CHAPTER 7: RISKY DECISIONS BY EXPECTED UTILITY: COMPLEX GOALS

In this chapter we will combine material from Chapters 4 and 6. In Chapter 4 we covered individual decisions under conditions of certainty and introduced a method of discovering the rational choice for decision problems having complex goals. Complex goals require using multi-criteria option evaluation. The result of this method of option evaluation is, for each option, a final outcome utility. The rational choice is the option having maximum final outcome utility. In Chapter 6, risk was substituted for certainty, and three methods were introduced to discover the rational choice for decision problems with simple goals: dominance, expected monetary value (EMV), and expected utility (EU), the latter being the most powerful method when there is no dominant option. In Chapter 6, solution by expected utility required assigning utility values to outcomes according to a single criterion.

Now we will work with decision problems that combine risk and complex goals: multi-criteria risky decisions. For these decision problems, instead of assigning utilities to outcomes according to a single criterion, the agent must form a set of criteria based upon a goal analysis in order to discover final outcome utilities. These final outcome utilities will then be used, along with the agent's degree of confidence concerning risk, to discover the option having maximum expected utility.

7.1

There is no new theory here, it is a matter of gaining practice solving more challenging decision problems than we have dealt with so far, by using a combination of practical reasoning methods that have already been introduced and practiced separately. Before turning to decision problems for practice, to those having single stage states first and then to decisions having multi-stage states, let's review the practical reasoning methods and the rational choice rule for making multi-

criteria individual risky decisions by expected utility. Imagine an agent who has a complex goal and has formed a menu of options.

There are six steps for framing and evaluating a multi-criteria individual decision problem under risk, making a decision by expected utility.

1. Clearly state the goal: identify agent, stakeholder(s), whether the goal has intrinsic or relative value for the agent.
2. Goal analysis:
    a) list objectives and calculate their values (normalize to 1.0).
    b) form attributes and calculate attribute values (normalize to 1.0).
    c) form criteria by multiplying objective value x attribute value.
3. Option analysis:
    a) frame options, states, and outcomes in a branching diagram.
    b) assign probabilities to states and calculate degree of confidence for outcomes.
4. Outcome evaluation:
    a) on a grid with criteria along top and outcomes along left side, evaluate each outcome relative to each criterion using a large enough positive – negative interval scale.
    b) multiply each of these utility numbers by the criterion value to get weighted outcome utilities.
    c) sum for each outcome = final outcome utilities.
5. Option evaluation:
    a) put these final outcome utilities on the branching diagram, using the largest final outcome utility for the best outcome an option might yield (given the goal) and reducing the utility amounts for the option's other outcomes in intervals that match their reduced goal achievement.

b) multiply final outcome utilities x degree of confidence for each option.

c) sum these for each option = expected utility of the option.

6. Translate expected utility numbers into a rational choice ordering of the options according to the rational choice rule:

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

7.2 Multi-criteria decisions under risk: single stage states.

Suppose you are about to get married, and you and your fiancée have to decide on a honeymoon. You both agree that your goal is to have an exciting memorable honeymoon. Your honeymoon will take place in mid-winter, and you both agree that one objective is to go to a warm place that has good beaches. With the exception of cost, a warm beach means a lot to both of you, the warmer the better. As to cost, you both agree that the honeymoon can't be overly expensive, for you hope to use most of the money you anticipate as wedding gifts for other things. You set a $3000 limit on cost, but would like it to be as much under this amount as possible. This is your top priority. There are two remaining objectives you two do not agree on: sightseeing experiences and night-life. For sightseeing, your fiancée favors visiting art museums, but you would like to experience nature preserves (your fiancée was an art major in college, but you majored in environmental science). As for night-life, you like casinos and the shows they put on, but your fiancée likes music and dancing very much. You both agree that your different sightseeing and night-life desires will have equal weight in the decision process. After looking over many brochures and honeymoon packages, you have narrowed your options to these three: a European resort that offers a honeymoon package right on the Mediterranean, a Hawaiian resort honeymoon package, and a Caribbean honeymoon package. Here is what you have found out about each possibility.

