Great expectations

Adam Morton

As you leave a bar in the early hours you are approached by a shifty character who offers you the following deal. If you first pay him a hundred dollars you can choose a number and he will roll a die, this die here. If it lands with your chosen number up he will pay you a thousand dollars. If it lands with any other number up, well, it's goodbye to him and to your hundred. Of course you worry a bit about whether the deal really is as described, even though he throws the diea few times and it doesn't seem loaded, and you have a quick glimpse of ten hundred dollar bills in his hand. On the other hand, it's a good deal. You are being offered a one in six chance of a thousand dollars, which you take to be worth one hundred and sixty seven, for only one hundred. And yet, a vague worry flashes through your mind, elusive enough that you cannot quite state it, but strong enough to make you decline the deal.

The next morning, sober and well-slept, though slightly hung over, you try to pin down the worry. Finally you see it. The disturbing thing about the gamble is that even if it is played honestly *most people who* *play it will lose*. Five sixths of them, in fact. And this is in spite of the fact that the odds suggested are strongly, almost worryingly, in the player's favour.

This chapter is concerned with the reasons we can give for our choices between risky options. It focuses especially on the phenomenon described in the previous paragraphs. I shall argue for a close connection between this phenomenon and some pervasive facts about human attitudes to risk. My aim is to say something helpful about the ways in which we ought to approach risky choices, the strategies and attitudes that can benefit us. There is an enormous literature in philosophy, economics, and psychology on decision making in risky circumstances. I am not aiming to improve or correct any of this tradition, but rather to bring out some points that it misses or hides. One usually neglected theme that becomes increasingly important as this chapter progresses is that of decision-making *virtues*, traits and abilities that a person has to have in order to make the most of whatever attitude to risk she has adopted. As I develop this theme it is closely related to another, that of the tensions between the interests of an actual person and the possible people who would have gained or suffered had the person's choice turned out differently. To see the relevance of this second theme to risk-taking suppose that you had taken the gamble offered by the shifty character and, as was most likely, lost. The next morning you are berating your self of the night before. ‘Why did you lose me the hundred dollars I need to take my friend to an impressive brunch?’ you ask. And in your imagination your past self replies ‘because although you in fact lost, in one sixth of all the possible futures you won big – wasn’t that worth it?’ To which your present impoverished self groans ‘I wish you were less considerate of our possible selves, when the probable cost to my actual self is so high.’

variability and expectation: three facts An agent, A, is facing a choice between two options. The options could be taking the right or left fork in the road. How things will turn out after her choice depends on some facts that she does not know. (Perhaps whether there are still bandits in this territory, who will attack travellers who take the short-cut path.) Suppose that one of the options is riskier than the other, in that it might turn out much better than the other, and also might turn out much worse, depending on the unknown facts. But suppose that she does have ideas about how likely these facts are to be true. (She knows that there have been no attacks for years, though not all of the known bandits have been caught.) Then she can ask ‘do the potential benefits of taking the first option outweigh its potential costs?' She has to ask herself this, because this is the essence of the situation she faces. But the concept of outweighing demands a lot. It asks her to count for an action its benefits and the probability that these will follow, to count against it its costs and the probability that these will follow, and to combine these probabilities and values in a way that allows them to be compared. The result of the combination is the *expectation* of the action: its benefits weighted by the likelihood they will occur reduced by its costs weighted by the likelihood they will occur.

