

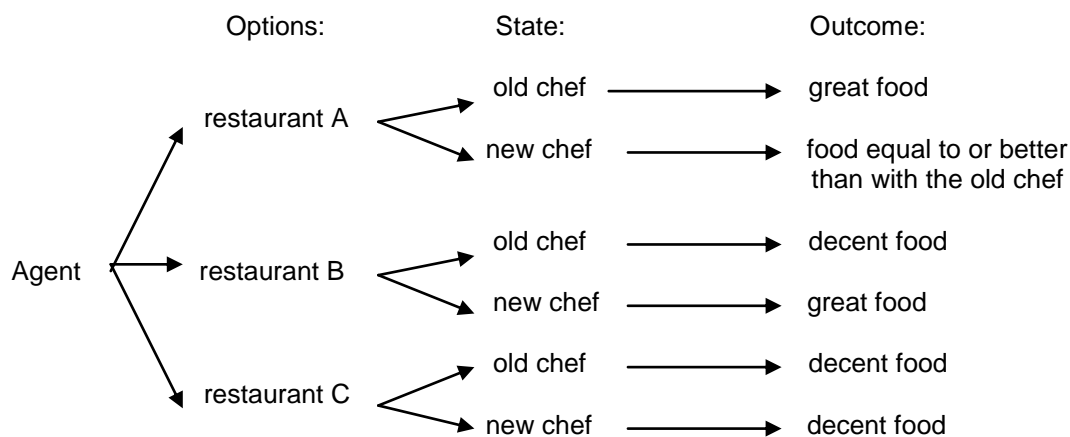
MAKING GOOD CHOICES: AN INTRODUCTION TO PRACTICAL REASONING

CHAPTER 6: SINGLE CRITERION INDIVIDUAL DECISION UNDER RISK: EXPECTED UTILITY

Chapter 4 presents practical reasoning methods for analyzing (framing) and evaluating (solving) multi-criteria decisions under certainty. In this chapter we make two changes in this topic. First, we give up multi-criteria decisions and change this back to single criterion decisions. This change simplifies one aspect of the decision situation. Next, we will give up decisions under certainty and change this to decisions under risk, applying the framework set up in Chapter 5 for representing risk as degree of confidence based on probabilities of states. This change adds an aspect of difficulty to the decision situation. We will begin with the simplest kind of risky decision problem and then advance to problems requiring the agent to evaluate options by the method of expected utility.

6.1 Single criterion individual decision under risk, single stage: decision by dominance.

Imagine an agent whose goal is to eat out at a nice restaurant, getting the best possible meal for herself. She has narrowed her choices to 3 restaurants. But the restaurant business in her town is very competitive, and restaurants frequently hire good chefs away from each other. The agent is not certain for each of the 3 restaurants whether the old chef or a new chef is creating the meals. She knows the restaurants well and the quality of chef each has been trying to hire, and so knows the quality of meal that she can expect depending on the chef situation in each restaurant. Which restaurant should she choose? Here is the decision diagram for this problem.



This is a decision under risk, for the agent is uncertain about the chefs at each restaurant and this affects the outcome of her decision. For example, if she decides on restaurant B, she would like the new chef to be working not the old chef, for this will gain her more of her goal than would B with the old chef. There will certainly be either the new or the old chef in restaurant B, but the agent is not certain which one it will be if she decides to have her meal there.

Before trying to assign a probability to the states of each option, let's see if this problem can be solved in a quicker way. Note that for one option, restaurant C, the best possible outcome is less of the agent's goal than the best possible outcome of each of the remaining options. The best possible outcome of restaurant C is only the agent's security level. There is no hope of doing better with this option, but there is hope of doing better with either of the other two options. Also, note that the worse outcome for the other two options is equal to or better than the best outcome for restaurant C. We can see by comparing the outcomes of each option that restaurant A and B each "beat" restaurant C. Thus, restaurant C is dominated by the other options. The agent would be irrational to choose restaurant C, given her goal and the availability of the other two options. So, it is safe to drop restaurant C as an option.

Now let's compare the outcomes of the two remaining options. Note that the best the agent can hope for in restaurant B is guaranteed in restaurant A no matter who the chef is. In other words, the worse outcome for restaurant A is equal to the best outcome for restaurant B. In effect, the agent can't go wrong, given her goal, by eating in restaurant A whereas she has the possibility of gaining less of her goal if the state "old chef" exists in restaurant B. Thus, restaurant A dominates restaurant B. It would be irrational for the agent, given her goal and the availability of restaurant A, to decide to eat at restaurant B.

By eliminating two options, the solution to this decision problem is the remaining option – restaurant A is the rational choice. In decisions by dominance, the agent need not assign probability estimates to the states and form degrees of confidence concerning the outcomes. Given a single criterion, the

agent compares the possible outcomes of the options, eliminating the dominated options. The dominant option will maximize utility. The rational solution to the above problem is: choose (A) over (B) and (B) over (C) (that is, these 3 options should be ordered in this way: Given the goal, restaurant A should be preferred to B, and restaurant B should be preferred to C).

Here is the general definition for dominance and the rational choice rule for this method of solving risky decision problem.

Definition of dominance:


For any two options x and y , and any goal G , x dominates y if-and-only-if:

- (i) each outcome of x is equal to or better than each outcome of y , given G ,
- and (ii) at least one outcome of x is better than any outcome of y , given G .

Rational choice rule – For any two options x and y : if x dominates y , then choose x over y .

6.2 Single criterion individual decision under risk, single stage: decision by expected monetary value (EMV)

Imagine an agent who is offered the chance to play a game. One game costs \$5 to play. The agent rolls a fair die and if the side with 5 dots lands up, the agent wins \$10. If a side with 1, 2, or 3 dots lands up, the agent wins \$5. If a side with 4 or 6 dots lands up, the agent wins \$7. A second game is very much like the first one with these differences. If the side with 2 dots lands up, the agent wins \$10. If a side with 1, 3, or 5 dots lands up, the agent wins \$2. If a side with 4 or 6 dots lands up, the agent wins \$3. This second game costs \$2 to play. Of course, the agent can decide to play neither game. Suppose the agent's goal is to maximize winnings with the minimum cost, what is the rational choice for this agent? Here is the decision diagram for this problem.



