of evading death while the other offers a greater likelihood of no impairment. Here, then, we have a conflict of risk.

My suggestion, then, is that the master principle is valid only with respect to cases of dominance, not with respect to cases of risk-conflict. Why am I saying this? It goes back again to the particularist methodology of "looking at examples." When we look at examples like cases I and II we find that it is not irrational to have a combination of preferences that violates the master principle. But when we look at examples like the Chinese-versus-Italian-food case we see that the master principle obviously does apply. I'm trying to explain what the relevant difference is between these cases. The reason why the master principle applies in cases of dominance is very simple. Suppose that alternative *X* dominates alternative *Y*, and one situation offers a certain probability of getting *X* while a second situation offers the same probability of getting *Y* (everything else being equal). Then it immediately follows that the first situation dominates the second. In other words, the very reason you had for favoring *X* to *Y* is a reason for favoring the first situation to the second. In such cases of dominance the master principle is utterly trivial. To put it another way: If your reason for favoring *X* to *Y* is simply that *X* is more likely than *Y* to be advantageous, then obviously you have that same reason for favoring a certain

⁶Several notions of "dominance" are defined in the literature. The notion I am introducing here is "probabilistic dominance," as in Weinstein, et.al., op. cit. p. 214, and Raiffa, op. cit., p. 245.

probability of getting *X* to that probability of getting *Y* (everything else being equal).

This simple argument for the master principle does not, however, carry over to cases of risk-conflict. If I favor *A1* to *B1* in case I it's not because I think that *A1* is more likely to be advantageous; in fact *B1* is more likely to cure me completely. In cases of risk-conflict I am juggling different probabilities of different levels of benefit. The nature of this juggling operation may, therefore, be drastically altered by introducing new probabilities as in the master principle. In the move from case I to case II we go from one kind of risk-conflict to a quite different kind. When I say that these are different "kinds" of risk-conflicts I mean that it may seem quite natural to describe them in quite different terms. In case I we can play it safe by taking a sure thing, whereas in case II we cannot play it safe in this sense so we may be tempted to go for broke. Given the difference in how these cases can naturally be described, why should one suppose that choosing the *A*-option in the first case commits one to choosing the *A*-option in the second case? Of course one could derive this commitment from the master principle, but the point I am making is that there seems to be no reason to suppose that the master principle is binding in cases of risk-conflict.⁷

This point spells doom for the program of the clinical decision analysts. Their agenda for educating physicians to be more rational revolves on the basic commandment of decision theory to maximize expected utility. But that commandment has no support—it is merely a dogma—given that the master principle does not apply to cases of risk-conflict. It would be a drastic mistake, therefore, for physicians to impose this commandment on their patients.

Let me try to make it clear that my "common sense particularist" standpoint doesn't imply that the first thing that comes into your mind about what you prefer has to be rational. There is certainly such a thing as wanting something when you really have good reason not to want it. The most obvious example is when you are making a factual mistake. In a famous recent philosophical paper Bernard Williams considers a man who wants to drink the liquid in the glass he is holding because he thinks the glass contains a gin and tonic, when really it contains poison.⁸ Certainly this man has good reason not to want to drink that liquid, though he doesn't realize this.

A slightly more subtle case concerns mistakes about dominance. When I was comparing options *A3* and *B3* in case III I mentioned the possibility of

⁷Both McLennan, op. cit. and Machina, op. cit., present interesting rebuttals of Raiffa's "dynamic" argument for the master principle, which is another dimension of the issue. In Robert Nozick, *The Nature of Rationality* (Princeton University Press, 1993), it seems to be implied that the master principle may fail to apply because of "symbolic value." My point, however, is that even if the patient has the most narrow "non-symbolic" interest in matters of health, there is no reason to think that the master principle applies to cases of risk-conflict.

⁸Bernard Williams, "Internal and External Reasons," in *Moral Luck* (Cambridge University Press, Cambridge, 1981), p. 102.

being confused by the fact that A3 affords the higher probability of achieving the intermediary outcome "survival with impairment." Certainly someone would be irrational who favored A3 over B3 because of this confusion.

Let me give another example of irrational preference found in the literature. The psychologists Amos Tversky and Daniel Kahneman describe a phenomenon that might be dubbed the "half-empty-half-full fallacy."⁹ They found in a number of studies that people are more likely to favor a treatment option (such as surgery) if it is described as giving an 80% likelihood of survival than if it is described as giving a 20% likelihood of death. Even surgeons are found to react in this way! Of course this is a gross fallacy, since an 80% likelihood of survival is identical with a 20% likelihood of death.