How can we calculate expectations? The standard model of a situation in which the calculations are unproblematic is given by games of chance. Suppose that instead of a fork in the road our agent faces a choice between two gambles, o_1 and o_2. In o_1 a fair coin will be tossed: if it lands Heads she wins $200, and if it lands Tails she looses $100. In o_2 she gets $30 whatever. (She might have paid to be in the situation where this choice is open to her. The person offering the gambles does not have to be benevolent.) If the gamble were repeated many times and she took o_1 every time her gains would approach $ 50 times the number of repetitions, since she would win roughly half the time and loose roughly half the time. And if she took o_2 every time her gains would approach $30 times the number of repetitions. (Or, equivalently, if a zillion duplicates of her were to take o_1 they would end up with 50 zillion dollars, and if they took o_2 they would end up with 30 zillion.) So this aspect of the gamble, the average amount that it would yield if repeated indefinitely, is clear. For each option it is the probability of Heads times the benefit from Heads for that gamble plus the probability of Tails times the benefit from Tails. Then it is just a small step to taking this expected or average value to be 'what the gamble is worth', to specify how the agent should rank it in comparison with other options. (Standard expositions of this idea are in chapter 1 of Jeffrey 1983, chapter 2 of Raiffa 1968, and chapter 3 of Resnik 1987. For some history see Hacking 1975, especially chapter 11.)

Suppose that the agent takes this small step, and evaluates the options by their expected value. Then the 'average' amount she will gain from the risky option will be \$50, which is better than the \$30 she will get if she takes the less risky one. 'Average' here means average over some indefinitely large set of possible occasions. It could be all the ways things might turn out if she takes the option, or it could be the way things would turn out if she (impossibly) were to repeat the choice over and over again. Thought of either way, the benefits to the assembly of her possible selves will be greater if she takes the risk. But there is also a price for this: some of these selves will end up richer than others. After the risky choice her possible selves will fall into two classes. Half of them will be $200 better off and half $100 poorer, and after the riskless choice all her possible selves will have the same profit of $30. So the downside for the greater expected outcome is the greater variability of actual outcomes.

This fact is completely general. Take one gamble to be riskier than another when its possible outcomes are more varied. (There are several ways of making this precise, and their differences do not matter here. For simplicity take a riskier gamble to have a greater variance of distribution of outcomes.) It will follow that people who make riskier choices will experience more varied outcomes than those who make safer ones. Consider the effects of this on a population of people in a gamble very similar to the one just described. A coin is tossed: if it lands Heads players gain $1, if Tails they lose $1. The game can continue, and then after the second toss each player may have $2 (after two successive Heads), $0 (after Heads-Tails, or Tails-Heads) or $-2 (after two successive Tails). And so on. Call this game g-risky and compare it to an alternative g-safe in which the players win or lose nothing. The two games have the same expected value, 0. Consider two sub-populations of players, one playing each game. After one round of g-risky roughly half of the first sub-population will be richer by $1 and roughly half will be poorer by $1. No one's wealth will have been unchanged. After two rounds about one quarter will be richer by $2, one half will have returned to zero, and one half will be poorer by $2. And of course the whole of the second sub-population, playing (or perhaps not-playing) g-safe will have remained at zero. And if we continue to play the game with more and more rounds then though the average wealth of the two populations is the same, it is distributed very differently. In the first population there are eventually some extremely wealthy people and some grotesquely indebted ones, while in the second population no one has changed their wealth more than anyone else.

Some very general facts are beginning to emerge. Fact number one: *when two gambles have the same expected value the riskier one will also produce a wider distribution of results*: more or greater winners and also more or greater losers.

There are social consequences of this fact, though they are not the focus of this chapter. A society in which people are free to take risks for their own benefit will often end up with a higher average wealth, but it will also end up with a greater variation in wealth, so that it is quite easy for many people to be worse off as a result of the choices that lead to an increase in the average well-being. (And the people who emerge well off will compliment themselves on their wise choices and their sense of opportunity, when often the fact will be that they are the few for whom the coin came down Heads many times.) More to the present point is the consequence that a riskier gamble can make it more likely that one does badly. Or, to put it more carefully, given two gambles with the same expected value, the riskier one will sometimes present a larger probability of emerging with less than the expected value of the gamble. And, more generally, sometimes though one gamble has a higher expected value than another it also makes it more likely that you will do worse than you will if you had taken the other. (Two ways: it can be more likely that you will do worse than the expected value of the other, and it can also be more likely that you will do worse than the most likely outcome of the other gamble.) This can be illustrated by variants of the g-risky game just considered or by cases like the shifty character story at the beginning of this chapter. Fact number two: *by choosing a gamble which has a higher expected value but also a greater risk, one can often increase the probability of doing badly.*