Given the agent's goal, games #1 and #2 both dominate the option of not playing. While the option not to play costs nothing, it also has as its outcome no winnings (an option under certainty, there is no risk involved). In game #1 and in game #2, the worse case for each is that the agent breaks even (winnings equaling the cost to play) and each option contains the possibility of coming out ahead. Thus, the option not to play is dominated by game #1 and likewise by game #2. Given the goal, the agent would be irrational not to play one of these games. Thus, the option not to play should be eliminated by the dominance rule. Note, however, that game #1 does not dominate game #2, for if the side with 2 dots lands up the agent wants to be in game #2 not game #1. Also, game #2 does not dominate game #1, for if the side with 5 dots lands up the agent wants to be in game #1 not game #2. So, the rational choice between these two remaining options cannot be decided by the method of dominance. How should the agent decide which game to play?

To discover the rational choice for this decision problem, the element of risk must be brought into the picture. The risk associated with each outcome of each option should be allowed to influence the value of the outcome (as always, given the goal). Let's do this in steps.

- 1) The first step in forming the degrees of confidence for each outcome is assigning a probability to each state in each option. In this problem, this is done by calculating pure probabilities. For game #1, side 5 has a 1/6 chance of happening: $P(.17)$. Sides 1, 2, or 3 has a 3/6 (1/2) chance of happening: $P(.5)$. Sides 4 or 6 has a 2/6 (1/3) chance of happening: $P(.33)$. Thus, the agent should have a .17 degree of confidence that playing game #1 will yield \$10 winnings, a .5 degree of confidence winning \$5, and a .33 degree of confidence winning \$7. Note that these degrees of confidence sum to 1.0 in keeping with the disjunction rule of combining probabilities

that was introduced in Chapter 5. Thus, the agent is certain that one of these three states must happen in game #1, and rationally distributes this certainty (= 1.0) among the three possible states as the risk factor of each.

Doing the same calculations for game #2 results in a .17 degree of confidence winning \$10, a .5 degree of confidence winning \$2, and a .33 degree of confidence winning \$3.

- 2) The next step is to adjust downward the value of the outcome in proportion to the risk. The central idea is: risk devalues the outcome because the higher the risk the less goal achievement, on average, the outcome yields. This reduction or discounting is done by multiplying the dollar amount of the outcome times the degree of confidence of achieving that outcome. For game #1, we multiply \$10 x .17 and this gives us \$1.70. A \$5 outcome x .5 = \$2.50, and finally a \$7 outcome x .33 = \$2.31

The same step for game #2: $\$10 \times .17 = \1.70 . $\$2 \times .5 = \1 . $\$3 \times .33 = \$.99$.

- 3) The third step is to add these risk-discounted outcome values for each option to see what the option is worth no matter which state happens. We get a total of \$6.51 for game #1, and a sum of \$3.69 for game #2. If this were the end of the decision process, it is clear that the agent should choose to be in game #1, for that is the option with the greater overall winnings and this was the agent's goal. In other words, the agent can expect an average winning of \$6.51 if game #1 were to be played over and over, given the probabilities of the states and the dollar amount outcome. Does this mean that \$6.51 will be won with each play? No, in fact \$6.51 will never be won! This amount is not one of the possible outcomes. Sometimes \$10 will be won (roughly one time out of six plays), and sometimes \$5 will be won (roughly one half the time), and sometimes \$7 will be won (roughly one out of three plays). The average winnings the agent can reasonably expect over many plays is \$6.51 for game #1. Likewise, the average winnings for game #2 is \$3.69. But this is not the end of the decision process. There is one more step, namely, factoring in the cost of playing each game.

4) The last step is to subtract from the average winnings of the game the cost the agent must accept in choosing an option; here the cost is the price of playing each game. For game #1 the cost is \$5 leaving a difference of \$1.51. The cost of playing game #2 is \$2 leaving a difference of \$1.69. \$1.51 is game #1's expected monetary value (EMV) and \$1.69 is game #2's EMV. So, when costs are accounted for the EMV of game #1 is less than the EMV of game #2.

The rational choice for this decision problem is: choose (#2) over (#1) and (#1) over (don't play).

We now define EMV in a general way and give the rational choice rule for solving decision problems by this method.

Definition of the **expected monetary value** (EMV) of an option:

For each option: (i) multiply the monetary value of each outcome by the risk (the degree of confidence) that the state required for that outcome will happen; (ii) add up these risk-discounted outcome monetary values; (iii) subtract any monetary value it costs to choose the option. The final monetary amount is the EMV of the option.

Rational choice rule for solution by EMV: Given the goal of maximizing monetary gain (and minimizing monetary loss), for any two options x and y :

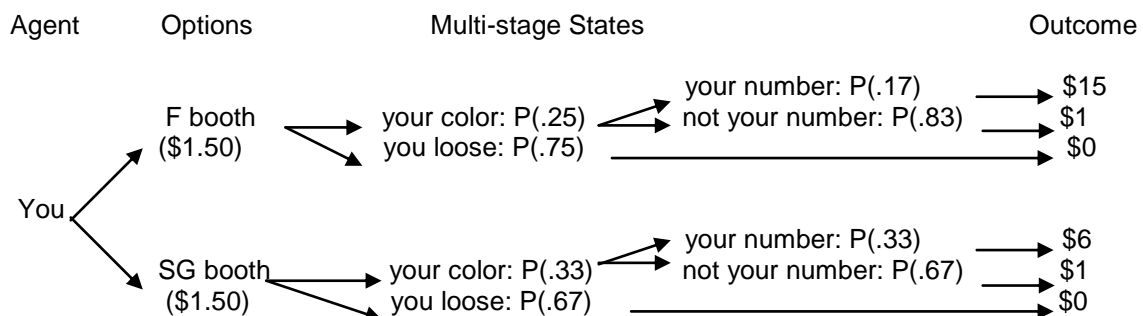
(i) if $EMV(x) > EMV(y)$, then choose x over y ,

and (ii) if $EMV(x) = EMV(y)$, then x and y are equally rational choices. For equally rational options, the agent has no practical reason to prefer one over the other and so should be **indifferent**. To be indifferent means that the agent has no problem substituting one option for the other (it does not mean that the agent is uncaring, unconcerned, or has apathy toward the decision).