In each of these cases of irrational preference I think it would be possible to say that the preference is based on a mistake about the facts in a broad sense. Perhaps a better way to put it is that the irrationality of a certain sort of preference stems from there being a higher level preference not to have a preference of that sort.¹⁰ In the example of dominance it can be assumed that the agent has the general preference to do what is most likely to be advantageous, but doesn't realize that this requires preferring B3. In the half-empty-half-full case the agent has the general preference not to wind up both preferring and not preferring the same thing under two descriptions, and that general preference is frustrated when the half-empty-half-full fallacy is committed.

No argument of this sort seems to show that it is irrational to prefer the A-option in case I and the B-option in case II. Might it be suggested that every normal person has the general preference to abide by the master principle? I find that completely fanciful. Imagine a doctor interviewing a patient faced with the medical problem that figures in cases I and II. "What are your basic concerns here?" the doctor asks. "The main thing for me," replies the patient, "is to come out of this alive." Of course if it's possible I'd really like to avoid having a speech impairment for the rest of my life....And, oh yeah, one other thing: I don't want to violate the master principle." The absurdity of this reply is not merely its tendentiousness. It seems absurd to suppose that a normal person faced with a medical crisis could care about the master principle as such. What he cares about is doing what is—in terms of his own life-priorities—most likely to be advantageous. But no one is claiming—and this is a point that has to be strongly emphasized—that conformity to the master principle improves the likelihood of a more advantageous outcome. To think that would be to ignore the difference between cases of dominance and cases of risk-conflict. Why, then, should this patient have the slightest concern about the master principle? He *might* have such a concern—if he were indoctrinated by the clin-

⁹Tversky and Kahneman, op. cit., p. 453, and "Can Normative and Descriptive Analysis be Reconciled?" *Philosophy & Public Policy Report*, vol. 8, 1988, p. 4.

¹⁰Compare with Williams, "Internal Reasons and the Obscurity of Blame," in *Making Sense of Humanity* (Cambridge University Press, Cambridge, 1995), p. 37.

ical decision analysts, for instance—and I'm not necessarily claiming that it would be irrational to have such a concern—though I would, frankly, be inclined to regard such a person as suffering from a kind of overbearing intellectualism—but my point is simply that it seems neither abnormal nor *irrational* not to have such a concern.

My main conclusion, then, is that physicians must not attempt to resolve cases of risk-conflict by imposing the demands of decision theory on typical patients. Insofar as these demands are not rationally compelling it would clearly be unethical for physicians in any manner to pressure or even encourage patients to meet those demands.

Is it my position, then, that physicians must never tell patients what to do in cases of risk-conflict? No, that is not implied by what I am saying. Let me give an example to show why this is so. Suppose that a patient is bitten by a rabid bat. I don't know what the probability is that the patient will acquire rabies. I think this depends on the number, location, and depth of the bites. Let's imagine that the probability in the kind of case in question is reckoned to be about 30%. Rabies is always a fatal disease, but, as I understand it, there is virtually no chance of getting it if a series of several rabies shots are begun before the onset of symptoms. Nowadays rabies shots are generally covered by insurance, they are given in the arms or buttocks, and they are no more painful than, say, a flu shot. Obviously this patient should go for the rabies shots. But, let us imagine, he declines to do so. "I'm not going to bother taking a series of shots just to avoid a 30% chance of dying. There's a 70% chance that I would just be wasting my time, that nothing would happen anyway. Forget about it." In any typical case (the patient is otherwise reasonably healthy, has no especially adverse reactions to shots, and so on) I think this is clearly irrational, and I would expect a physician to do whatever she can to get this patient to change his mind.

I say that the patient is being irrational, but let me emphatically add that I have no way of *proving* this. It's important to see that we are dealing here with a case of risk-conflict. This is not a case of dominance. We cannot simply say to the patient, "You'll probably be better off if you take those shots." On the contrary, probably it won't matter whether he takes them or not. But the risk of not taking them seems obviously too high. The patient has compelling reason to get the shots, whether he recognizes it or not.¹¹ Here I am simply appealing to common sense intuition. There seems nothing else to appeal to. It should be understood that even the decision theorist would agree that in this kind of case there is nothing else to appeal to, for the patient in our example need not be violating either the master principle or the commandment to maximize expected value.

¹¹I am not able to tell whether this reason is "internal" or "external" in Williams's sense. Given the ubiquity of risk-conflict this doubt seems to cast a significant shadow over Williams's dichotomy. Cf. Scanlon's criticism of the dichotomy in the appendix to Thomas Scanlon, *What We Owe to Each Other* (Harvard University Press, Cambridge, Mass., 1998).