My examples have been gambles with money, and I have been measuring expected value in money. Most choices are not like this, and the person approaching the fork in the road in bandit country, for example, will not frame her problem in money terms. But the same facts apply. The real dilemmas of risk will occur when two options are roughly equal in their overall attractiveness, balancing possible benefits against possible losses, but where one, the riskier one, offers greater possible benefits at the price of greater possible losses. (Or, to put it more subtly, the element of risk only arises as a factor in comparing options when one can make some sense of the balancing of benefits and losses, but the possible losses of one option are greater than those of another otherwise equally attractive one.) And then analogs of the two facts just mentioned above will be true. People who take more risks, in order to get more benefits, will get more variable results than those who take fewer, so that over time risk-takers will be both better and worse off. And if they make clever choices among a wide range of options people who take more risks will be more likely to do badly, though on average they will do better.

what you really prefer In the next section I will consider the pros and cons of some ways of steering between the competing pulls of safety and benefit. But first I shall formulate things in slightly greater generality, in a way that allows me to bring out a third important fact.

It also allows me to engage with a standard economists' line about risk.

An option is risky, I said, in comparison with another, when it might turn out better but also might turn out worse. I have been illustrating risk in a traditional way with gambles between sums of money. But it has also been clear that one cannot react to such gambles without considering *how much* one wants the possible gains and losses. And it is pretty clear that we cannot measure how much one wants things, not even amounts of money, in direct proportion to their monetary values. The standard example is that a gift of $1,000 is wonderful for a street person, but almost irrelevant to a billionaire. So for the street person a gain of $1,000 will be something like ten times as desirable as a gain of $100, while to the billionaire it is hardly better at all. So consider someone for whom, say, $10,000 is about ten times as good as $1,000, but $20,000 less than twice times as good as $10,000. (This might be because for each extra thousand up to about ten thousand the person can improve their life in basic ways, but after that the money would go for luxuries.) For that person a heads/tails gamble between $20,000 and $0 will not be the same value as a certainty of $10,000, since the $10,000 that would be lost if the coin comes down the wrong way would be missed more than the $10,000 that would be gained if it came down the right way. It is standard in economics to express this by describing agents in terms of 'utility functions', measures of the benefit to them of different amounts of goods, in particular money. (I shall use the terms 'benefit' and 'utility' interchangeably: 'benefit' when it makes more natural English, 'utility' when it makes a clearer connection with the literature.) It is usually assumed that money has a diminishing marginal utility, that is, that all people's utility functions give higher values for more money than for less, but increase less and less as the amounts get larger. (For expositions of the utility of money see Chapter 1 of Hargreaves-Heap et. al. 1992, chapter 4 of Raiffa 1968, chapter 6 of Morton 1990, and, going back to a source of much later work, Friedman and Savage 1948.)

A person whose utility function is of the diminishing marginal utility kind, and who evaluates gambles in terms of their expected utility or benefit rather than their expectation in money or other goods, will automatically show a certain kind of risk-aversion. That is given two money gambles with the same *monetary* expectation, one of which has a wider spread of possible outcomes than the other, the person will prefer the one with the less wide spread. Such a person will for example usually take a heads/tails gamble between two amounts of money to be less valuable than a certainty of an amount half way between the two, since the value to the person of increasing by that half way amount is less significant to them than the disvalue of decreasing by it.

There is clearly something right about this. The value of money, and most goods, does increase at a decreasing rate. In the hands of some economists, though, the point has got expanded into the following thought 'whatever choices someone makes, unless they are blatantly irrational, we can find a utility function so that the person is choosing the options with the greatest expected value in terms of that utility function. So in a way, no one is risk averse in terms of their own preferences.' This view is sometimes called the 'revealed preference' theory. (A person's 'real' preferences are revealed by their choices, not by what they say or think they want.) I think it is a misguided position. Though this is not the place to argue the issue, I do not think the resulting theory can give real explanations of why individual people make particular choices. And, quite evidently, it is no use to people trying to think their way through risky situations, since it says "whatever you eventually do will have been the rational choice, in terms of the preferences you revealed in choosing it." In focussing on which are the rational actions given a person's preferences, it loses sight of the more interesting question of what are the rational ways for people to think out what their preferences and hence their actions should be. (For a brief discussion of revealed preference see chapter 2 of Hargreaves-Heap et. al. 1992. A source of the view is Samuelson 1947.)