We'll now look at another example of a decision under risk that is solved by EMV. This example will involve the agent losing money in order to achieve the goal, and will be multi-staged. Here is the decision problem.

Suppose the Student Government at your college puts on a fund-raising event. Various campus groups have booths at the event, and the money these booths raise goes to Student Government. Some booths sell things, and some booths have contests and games designed to make money. You attend this fund-raising event and your goal is to help make the day a financial success for Student Government. In other words, your goal is to spend and to lose money! The stakeholder, the intended beneficiary of your decision, is the Student Government, not yourself. Two booths attract you, the Faculty booth (F) and the Student Government's own booth (SG). The SG booth offers a game you can play. For \$1.50 you get to pick a number from 1 to 9 and spin a wheel that stops at an arrow. The wheel is evenly divided into three colors: red, green, and blue. Each colored part has three numbers evenly spread around the wheel from 1 to 9. So, picking a number automatically also picks a color. You spin the wheel. If it stops at your color but not your number, you win \$1. If the wheel stops at your number, you win \$6. If it stops at a color other than the one your number is in, you get nothing. The F booth is very much like the SG booth, only the wheel is evenly divided into four colors, it is evenly numbered from 1 to 24, and each color contains 6 evenly spaced numbers. For \$1.50 you get to pick a number (which automatically picks a color) and spin the wheel. If it stops at your color but not your number, you win \$1. If it stops at your number, you win \$15. If it does not stop at your color, you win nothing. You want to play one of these games. Which one is the rational choice for you to play, given your goal? Here is the decision diagram:

Goal: Help make the Student Government fund-raising event a success.



To solve this decision problem, we must:

- (1) find the rational degree of confidence (risk) of the belief in each outcome the agent should have,
- (2) reduce the value of each outcome in proportion to the risk involved in achieving it,
- (3) sum these “risk-corrected” outcome values for each option to find what each option is worth, and
- (4) subtract the cost. The rational choice will be the option that gains the agent more of the goal than the other options.

(1) For the F booth, there is a $1/4$ or .25 probability of hitting your color. Applying rule 2 for combining initial probabilities (see Chapter 5 for review) leaves a $3/4$ or .75 probability of losing. In keeping with the disjunction rule for combining probabilities, these two possible states sum to 1. Given that the wheel stops at your color, there is a $1/6$ or .17 probability of hitting your number. (Note: here we have a dependent event. Independent of color, there is a $1/24$ chance of hitting your number.) By rule 2, this leaves a $5/6$ or .83 probability of missing your number, given that the wheel has stopped at your color. Again, in keeping with the disjunction rule, these two possible states sum to 1.

To find the rational degree of confidence with which a \$15 outcome should be expected in the F option, the conjunction rule for combining probabilities must be used, for there is a two-staged state to deal with. Multiplying $P(.25) \times P(.17)$ yields a .04 degree of confidence. Doing the same thing for the \$1 outcome $P(.25) \times P(.83)$ gives a .21 degree of confidence. Because the \$0 outcome is single-stage, its degree of confidence equals its probability: .75.

For the SG booth, there is a $1/3$ or .33 probability of hitting your color, leaving a $2/3$ or .67 probability of a \$0 outcome. Given that the wheel stops at your color, there is a $1/3$ or .33 probability of hitting your number, leaving a $2/3$ or .67 probability of missing your number. You can easily verify by now that the disjunction rule for the alternative possible states in this option has been obeyed. The rational degree of confidence with which the agent should expect the \$6 outcome under this option is

.11 ($P(.33) \times P(.33)$). Doing the same for the \$1 outcome gives a .22 degree of confidence ($P(.33) \times P(.67)$). The \$0 outcome, being single-staged, has a .67 degree of confidence.

(2) To reduce or discount the value of each outcome of each option in proportion to the risk (degree of confidence) we must multiply the value by the risk. For the F booth option: $\$15 \times .04 = \$.60$; $\$1 \times .21 = \$.21$; and $\$0 \times .75 = \0 . For the SG booth option: $\$6 \times .11 = \$.66$; $\$1 \times .22 = \$.22$; and $\$0 \times .67 = \0 .

(3) Now we must sum these risk-adjusted outcome values for each option. The F booth option: $\$.60 + \$.21 + \$0 = \$.81$. The SG booth option: $\$.66 + \$.22 + \$0 = \$.88$.

(4) Finally, we must subtract what it costs the agent to choose each option. For the F booth it is $\$.81$ minus $\$1.50 = \$.69$. So, playing the Faculty booth game over and over loses you, on average, $\$.69$, and thus earns for the Student Government at your College, on average, $\$.69$ each time the game is played. Doing the same for the SG booth we arrive at $\$.88$ minus $\$1.50 = \$.62$. Playing the Student Government booth game over and over loses you and earns the Student Government, on average, $\$.62$ per play.

Now we can state the solution to this decision problem; but first recall the goal. If your goal had been to go to the Student Government fund-raising event and *win* as much money as you can (losing as little money as possible), then the rational choice is for you not to play the game having the biggest possible win ($\$15$ at the F booth). The rational choice would have been to play the game that minimized your loses, once you realized that both of your options were “rigged-games” designed to make money for the Student Government. So, *if* your goal had been to leave the fund-raising event with as much money as possible, you would head for the SG booth to spend your time and money there. Because $EMV(SG) > EMV(F)$, given this *other* goal, the rational choice would have been: choose (SG) over (F). But in this decision problem you had a very different goal in mind. Recall that you wanted to help make the Student Government fund-raising event a financial success.

Your *loss* is their success. Thus, given your goal, $EMV(F) > EMV(SG)$. The rational choice is: choose (F) over (SG).

It may seem odd at first to think that, in this example, you help the Student Government at your college most by bypassing the Student Government booth in favor of the Faculty booth. You can easily imagine many people with the same goal that you had in this example making an irrational choice. They would believe that the best way to help the Student Government raise money is to play the game at the SG booth. Unfortunately, this is poor practical reasoning and ends up financially hurting Student Government. The clear rational choice in this example is to give your business to the Faculty booth, given your goal. In fact, if everyone who attended this fundraising event played the F game and no one played the SG game, the Student Government would be better off than if even one player switched from the F game to the SG game.