(1) The European package costs $2750. It features excellent sightseeing in world renowned art museums containing works of the great artists, but poor nature preserve experiences. The average temperature during your stay varies between 65 and 75 degrees on the coast, but the beaches in the area where you will be are poor to average in quality. There is great night-life, many casinos offer wonderful shows and there are numerous places offering good music and dancing. The only risk that needs to be seriously considered is that the Hotel and Resort Union in the area has threatened a labor strike exactly during the time you would be honeymooning there. If a strike happens, and there is a realistic 25% chance that it will take place, your honeymoon will be pretty well ruined, except for enjoying the beach facilities.

(2) The Hawaiian package costs $3000. It features excellent beaches, with an average temperature during your stay between 75 and 85 degrees. The sightseeing, as with the European package, is mixed: there are world renowned nature preserves to experience, but few art museums in the area and they all specialize in local folk art. The night-life offers very good music and dancing opportunities, but no casinos and thus no shows of the kind usually found in casinos. The major risk that worries you both is that there has been a new Asian flu outbreak in Hawaii for which there is no vaccine yet. It strikes quickly, will spread to both of you if one catches it, and is pretty serious. It will ruin your honeymoon completely, keeping you both bed-ridden and miserable for your whole stay. You have found out from experts that for people in your health category, there is a one in five chance of coming down with this flu upon landing in Hawaii.

(3) The Caribbean package costs $2500. The beaches are not quite as good as in Hawaii, but much better than at the resort in Europe. The average temperature during your stay will be between 70 and 80 degrees. The sightseeing offers a balanced picture with decent nature preserves as well as decent art museums. Likewise, the night-life offers a balanced set of opportunities: there are good casinos with pretty entertaining shows, as well as good music and dancing clubs. The big risk for the Caribbean during the time you would be there concerns the weather. It will be storm season and the weather could get violent with conditions approaching

hurricane severity. Should this happen, all the damage and flooding would ruin your honeymoon in all respects. You have carefully looked up the probability of such severe weather and found that the risk is about a third.

You and your fiancée are pondering all this information, trying to balance everything out against the backdrop of your dreams for a wonderful honeymoon. Which honeymoon package should you go with?

The first step in solving this decision problem is to form a set of criteria by doing a goal analysis. Make sure to verify that the values match the information provided in the narrative.

Goal: For my fiancée and I to have an exciting memorable honeymoon. (Stakeholders: this couple. Outcome: quality of honeymoon. Value: intrinsic.)

Objectives	qualitative rank	ordinal rank	interval (1-10)	normalized
Warm beaches	very important	2	8	.30
Under $3000	most important	3	10	.38
Good sightseeing	equal to night-life	1	4	.16
Night-life	equal to sightseeing	1	4	.16
			26	1.0

Recall that the set of attributes must meet the six conditions set forth in Chapter 4. (This is a good opportunity for you to go back and review these six conditions, and to check that each of the following attributes conforms to all six conditions.) Under warm beaches, let's put two attributes: beach temperature and beach quality, the first being most important for our imagined engaged couple. For the objective of a $3000 limit, we have just one attribute: price, or amount under $3000. For good sightseeing opportunities there were two attributes: art museums and nature preserves, and these are to count as equally desirable. Finally, for entertaining night-life there are two attributes: casino shows, and music and dancing nightclubs, and these also are to be given equal value.

Attributes	ordinal rank	interval(1-5) rank	normalized rank	attribute value
1) beach temperature	2	5	.63	x .30 = .19
2) beach quality	1	3	.37	x .30 = .11
		8	1.0	
3) price		5	1.0	x .38 = .38
4) art museums	1	5	.5	x .16 = .08
5) nature preserves	1	5	.5	x .16 = .08
		10	1.0	
6) casino shows	1	5	.5	x .16 = .08
7) music and dancing nightclubs	1	5	.5	x .16 = .08
		10	1.0	

Next, we will use these criteria to evaluate outcomes and arrive at utility values for each. But before this can be done, we need to know what the outcomes are, so we must form a decision structure, our left-to-right diagram, and clearly describe each outcome, given the goal.