I don't want to minimize the puzzles involved in claiming that the patient is irrational in this kind of case. Consider for a moment that if the probability of getting rabies were one in a trillion instead of 30% the decision not to get the shots would seem unexceptionable; on the contrary, the decision to get shots in that case might seem irrational. Who has laid it down that 30% is too high a risk but one in a trillion can be ignored? Is this judgment somehow rooted in human instinct, or is it in some way culturally relative? I don't know the answer to these questions; as far as I know, no one does. But I am still prepared to appeal to what seems intuitively obvious in these cases, and I would certainly want the ideally rational physician to do the same.

In fact physicians are always telling patients what to do in cases of risk-conflict. If you take medication for hypertension there is some minute probability that it will cause you to fall asleep within the next week at the wheel of your car and lead to a fatal accident. There is no chance (or a much lower chance) that leaving your hypertension untreated will cause you to die within the next week. Your chances of dying in the next week are therefore increased to some degree by your taking the medication. Physicians tell you to take the medicine, anyway. And they are, I think, right to do so, though there could not possibly be any way to prove this.

In many cases, we see, it will seem intuitively obvious to a physician how a risk-conflict ought to be resolved. In those cases I hold that the ideal physician will trust those intuitions and not look for fancy arguments that don't exist. The guiding principle in this kind of case must be Aristotle's profound remark that "one must not seek more precision than a topic allows."

But what about the kind of cases in which there is no intuitively obvious resolution to the conflict? Those were the cases for which the clinical decision analysts hoped to introduce the apparatus of decision theory. Since I have repudiated that approach, what can be done in that kind of case?

Let's look again at the case we started with, the case of the diabetic patient with the infected foot. Should the patient have immediate surgery or wait? Waiting holds out the hope of escaping without any disability, but increases the risk of either death or a more extensive disability. How should we advise the patient? We are not going to try to specify utilities and then calculate expected values, because we don't believe in that. Should we perhaps tell the patient simply to flip a coin? But the action of flipping a coin will merely give us new probabilities of the medical outcomes. Having immediate surgery gives us one kind of gamble; waiting gives us another kind of gamble; and flipping a coin gives us a third kind of gamble. Which of these gambles should the patient take? The idea of flipping a coin has merely complicated the risk-conflict without resolving it. Should we now flip a coin to decide whether to flip a coin? ... Maybe we should forget about flipping coins.

When I project myself imaginatively into this patient's situation I find that I want to try to figure out what to do. It seems to be central to the phenomenology of this kind of risk-conflict—I mean the kind of risk-conflict in which

no intuitively obvious answer immediately suggests itself—to feel that there is something that needs to be thought about. I think it is liberating to realize that in a very important sense there is *nothing* to think about here. Of course one has to review the various probabilities; and one has to consider what it means to lose one's leg below the knee, and what it means to lose one's leg above the knee (and perhaps what it means to lose one's life). Those are things that certainly do have to be thought about. But my point is that once one has gone over those facts there is nothing left to do but *decide*. It's not as if there is some *right answer* that one can be looking for. I think it is in a way liberating to free oneself from the illusion that there is a right answer. It's not as if I carry around within me some specific degree of concern for my foot below the knee, and for my foot above the knee, and for my life—and then those concerns somehow should get calculated into the right answer about what I should do in this case. That's the picture generated by the clinical decision analysts, and it is, I think, an extremely misleading picture. The patient has to decide, and that means doing something more like creating the right answer than discovering it. To resolve the conflict is to make a decision that feels somehow acceptable, that puts the conflict to rest. There can't be anything more to it than that.

The patient does need to think over the various facts of the case. That's an essential part of reaching a decision that feels acceptable. "The mind goes about its business"—a line from a play by Edward Albee—and part of the mind's chief business is to go in circles, that is, to go round and round the facts of the case until something feels acceptable. The ideal physician will help the patient to review the facts of the case, and will listen compassionately as the patient's mind goes in circles.

Suppose the patient asks the physician to decide for him. I think the physician should be open to that. Why not? It's not as if there is some worry about getting the wrong answer. Remember, there is no right or wrong answer in this sort of case. Of course the physician who takes this burden on herself will have to face the conflict of risk and try to put the conflict to rest in her own mind. But she may be in some ways more adept at doing this than her patient. Perhaps she has participated in (or read about, or observed) this kind of conflict-resolution many times before. That may make it easier to draw the circle to a close and decide.

But my picture of how it should ideally work is that the patient and physician talk the thing out and come to a resolution together. Resolving this kind of conflict of risk is a leap in the dark. The ideal physician will take the patient by the hand, and they will leap together.