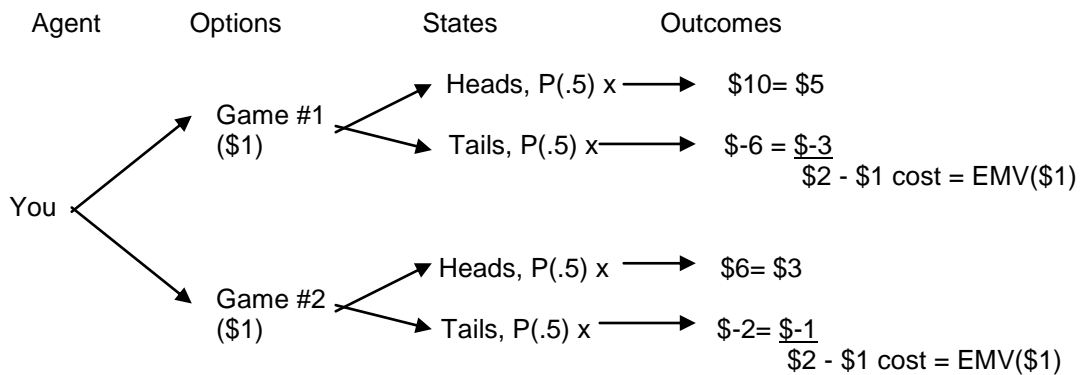
6.2.1 Limits to decision by EMV

Solving a decision problem by EMV is a powerful method of practical reasoning. It works well for evaluating options under risk when dominance reasoning can't apply and the goal is financial profit and outcomes can be assigned monetary values. This is very often the case, for example, with decisions in gambling, in games involving money, in business and finance, and in buying and selling. But solution by EMV has two weaknesses that limit its use. First of all, not all decisions under risk have monetary goals or have options with outcomes that can be valued in dollar amounts. Most do not. Thus, EMV as a way of solving decisions under risk is very restricted in application.

The second weakness is more troubling, for it is a weakness that can come up in cases where solution by EMV seems appropriate. The problem is that money has relative value, and this value is not always equal to the actual dollar amount. For example, \$10 will be valued (desired) highly by a person who is very poor, but the same \$10 will have little value to a billionaire. The dollar amount has not changed, \$10 is financially worth \$10; it has not changed in monetary value due to, say, economic inflation or depression as we go from the case of a poor person to a billionaire. Yet relative to poverty \$10 has great value, whereas relative to vast wealth \$10 amounts to nothing.

Let's look at an example where the relative value of money makes EMV fail as a way of solving a decisions problem.

Suppose you desperately need \$10 right away, a sudden emergency has come up. You don't have \$10, all you have is \$7. On the street corner in front of you two simple games of "flip the penny" are going on. Each costs \$1 to play. In game #1 a fair penny is flipped and heads wins \$10 for the player, but tails means the player must pay \$6. In game #2 a fair penny is flipped and heads wins \$6 for the player, but tails means the player must pay \$2. Suppose that you can only play one of these games once. Which one should you play? Here is the decision diagram and the solution by EMV.



Solution by EMV: $EMV(\text{game \#1}) = EMV(\text{game \#2})$, so (game #1) indifferent to (game #2).

According to the EMV rule of rational choice, you should treat these two flip-the-penny games as equally rational options. Yet this seems wrong. Because you need \$10 right away, you should clearly choose game #1 over #2. Game #1 gives you a 50/50 chance of achieving your goal, but game #2 gives you no chance whatsoever of achieving your goal in one play. The *utility* of the heads outcome in game #1, given your goal, is greater than the dollar amount of \$10 (relative to your sudden emergency). But EMV does not have a way to represent the relative value of money, and this weakness limits its power to solve decisions under risk even when monetary values are involved.

Here is a variation of this same decision problem showing that solution by EMV makes indifference rational, when clearly it isn't. This variation is designed to show the first weakness of solution by EMV. Suppose you approach these two flip-the-penny games taking place on a street corner, this

time however there is no emergency need of \$10. You notice that game #2 looks more pleasant. The players are having fun and they seem peaceful types. In contrast, game #1 look tense and some of the players seem to be violent types on the verge of anger. You are feeling lucky and want to play flip-the-penny to win some money. Should you treat these two games as equal because they have equal EMV's? Is indifference rational as solution by the EMV rule requires? Clearly not! Game #2 is the rational choice given such a decision situation. Yet solution by EMV can't capture the important non-monetary difference between these two games that makes the decision to play #2 the better choice.

Because of these two problems: (i) restricted application to monetary decisions and (ii) the relative value of money, we need a more powerful way to discover the rational choice in decisions under risk. In the next section, we will present a form of practical reasoning called solution by expected utility, and we'll see that this method of solving decision problems under risk does not have the limitations that EMV reasoning has.

EXERCISE: For the following decision problems, frame each into a decision diagram and solve it by the appropriate method of practical reasoning: either dominance or expected monetary value (EMV).

1) Your goal is to win money and you must choose between two games. Game #1 costs \$5 to play and goes like this. You flip a fair coin. If it comes up tails, you get one more flip. If the second flip is also tails, you win \$5, but if it comes up heads you win nothing. But if the first flip comes up heads, you must flip again. Tails means you win nothing, but heads on the second flip means you flip a third time. If its tails, you get nothing. But if third flip is heads, you win \$50. Game #2 costs only \$1 to play and it goes like this. You pick a card from a standard well-shuffled deck of cards (52 cards). If it is a number card 2 – 10 (there are 4 each) you win nothing. If you pick a Jack, Queen, or

King (there are 4 each) you win \$5. But if you pick an ace (there are 4), you put it aside and get to pick again. If you don't get a second ace (there were 3), you win nothing. But if you pick another ace, you put it aside and pick a third time (now there are 2 left). If you pick a third ace, you win big: \$10,000! But if you don't get an ace on the third pick, you win nothing. Which game is the rational choice for you?