One complication that makes it harder to get a revealed preference theory off the ground, and which is directly relevant to the main themes of this chapter, is the fact that a person will sometimes be best off *not* evaluating gambles in terms of her utility functions, but in terms of money. This is true even if 'best off' is understood in terms of her personal utility rather than money or other goods. This paradoxical-sounding situation can be best understood through an example. Suppose that a person is facing a series of choices between gambles in a fixed period of time, all with monetary outcomes. Suppose to make it simple that the person has paid a fixed amount for the opportunity of making the choices, and that none of them involve losses. Suppose that the person wants to come out of the process doing as well as possible. Then she has two strategies, consistent with the general aim of trying to maximize expected utility. (a) take each gamble as it comes, evaluating them in terms of her utility function. (b) decide in advance how she will choose at each possible stage of the process, in such a way that her expectation of gain at the end is maximized. One might think that these amount to the same, but they do not have to. Our person is not extremely rich, so she regards sums up to $10,000 as serious amounts and values them linearly, so that gaining twice as much money within this range is twice as good. Above $10,000, though, the importance of increases in money begins to tail off, so $20,000 is just more than half again as valuable as $10,000. I'll express this by assuming that she has a utility function u such that u($0) = 0, u($10,000) = 1, u($20,000)=1.6, and that u follows a straight line between the values of $0 and $10,000 and then curves smoothly to the value of $20,000 and onward, increasing at a decreasing rate. (The exact numbers could be changed quite a lot while still supporting the point I shall make. By the utility of x dollars I mean the value of an increase of her pre-game wealth by $x, not the value of having a wealth of $x.) She is faced with a series of twenty choices between gambles, in rapid succession. Each will consist of a choice between the risky option of 'Heads $20,000; Tails nothing' and the safe option of $9,000 for sure. What should she do? Suppose that in accordance with the suggestion (a) just above she first considers the first choice. The expected value in utility terms of the risky gamble is ½u($20,000) + ½u($0) = 0.8, and the expected value of the safe gamble is u($9,000) = 0.9. So she will choose the safe gamble, and then for the same reasons choose the safe gamble on the following choices. And at the end of the series she will have $9,000x20 = $180,000. Now contrast this with what happens if in accordance with suggestion (b) she thinks of the whole series as a single unit. A little calculation shows that the expected monetary value of the gamble is $200,000, which is definitely more than $180,000, measured either in utility or in money. (In this case, then, it makes no difference whether she goes for money or goes utility.) So if she thinks in this second way and values gambles in terms of their expected values she will commit herself to taking the risky gamble each time. So (a) and (b) are not the same; evaluating a series of gambles as they come along leads to different choices as evaluating them as a single process. Moreover, it is clear which is the better way to think. If she makes the (a)-style series of choices she ends up with $180,000. Suppose on the other hand she makes the (b)-style single choice. Then she will on average end up with $200,000. Moreover – as slightly longer calculations will show – there is an approximately 0.6 chance that she will gain $200,000 or more and an approximately 0.4 chance that she will gain $180,000 or less. (At 20 trials the next worst possible outcome after $200,000 is exactly $180,000.) So she is more likely than not to do better by following (a) than (b). So – however she balances utility against money and whichever way she evaluates the advantages of a decision-making method – she is better off thinking of the series of gambles as a single choice[1].