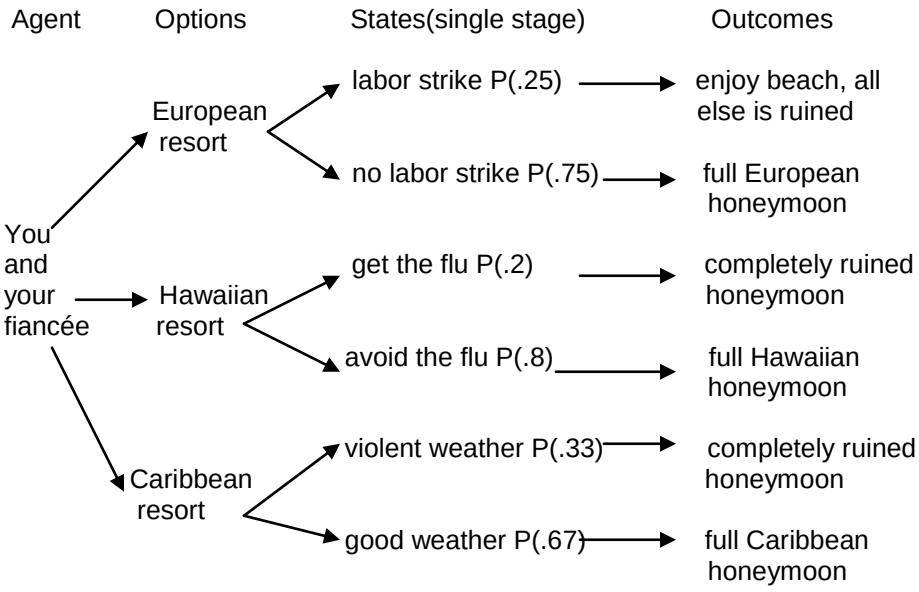

Now we are ready to evaluate these outcomes with the above seven criteria. Instead of writing out the description of each criterion, I will use its number and value along the top row of the grid. The outcomes are on the left column. Again, be sure that you verify that the positive and negative utilities assigned to the outcomes (values to the left of the slash), going down and evaluating

each outcome criteria-by-criteria, are in keeping with the information in the narrative. The values to the right of the slash are these utilities weighted by the criteria values.

(-10…0…10) Criteria:	1 .19	2 .11	3 .38	4 .08	5 .08	6 .08	7 .08	Final outcome utilities
Outcomes:								
Europe	6/ 1.14	-3/ -.33	5/ 1.9	10/ .8	-10/ -.8	10/ .8	10/ .8	= 4.31
Hawaii	10/1.9	10/1.1	1/.38	-10/-.8	10/.8	10/ -.8	8/ .64	= 3.22
Caribbean	8/1.52	6/.66	10/3.8	5/.4	5/.4	5/.4	5/.4	= 7.58
Completely ruined	-10/-1.9	-10/-1.1	-10/-3.8	-10/-.8	-10/-.8	-10/-.8	-10/-.8	= -10
European enjoy only beaches	6/1.14	-3/-.33	-6/-2.28	-10/-.8	-10/-.8	-10/-.8	-10/-.8	= -4.67

The final step is to use these outcome utilities (and disutilities) on the decision diagram with the degrees of confidence to arrive at the expected utility of each option.