2) Your family is in trouble; they desperately need money right away for a medical emergency and you have promised to help. Your goal is to sell your car by the end of the day, coming away with as much money as you can get by the end of the week. The car is fairly new; you are asking \$12,500 for it. There are 3 serious buyers and time is of the essence; you must decide on one by the end of the day, and the other 2 who don't hear from you will be retracting their offers and buying other cars. Here are the problems connected with each offer. Buyer #1 has offered you \$12,000, but requires that you make \$1000 in several repairs (this is your cost). However, this buyer must get a bank loan for the \$12,000 and tells you that his credit history is "shaky" due to a recent bankruptcy. He tells you, and you believe him, that he has only about a 1 in 3 chance of getting the loan. If he doesn't get it, the outcome is no sale (think of this as \$0); but if he gets the loan, he'll pay you \$12,000 once the repairs have been done. Buyer #2 has offered you \$10,500, but requires \$500 in repairs. She likewise must get a bank loan or the outcome is no sale (= \$0). Her credit history is not as "shaky" as buyer #1; there is a 50/50 chance she'll get the loan and pay you \$10,500 once the repairs are done. Buyer #3 has offered you \$8,500 for your car, and will take it "as is", no up-front repairs required. He tells you that he plans to borrow the \$8,500 from his family, and there is a strong likelihood, say 80%, that they will lend him the money once he tells them about your car at this great price. But if they don't lend, the outcome is no sale (= \$0). Given your goal, which potential buyer of your car is the rational choice for you to make?

3) A family has recently moved into town and must decide which of 3 elementary schools to put their 3rd grade child into. Their goal, of course, is the best 3rd grade education for their child. One option is the local public school, for which they see 3 possible outcomes depending upon the teacher situation. If the existing teachers remain for the year, the outcome is that their child will get an

average 3rd grade education. If some teachers leave and new teachers are hired, their child will receive a slightly below average 3rd grade education due to inexperienced teachers. But if some teachers leave and none are hired to fill the spots, the outcome is that their child will get a poor 3rd grade education. A second option is to put their child in a nearby private school. In this case, they foresee two possible outcomes, again depending on the teacher situation. With the current teachers in place, the outcome will be an average 3rd grade education for their child. But if the private school hires some new teachers, they will be experienced; the outcome is that their child will receive an above average 3rd grade education. The last option for their child's 3rd grade is a new experimental school designed and funded by the large state university. If the financial support remains at the current level, the outcome is that their child will receive an excellent 3rd grade education. But if the financial support declines, which is likely to happen, their child will receive only an above average 3rd grade education. What is the rational choice this family should make on behalf of their 3rd grader? Is it necessary to estimate the probabilities of the various teacher situations and the financial support situation before a rational choice can be made?

6.3 Expected Utility

Recall that *utility*, in our special use of this term, is a measure of how strongly an outcome gains an agent the goal. Monetary value fails as a measure of utility, even in some cases where a monetary value can easily be assigned to outcomes and to goals. Assigning a utility number to an outcome is a way of estimating the value of that outcome, given the goal, which avoids the problems described above that arise with using money as a measure of outcome value. Before solving sample problems by this more powerful method of practical reasoning, here is a general outline for assigning utility values to outcomes and for calculating the expected utility of an option. The agent wants to discover which option has maximum expected utility; the method will be familiar from Chapters 2 on complex goal analysis and objective ranking.

- 1) Using the goal as a single criterion, qualitatively rank each outcome with respect to the single attribute the criterion isolates. The agent might ask: for any two outcomes, which one would gain me more of the goal and which less of it? Once two such outcomes are ranked, the agent could go to another outcome and ask: would this outcome gain me more of my goal than the greater of the two, less than the least of the two, or does it fit somewhere between the two? This process would continue until all outcomes are qualitatively ranked; the resulting descriptions are verbal indicators of utility that now have to be transformed into quantitative units.
- 2) Ordinally rank these outcomes using "1st" for the outcome with *least* value and the largest ordinal number for the outcome having *most* value for achieving the goal.
- 3) Select a sufficiently large interval scale. The scale should include negative numbers to represent disutility, e.g. (-10...0...10), or perhaps (-25...0...25).
- 4) From the interval scale, assign a number to each outcome in such a way that (i) the ordinal rank of outcomes is preserved, and (ii) the intervals between the numbers reflect the different degrees to which the outcomes gain or lose the goal as this information was contained in the qualitative ranking. These interval numbers are now the **outcome utilities**.

Once utility numbers have been assigned to all outcomes, the expected utility of each option is calculated the same way EMV is calculated. For each option:

- 1) Multiply the agent's degree of confidence that the outcome will result, times the utility of the outcome. (Expressed more simply: multiply risk x utility. Recall what these numbers represent, and you should be able to see that the agent is multiplying strength of belief times strength of desire, in effect discounting the value of an outcome (desire) by the degree of risk (belief) that the outcome will happen if the option is acted on.)
- 2) Sum these products. This sum is the **expected utility** (EU) of the option. Note that it is not necessary to subtract any financial cost connected with the option from this sum, because automatically any cost gets factored into the utilities assigned. If an option has a huge cost attached to it, then the possible outcomes would not be given as high a utility rating as each

would get if no cost were connected to the option. So, the issue of cost is already taken care of in assigning utility numbers to outcomes.

The rational choice rule for solving risky decision problems by expected utility is very much the same as we saw for expected monetary value.

Rational choice rule for solution by EU: Given the goal as single criteria, for any two options x and y :

- (i) if $EU(x) > EU(y)$, then choose x over y , and
- (ii) if $EU(x) = EU(y)$, then be indifferent between x and y .