The point here can be summed up as another fact. Fact number three: *avoiding risk by valuing money or other goods in a risk-averse way – one that makes the value of greater amounts less than proportionately greater than that of smaller amounts - can, if one does not carefully pick the way one frames one's choices, lead to choices that are less valuable, even in terms of the chosen valuation.* Canny expectation-oriented agents will choose their utility functions carefully, with an eye to the choices that are imminent[2].

choosing how to choose The main character in the story so far has been the rule 'evaluate gambles – or risky options in general – in terms of their expected values'. (Henceforth often 'the expectational rule', or 'expectational thinking'.) This rule can be stated in considerable generality as a basic principle of decision-making. In fact formal decision theory is essentially an elaboration of the expected value principle into a description or prescription applying to all aspects of decision. The rule is not easy to use, though. It requires us to put our preferences into numbers: and besides the difficulties of knowing and quantifying how much one wants something fact number three shows a deep complication in doing this. It also requires us to express our estimates of likelihood or degrees of belief as numerical probabilities, also very demanding in terms of self-knowledge and the assessment of evidence. And it requires us to make probability and utility calculations of a kind that can easily become very complicated and easy to bungle. I think – I just don't know whether the suggestion would meet with general agreement – that to recommend thinking of risk in terms of expectation is to suggest an intellectual virtue, that of being able to produce numerical estimates, of how likely events are and how much one would gain from outcomes, which give one an accurate grasp of the relative frequencies and acceptabilities of potential outcomes. (Or at any rate a virtue that will give a grasp of these things that leads to decisions that one tends not to regret.) There's not much point deciding by expected values unless you have acquired this virtue, and acquiring it may require a major reorganization of your thinking. One of the signs of being a member of a society with a market economy, perhaps.

We may well ask: why follow this rule? We might ask this with respect to the choices we will have to make in a particular situation – 'how am I to think this one through?' – or with respect to the choices people should make generally. Here are three reasons that might be given for valuing gambles by expectations, together with problems with the reasons.

<u>reason 1</u>: if an agent values gambles by their expectations then shewill probably do better than if she chooses in accordance with some other rule.

<u>explanation</u>: if a person makes a long series of choices between gambles, and always chooses the gambles that have the greatest expected value, then her total gains will approximate, the more so the longer the series, to the sum of expected values of her choices, and this will be greater than any other policy is likely to result in. Therefore evaluating risks by expectations is the best possible policy.

<u>Problem</u> (1): *it may not be true*. Agents operating on perfect information, with perfect calculating ability, and perfect self-knowledge may make the best decisions this way, but it just is not evident that an imperfect agent facing a long series of choices will come out best by thinking expectationally. This is so for two reasons. First, the effect of stakes in evaluation and calculation. Second, fact number three above: if by expectational thinking we mean valuing each gamble one meets in terms of its expected value according to one's preferences then it is not true that this will give one the best results one can get. (And if we mean to include just the right combinations of acts and choices of utility functions, then we are adding a completely new level

of considerations, an additional set of rules, which have never yet been carefully formulated.)

Problem (2): *death*. No real agent goes through more than a limited number of choices before dying. And some of the gambles we consider have death (or the commercial equivalent, bankruptcy) as a possible consequence. So the problem of getting from what would be best on an indefinitely long series of choices to what would be best in a finite life span is as difficult as that of getting from a range of possible choices and outcomes to an actual one.

Problem (3): *equivocation*. As fact number two shows, in some situations most people who choose in terms of expected value will do badly (though a minority will do well enough to raise the average.) So it is simply not true that to choose by expected value is always to make it more probable that you will do well. Will people who choose by expected value in general do well? That depends on what gambles the human race happens to be faced with, and in what proportions. Will possible people meeting all possible gambles do best if they think in expectational terms? In the absence of some probability distribution over all possible occasions of person-meets-gamble the question is not well defined.

reason 2: if you evaluate risks in accordance with their expected values the probability that you will regret your decision is minimized.

explanation: the benefits of expectational decisions tend to be greater than those of decisions made in other ways, and so the difference between the outcome of such a decision and the outcome had an alternative been chosen tend to be in the favour of the expectational decision.