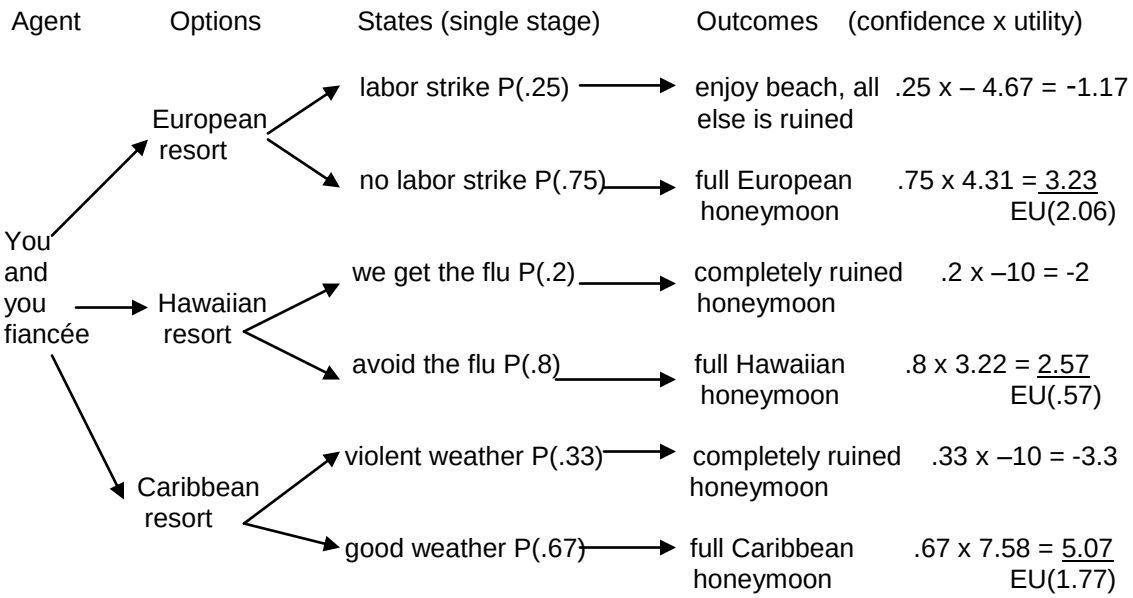

As you can see, these options come out very close in value. In such cases, the agent should go over the steps and try to use interval scales with wider ranges to see if the values might represent

the agent's beliefs and desires better with intervals assigned having larger gaps. If not, then it is a pleasant decision problem (assuming positive expected utility values for the options) in which the agent can't really go very wrong with any option. Given these expected utility numbers, the rational choice order is: choose (Europe) over (Caribbean) and (Caribbean) over (Hawaii).

7.3 Multi-criteria decisions under risk: multi-stage states.

You live in the suburbs of a large city and today you want to go into the city to meet with a friend and together go to a concert. The trouble is a snowstorm has just started. So, your immediate goal is to get to the concert hall the best way, which means for you the quickest, safest, and least troublesome way. These are your objectives, with safety being your greatest concern and ease of trouble your least worry. Your options are to drive, take the bus, take the train, or maybe it is best not to go at all. If you drive, there is an even chance that you will be stuck in backed up traffic. But if you do not get stuck in traffic, there is a slight chance, say one in ten, that your 4-wheel drive vehicle will get stuck in the snow. Should either one happen, stuck in traffic or in snow, you will miss the concert. If there are no problems getting stuck, driving is the quickest least troublesome way to go, but not the safest because other drivers can easily slide into you causing a minor accident. If the drive goes well, you will meet your friend and make the concert. If you go by bus, there is a chance, say one in five, that it will get stuck in traffic, and a greater chance, like three in five, of getting stuck in the snow. Should either happen, you will miss the concert. The bus stop is fairly close to your home, you will have to wait a bit in the snow for the bus, and you will have the inconvenience of getting from the city bus terminal to the concert hall by cab. The bus is safer than driving, but very slow because of bus stops along the way. If you take the train, by far the safest way to travel in a snowstorm, there is only a minimal chance, maybe one in fifty, of getting stuck in snow, and you are certain that you will not be stuck in backed up traffic. But the train is very inconvenient. The train station is quite far, requiring a long walk in a snowstorm. The trip will be faster than by bus, but not by much because of several train stops along the way. The

city train station is about the same distance from the concert hall as the bus terminal, so the inconvenience of the cab is the same with train as it would be by bus. Finally, you can always opt not to go into the city at all today, but that means missing the concert for sure.