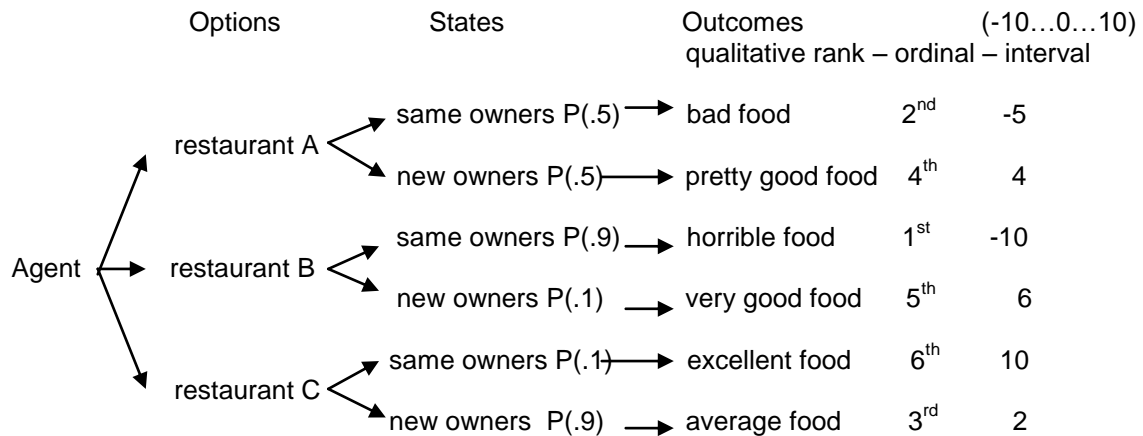
Now we'll use the method of expected utility to solve a variety of decision problems.

6.3.1 Single criterion individual decision under risk, single stage: decision by expected utility

Let's start with the example that was used above in section 6.1 for solution by dominance and add a bit more detail. Imagine an agent whose goal is to go out for dinner getting the best food for herself, never mind price. She can't travel far and as a result has just three restaurants she could go to: A, B, or C. But these restaurants frequently change hands so that for each there might be new owners or there might be the same owners. From past experience and from reading the restaurant critic's column in her local newspaper, the agent has a good idea about the quality of the meal she will get depending on the owners. There is an even chance that A will have new owners, in which case she will get a pretty good meal, but if the same owners are still there she will get a bad meal. For restaurant B, there is a very strong chance that it is in the hands of the same owners and if so she knows she will get horrible food. But there is a slim chance that new owners have taken over and these owners turn out very good food. Finally, for C there is a small probability that the same owners are still there, in which case excellent food is assured. But much more likely new owners have taken over C and they are well known for serving only average quality food. Which restaurant should this agent go to?

Here is the decision diagram. The agent's degrees of confidence, represented by the probabilities of the states, have been inserted. You should verify for yourself that these degrees of confidence correctly match the descriptions given in this narrative. Likewise, the qualitative, ordinal, and interval ranking of the outcomes are given. Again, you should check for yourself that they follow the 4 steps presented in the previous section.

Goal: to eat out getting the best meal for myself.



The first thing to note is that this decision problem can't be solved by dominance. No restaurant is dominated by the others, and so none can be eliminated. No restaurant dominates the others, so none can be chosen as superior just by comparing outcomes relative to the goal. The next thing to note is that two outcomes distance the agent so far from the goal as to warrant using negative numbers. Third, you should not think that these exact interval numbers are correct in the sense that different interval values would somehow be wrong. We could easily have used a scale (-5...0...5), or (-15...0...20) to capture the information concerning the relative utilities of these outcomes contained in the narrative, and this would have yielded different outcome utility numbers (but not a different rational choice solution to this problem!). Finally, note that the agent is certain that either the same or new owners exist for each restaurant. Thus, the degrees of confidence for each option sum to 1.0. This is in keeping with the disjunctive rule set out in the last chapter.

Now let's calculate the expected utility of each option. Multiply the degree of confidence times the utility for each outcome (this adjusts downward the value of the positive outcome – and upward the value of a negative outcome – for gaining the goal by the risk factor), and sum these products for

each option. So, for restaurant A we have $P(.5) \times -5 = -2.5$, and $P(.5) \times 4 = 2$. These products sum to: $EU(-.5)$. For restaurant B we have $P(.9) \times -10 = -9$, and $P(.1) \times 6 = .6$. These sum to: $EU(-8.4)$. For restaurant C we have $P(.1) \times 10 = 1$, and $P(.9) \times 2 = 1.8$, which sum to: $EU(2.8)$ According to the rational choice rule for solution by expected utility, we get this order of choice: restaurant C should be chosen over A, and A chosen over B; that is: The solution to this decision problem is to choose restaurant C; it has maximum expected utility.

EXERCISE:

Structure and solve the following single-stage state decision problem by expected utility.

Remember that the probabilities for the alternative states within each option must sum to 1.0. You are distributing the unit 1.0 (= certainly) to the states in a way that captures the informal information about the likelihood of the states happening as provided in the narrative.

It is Saturday night and you want to go out and have some fun. You could go to the movies. A science fiction film is playing. You are not into science fiction very much, but there is a slightly better than 50/50 chance that a good friend will meet you at the theater who really enjoys science fiction movies. If you meet your friend, you'll have a pretty good time at the movies, but if she doesn't show up the result will be a mildly disappointing evening. Another option you have is to go to a party. The problem is that you just broke up with someone and this person might also be at the party. You know that there is a small chance this person will arrive at the party early. If so, you will have a miserable time and will end up going home early very upset. On the other hand, it is more likely than not that this person will show up at the party late, in which case you will have a great time until the person arrives and a miserable time only at the end of the evening. Finally, there is a small chance, equal to the likelihood of your former partner arriving at the party early, that this person will not show up at the party at all. This means as outcome a really fun time for you for the whole party. Your last option is a concert. Tickets are expensive, and this detracts a bit from the enjoyment of the concert. But getting a decent seat is your main worry. You love the band and the show will be a great time