problem: the argument just given does not support expectational thinking. Instead it supports the quite different idea that the best decision is the one that will probably give the best results. That is a different idea because, for example in the midnight gamble at the beginning of this paper the option with the greater expectation will probably give the worse results. Some people might take it as a better idea, but there are cases where it is clearly wrong. Consider a choice between on the one hand a gamble which leads to $1,000 nine hundred and ninety nine times out of a thousand and death one time in a thousand, and on the other hand a “gamble” that gives no money and no danger. Then the first gamble will probably give better results, but the risk of being in the minority of dead people is clearly not worth it.

reason 3: people who evaluate risks in accordance with their expected values will generally do better than those who do not, so you should encourage this mode of evaluation.

explanation: one way of thinking of what you should do is as an instance of a general pattern which it makes sense for people to adhere to. And you can hardly recommend it for others without following it yourself.

problems: (1) as with the first problem with reason 1 above: it is not obviously true that people who go by expected value do generally do better.

(2) assuming that the assumption is true, it must mean that people on average do better if they evaluate by expectations. Grant the further assumption that you should encourage the general adoption of practices that on average are good for people. It does not follow that these are the practices that will do best for you on this occasion. (Compare: you should encourage people to be truthful and cooperative, and your own truthfulness and cooperation may be a necessary means to this. But on any given occasion a lie or a cheat may be your best choice.)

Since I have been finding holes in justifications for the expectational rule. I should mention an objection to it, which also has problems.

the reliability objection: Grant that in many circumstances if we evaluate gambles by their expectations we will get results that on average are in our best interests, *if* our estimates of probability are accurate. It does not follow that, given our human propensity to mis-estimate and mis-calculate, gambles evaluated in terms of the consequently inaccurate probabilities will be similarly well-chosen. Expectational thinking may be fine for ideal agents, but we need something that works for us human bunglers.

problems: (1) Expectational thinking is surprisingly tolerant of approximate probabilities. Consider for example a case in which one is pondering a heads/tails gamble with a coin that one takes to be fair, with payoffs of \$1 for heads and \$0 for tails, so that one evaluates the gamble as worth ½\$1 = \$0.5. Suppose that in fact the coin is not fair, but biased so that the probability of heads is 0.7. Had one known this one would have given the gamble a value of \$0.7. The difference in the expected value of the gamble is linear in the error: if one's probabilities are out by 30% one's valuations are out by just 30% The expectational rule is even more tolerant when we consider embedded or repeated gambles, in which the expectation depends on the probability of an outcome which is itself probabilistic. In a simple embedding the expectation will be a function of the square of the probability, so if that is out by, say 30%, then the expectation will be out by only 11%. (Remember that probabilities are less than one, so that when raised to a power they become smaller rather than larger.) (2) Moreover estimates of risk are also quite tolerant of errors in probability. One's estimate of the probability of the worst case outcome of a gamble will be inaccurate only to the degree that one's estimate of the basic probabilities is inaccurate. It will not amplify the error. (And, more generally, the fact that one's estimates of probability may be in error increases the variance of gambles. But, considering them now as weighted sums of gambles where the weights are proportional to a distribution of errors assumed to have one's original estimate as its mean, the variance of this new distribution is not so different from that of the original one. I find this fact surprising.)

And, to end this section, here is a direct objection to using the expectational rule, that does have some force, though it could be exaggerated.

<u>the unpredictability objection:</u> We need to be able to predict the future, and riskier choices, even when accompanied by higher average payoffs, are less predictable in their outcomes.

<u>explanation</u>: Consider someone facing a series of choices like the

ones that illustrated fact number three. (Heads/tails between $20,000 and $0, versus certainty of $9,000; repeated several times.) If he takes the risky option each time his wealth will probably increase, but will do so in an unpredictable way. If he takes the safe options he can know in advance that he will be $9,000 richer at the end of each play. This might for example allow him to make plans to invest the money, or to keep it safe from thieves, or to purchase things expecting that the money will be available when the bills arrive.