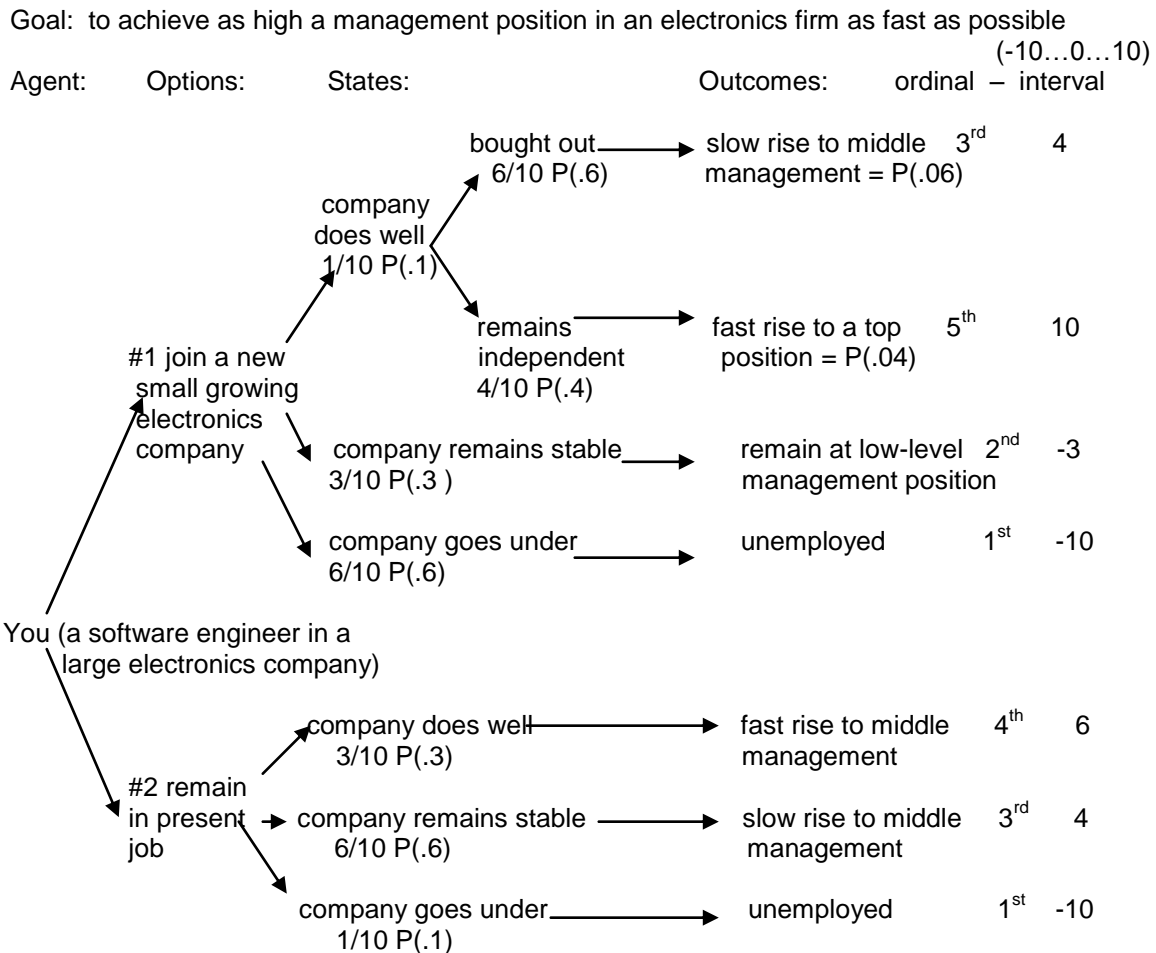
whether or not you connect up with friends, but only if you get a good seat. At this late date, there is only a small chance that a good seat is available. There is an even likelihood that you will get an average seat or a poor seat. An average seat will mean that you'll have a good time at the concert, and a poor seat will mean that you'll have a pretty disappointing time. Given your goal and these options with accompanying risks, what is the rational choice for you – the movies, the party, or the concert?

6.3.2 Single criterion individual decision under risk, multi-stage: decision by expected utility

Now let's practice solving a decision problem by expected utility that involves calculating the agent's degrees of confidence with respect to multi-stage states-of-the-world. Recall from Chapter 5 that the conjunction rule (multiplying probabilities) will be needed to form a reasonable degree of confidence that an option will yield a given outcome.

Pretend that you are a software engineer in a large well-established electronics company. Your goal is to achieve as high a management position as fast as possible in the electronics industry. Recently you have been offered a job in a new small growing electronics company. Of course you can also remain with your present employer. You have been researching the industry, trying to make up your mind whether to stay or to go to the new company. Here is what you have found out. For large electronics company like the one you work for, roughly 1 in 10 goes out of business. If this happened to your company, you end up unemployed. 6 in 10, on average, remain stable, neither going out of business nor having rapid growth. If this happens to your company, you'll have a slow rise to middle management. But 3 in 10 companies like yours experience healthy growth. Should this happen you would have a fast rise to middle management in your present job. Here is the data you have gathered about new small growing electronic companies. On average, 6 in 10 don't make it and go out of business. If this happens, you would be unemployed. 3 in 10 remain stable, with growth slowing down to a flat rate. If this is the case, you would remain at a low-level management position. But 1 in 10 such new small companies do well. When they do well, roughly 6 in 10 are bought out by

larger companies. If this should happen, you would have a slow rise to middle management. On the other hand, 4 in 10, on average, remain independent. Should this happen, you would have a fast rise to a top management position. What should you do, given your goal, your options, and the information you researched, stay or change jobs? Here is the decision diagram for this problem with the probabilities of the states, based on the above narrative, filled in.



First, verify that this decision problem can't be solved by dominance. If each outcome utility of one option were equal to or greater than each outcome utility of the other option, and at least one outcome utility were greater, then one option would dominate the other yielding a solution. But this is not the case. So, we must turn to solution by expected utility. Next, note that each option contains the same outcome – unemployment. These similar outcomes should receive the same ordinal rank, and the same interval disutility number, for they equally distance the agent from the goal.

In this example, most states are single stage. Thus, the agent's degree of confidence for each outcome equals the probability that the required state will happen. These values have been inserted bases upon the narrative. But in the case of option #1, if the company does well there is another state that is required for the best outcome to result. But the risk is that this additional state will not happen, a different one will, in which case a less than best outcome will result. To form a degree of confidence in these two multi-stage states, we use the conjunction rule for dependent states: $P(a \text{ and } b) = P(a) \times P(b/a)$. So, for the state in which the new company is bought out, we have $P(.1)$ that it does well times $P(.6)$ that it is bought out given that it does well = $P(.06)$ as the degree of confidence that you will have as outcome a slow rise to middle management. Likewise, for the state that the new company will remain independent: $P(.1) \times P(.4) =$ a degree of confidence $P(.04)$ that changing jobs will result in a fast rise to a top management position. These degrees of confidence values have been inserted at the relevant outcome descriptions on the decision structure.