<u>qualification</u>: The main disadvantage of expectational thinking that this reveals is the bad combination of risk-taking and unpreparedness. Someone who chooses with an eye to average rather than probable benefits should also have contingency plans, both for winning and for losing, and ways of using winning occasions to palliate the effects of losing occasions. The disposition to do these things is a practical virtue that has to accompany the expectational mentality, rather like the virtue of ordering one's beliefs and desires to fit that mentality, mentioned earlier.

against rules What we have seen is a tension between three intuitively appealing ideas about good decision-making. We can't hold all three of these as general rules.

(1) choose so that you are likely to get what you want. (Maximize the chance of good consequences.)

(2) choose so as to minimize the probability of bad consequences.

(3) choose so that you will on average get as much as you can of what you want.

(3) is the expectational rule. (2) is a general principle of risk-aversion, in the same general spirit as what is sometimes called the precautionary principle (see the chapter in this book by Per Sandin). (1) is a general principle of practical reasoning, which we often appeal to as if it were obviously true. But it is clear that no sensible person would always obey (1) or (2). (1) ought to be violated – for people with normal preferences about the kinds of choices they want to make - when there is a small but unlikely chance of something extremely desirable. (You are trapped in a canyon as the river rises. Your only hope is to climb a fragile vine. Probably this will fail, but it might just succeed. The alternatives are staying where you are and drowning in five minutes or moving to a ledge and drowning in ten minutes.) (2) ought to be violated when the same course of action leads to an extremely good outcome and an extremely bad one. (If the vine breaks the crocodiles will get you – but then if you don't take the chance you'll drown.)

When should (3) be violated? We have seen a number of situations where sensible people would be very reluctant to choose the act with the greatest expected value. They are all risky situations, of course, but I think it helps to distinguish two kinds of risk.

The first – call it *wide consequences risk* – comes when an act has both good and bad consequences. Since the interesting cases, the ones that make for hard decisions, are those where the risky option is a serious rival to a safer one, the hard problems involving wide-consequences risk arise when one option gives a higher expected value than another, but also gives a possibility of a worse outcome. Is the chance of the better payoff worth the risk of the worse one? The other kind of risk – call it *probable trouble risk* – comes when an act will most likely lead to bad results. The hard problems with probable trouble risk arise when one option will on average give a better result than another, but more often than not will do worse. We can have wide consequence risk without probable trouble risk when it is most likely that one will come out with the average or even the best result, but there is a small yet non-negligible chance of a much worse one. (See the diagram.) We can also have hard choices between risks of the two kinds, in which one option will probably turn out well but might well turn out disastrously, and the other option will probably turn out badly, though not disastrously, but might turn out well. These two options might be identical in expected value, but one person might be more averse to one and another person to the other. This is something more complex than risk-aversion; it is an aversion to a particular kind of risk, to a particular profile of desirable and undesirable consequences.

I do not think there is anything rationally wrong with someone who chooses in a way that avoids particular kinds of risk. There is nothing wrong, that is, as long as the person also exhibits other traits, virtues as I have been calling them, that are necessary in order to make a success of that style of choice. And as long as the person finds the general consequences of choosing in that way acceptable. This applies to choosing in terms of expected value too. In the course of this chapter I have mentioned a number of virtues that expected-value-choosers should have: the ability to think of one's preferences and one's degrees of confidence in numerical terms, the ability to make contingency plans for the inevitable times when a gamble with a high expected value has a low actual one, and the ability to schedule and gather together one's choices for the best overall outcome. If you don't have these virtues, then you should stay away from expectational thinking. You should also stay away from expectational thinking if you are not prepared for frequent losses which will accompany your gains.

In fact, I put this last point too gently. Suppose that the gambles available to you include many probable trouble risks. Then if you choose in purely expectational terms you will have more losses than gains, although if you are lucky you will have the occasional big gain. If you are not lucky you will have just the losses, while just a few of your like minded friends get the big gains. I doubt that very many people are really prepared for this. Most people's general preferences about the outcome of their choices, I am sure, countenance a mixture of losses and gains, but resist the possibility of unmitigated losses. That suggests that to the extent that people reason expectationally, they do so in contexts in which they are reasonably sure that the risks tend to the wide consequence rather than the probable trouble profiles. And this suggests yet another virtue, or rather yet another combination of a decision-making style, a preference for kinds of outcome, and a virtue. If you want to do well on average, are prepared to take losses along with your gains, but are not prepared for the prospect of far more and more likely losses than gains, then you must learn to find wide consequence risks, and avoid situations in which they are not to be found.

There are virtues associated with other attitudes to risk, too, though the focus of this chapter is on expectational thinking. If you are averse to losses even as the price to pay for the possibility of gains then you will need the virtue of equanimity as gambles you turned down pay off grandly. (And it would help to have the virtue of tact when gambles you turned down prove disastrous for others.) And above all you will need the virtue of finding non-risky options with reasonable payoffs: you will have to spend a lot of time looking. If you want to gain most of the time, accepting the possibility of frequent losses as a price for this, in effect submitting your choices to principle (1) at the beginning of this section, then too you will have to acquire the virtue of finding wide consequence risks where at first you see only probable trouble ones.

The central point is this: don't think in terms of rules alone. They are not that important. What matters is the fit between a person's general aims in decision-making, the procedures they have for comparing and deciding between options , and the skills of judgement, perception, and long-term self-control that they have or can acquire. These last, what I have been calling virtues, play a much larger role than is generally appreciated. One reason they are easily ignored is that we often frame questions about decision-making as if people were faced with a set of options, over whose composition they have no control, from which they must choose. But in fact we *search* for options, digging out facts and using our imaginations, and one of the central virtues of a good decision-maker is the capacity to find the right set of options to choose between.   A situation studded with probable trouble risk, for example, is not one in which comparison by expected value is very appealing. So what should you do if you seem to be in such a situation? The first thing you should do is look for more options! The main focus of advice concerning risk should be not how to compare the options that you have but what sort of options to search out before comparing them[3].

DIAGRAM   [see reference in text – this diagram needs to be professionally drawn]

REFERENCES

Broome, John (1991) *Weighing Goods*. Oxford: Blackwell

Friedman, M and Savage, L (1948) ‘The Utility Analysis of Choices Involving Risk’, *Journal of Political Economy* 56, 279-304.

Hacking, I. (1975) *The Emergence of Probability*, Cambridge: Cambridge University Press.

Hargreaves Heap, S., Hollis, M., Lyons, B., Sugden, R. and Weale, A. (1992) *The theory of Choice*. Oxford: Blackwell.

Jeffrey, R. (1983) *The Logic of Decision*, second edition. Chicago: Chicago University Press.

Morton, A. (1990) *Disasters and Dilemmas*. Oxford: Blackwell.

Pollock, J. L. (2002) Rational choice and action omnipotence *Philosophical Review* **111**, pp. 1-24

Raiffa, H. (1968) *Decision Analysis*. Reading, MA: Addison-Wesley.

Resnik, M. D. (1987) *Choices*. Madison WI: Wisconsin University Press

Samuelson, P. A. (1947) *The Foundations of Economic Analysis*. Cambridge, MA: Harvard University Press

Savage, L. J. (1954) *The Foundations of Statistics*. New York: Wiley

NOTES

1. The point could be made with a shorter series of choices by using gambles in which the expected monetary value is just greater than that of the certain alternative, but in which the agent is more likely than not to get more than the expected value. Though discrepancies between expectation and most likely outcome are important in this paper, I thought it would be simpler to make the point with more familiar-shaped gambles.

2. Really canny agents can have very complex utility functions. One aspect I have not discussed is that they may make their evaluation of an outcome depend on whether the process that led to it was a risky one. See Broome (1991), especially chapter 5. I do not think this affects the point I am making here. My point is related to what Savage (1954) called the 'small world problem'. See also Pollock (2002).

3. Risk-management in terms of procedures for searching for options

is discussed in chapter 7 of Morton (1991).

4. James Hawthorne and John Simpson gave me good advice about this paper.