Now we are ready to discover the option that has maximum expected utility. For each of the 4 outcomes in option #1, multiply the degree of confidence times the outcome utility, and sum these 4 products. We do the same for the 3 outcomes in option #2. Option #1: $(.06 \times 4 = .24) + (.04 \times 10 = .4) + (.3 \times -3 = -.9) + (.6 \times -10 = -6) = EU(-6.26)$. Option #2: $(.3 \times 6 = 1.8) + (.6 \times 4 = 2.4) + (.1 \times -10 = -1) = EU(3.2)$. The rational choice for you in this decision problem is clear: $EU(\#2) > EU(\#1)$, so choose option 2 over option1. Note that remaining in your present job should be strongly preferred over accepting the new job, given your goal, for the amount of goal achievement of these two options have a wide difference in value. While the small electronics company has the potential of your gaining your full goal (an outcome utility 10 is possible) the risk of unemployment is great. In this story, it is true that your present job will not gain you your full goal, but it has the potential of bringing you reasonably close (an outcome utility 6 is possible) while minimizing the risk of unemployment (= total goal loss!).

Solution by expected utility is the most powerful method of practical reasoning about risky decisions. It clearly balances the agent's desire of the outcomes, given the goal, with the agent's beliefs

concerning the likelihood of the outcomes happening. But for solution by expected utility to work, it requires that:

- a) The agent's degrees of confidence be well-based. This means that the agent must have accurate and reliable evidence in assigning initial probabilities to the states the outcomes require.
- b) Outcome utility values accurately represent the strength with which the outcomes gain or lose the goal for the agent. This means that the outcomes must be correctly described with respect to the simple goal whose single objective is the basis for forming a single criterion (that is: a single attribute whose value or weight is 1.0).

EXERCISE: Structure and solve the following single criterion multi-stage decision problem by expected utility. Remember that the probabilities for the alternative states within each option must sum to 1.0, and that you must use the conjunction rule for combining probabilities in each case of multi-stage states. The probabilities you assign should capture the informal information about the likelihood of the states happening as provided in the narrative.

The government of a nation (A) has come to realize that it is both in the national interest as well as in its political interest to do something about another nation (B). B has violated the borders it has with a neighboring nation (C) whose resources are valuable to A, and with whom A has both protection and trade treaties. A's goal is to restore the former borders between B and its neighbor country C in such a way that it boosts its own political popularity with voters at home. The problem is: what should A do? Government leaders and advisors have discussed many possibilities and scenarios, and many of these were ruled out by option disqualifying rules: actions like sending an assassination team to B to kill its leader, trying to overthrow B's government, or paying a third country to intervene and restore the former border. A has narrowed its acceptable options to these four: either send troops to B, or establish a blockade around B, or send a diplomatic negotiation team to B, or finally

announce that A will do nothing for now and take a wait-and-see position to discover if the situation in B will change for better or worse. After researching these options, here is what the advisors and analysts have come up with. (1) If troops are sent, there is a pretty good chance that a battle will take place. If a battle does not take place, a lucky event for A given that troops have been sent, the outcome will be restored borders with only minor criticism from voters, for sending troops is perceived by some voters as an overly aggressive course-of-action. But if a battle takes place, there is a high probability that some of A's troops will be killed, and a very small chance that there will be no casualties. In the event of no casualties, the outcome will be the same as the state in which no fighting takes place. If A's troops are killed, there could be high casualties or low casualties. Fortunately, there is not much chance for high casualties, but if there were, the outcome would be restored borders but A's voters would oust the government in the next election. If casualties are low, most likely to happen given a battle, the outcome will be that the borders are restored but the government will be strongly criticized by vocal voters, the media, and political opponents for the deaths. (2) What about the option of blockading B? A blockade could be violated by B. The outcome would be the status quo as to the borders between B and C, and the government widely criticized for failure. On the other hand, the blockade might be respected, in which case the borders would be restored and the government would receive much praise from voters, for a blockade both shows strength and yet is not perceived to be overly aggressive. Unfortunately, the chance that the blockade will be respected is only slightly greater than the likelihood that it will be violated by B, so the latter is a very real danger. (3) If diplomats are sent to B, option #3, there is only a small chance of successful negotiations. The outcome here is that the borders are restored, but the government will appear weak to the majority of A's voters. In all likelihood, however, the diplomatic team will be rebuffed. The outcome in this case is that the status quo in B remains with respect to the borders with C, as well as the perception among A's voters at home of a weak government. (4) Finally, A could opt for a wait-and-see course-of-action. In this case the outcome is certain to be the status quo in B with respect to the border problem, and the government is sure to be strongly criticized by some political opponents for "doing nothing". However, it is also certain that the majority of A's voters will neither criticize nor praise this action, for they will also take a wait-and-see attitude.

You are a practical reasoning expert and the government of A calls you for advice about which option is best. They send you the above information, and are ready to send you a large check for your services. Given A's goal and these 4 options, what is the preference order and rational choice you would be recommending?

Sources and Suggested readings:

Making risky decisions, as you can imagine, is a large and important part of practical reasoning. There are many presentations of decision under risk, elementary to advanced, that are either geared toward business decisions, toward military decisions, toward political decisions, or toward personal decisions studied by psychologists. This chapter draws on: Jeffrey (1983) Chapter 1, Luce and Raiffa (1957) Chapters 2.4 and 13, Mullen and Roth (2002) Chapter 6, and Resnik (1987) Chapter 4. Both Skyrms (2000) Chapter VI.5 and Hacking (2001) Chapters 8, 9, and 10 offer clear presentations that are philosophically oriented. Allingham (2002) Chapter 3 is quite compact, but Chapter 4 extends his presentation in the context of gambling and insurance. For more advanced presentations of risky decisions, especially in the context of business and policy decision making with special emphasis on risk, see Keeney and Raiffa (1993) Chapter 4 (but note the change in terminology), and Raiffa (1997) – considered a classic – Chapters 0 to 4.