You have been looking forward to this concert with your friend for a long time, and will be very disappointed to miss it. You checked and the concert is not cancelled due to snow, so you will lose your tickets at $40 each, which you bought and have in your possession, if you don't make it. Should you go to the concert, and if so by what means of transportation?

We begin to structure and solve this decision problem by forming a goal statement and doing a goal analysis to find the values for the objectives. As with the decision problem in the previous section, we extract the important information from the narrative, while realizing that we must project certain details of belief and desire strengths into the mind of our supposed agent, whom we are imagining to be you. Of course, any agent with a real decision problem of this sort would have the required beliefs and desires to draw upon as he or she reasoned practically to solve the problem. Be sure to check the details the analyses and evaluations below with the narrative above, and note where you might use different values had this really been your decision problem.

Goal: I want to attend a concert in the city with a friend for my own enjoyment getting there by the best means of transportation in this snow storm.

Objectives	ordinal rank	(1 – 10) interval rank	normalized.
Arrive quickly	2	7	.33
Arrive safely	3	10	.48
Minimum trouble	1	4	.19
		21	1.0

Next we must form a reasonable set of attributes for each objective, attributes that keep to the six conditions set forth in Chapter 4. Arriving quickly is clearly a matter of time, so let's use these

three attributes: time to transportation, time to concert hall (once in the city), and travel time from suburb to city. Let's give these attributes equal values. We will make the objective "arrive safely" its own attribute, for it is unnecessary to break this attribute down into the relative safety of this or that part of the trip to the concert hall. Finally, let's use three attributes for the objective "minimum trouble": walk in the snow storm, wait in the snow storm, and getting a cab. These would seem to be three reasonable areas of inconvenience in the narrative, and let's suppose that you (the imagined agent) rank the first as most important, the second as next, and the third as least important to you. With this information we are ready to form criteria from these attributes and evaluate the outcomes.

Attribute	ordinal rank	interval (1 – 5) rank	norm.	final attribute value
1. Time to transportation	3	5	.33	x .33 = .11
2. Time to concert hall	3	5	.34	x .34 = .11
3. Travel time	3	5	.33	x .33 = .11
		15	1.0	
4. Arrive safely			1.0	x .48 = .48
5. Walk in snow	3	5	.46	x .19 = .09
6. Wait in snow	2	4	.36	x .19 = .07
7. Getting a cab	1	2	.18	x .19 = .03
		11	1.0	

Next we do an outcome evaluation on a grid with the criteria values along the top row and the outcomes alone the left column, using the outcomes described on the decision structure below. The utility scale is (-10…0…10). The utility number you see in each cell should match the information in the narrative. For example, concert by train should get high marks under the criterion safety, while concert by car should get high marks for the three criteria under the objective of minimal trouble. Note that each criterion's number will be used on the grid in place of writing out the criterion's name. In each cell, the left value is the utility assigned according to the criteria and the value to the right of the slash is that utility weighted by the criteria value.

(-10…0…10) Criteria: / Outcomes:	1 (.11)	2 (.11)	3 (.11)	4 (.48)	5 (.09)	6 (.07)	7 (.03)	Final outcome utility
Concert by car	9 /.99	9/ .99	9/.99	-7/-.34	9/.45	10/.7	9/.27	= 4.05
Concert by bus	-2/-.22	-3/-.33	-5/-.5	3/1.44	-2/-.18	-2/-.14	-3/-.09	= -.02
Concert by train	-5/-.55	-3/-.33	1/.11	10/4.8	-9/-.8	10/.7	-3/-.09	= 3.84
No concert								-10

These final outcome utilities are now used for the outcomes on the decision structure which, together with the agent's degrees of confidence, gives us the expected utility of each option.

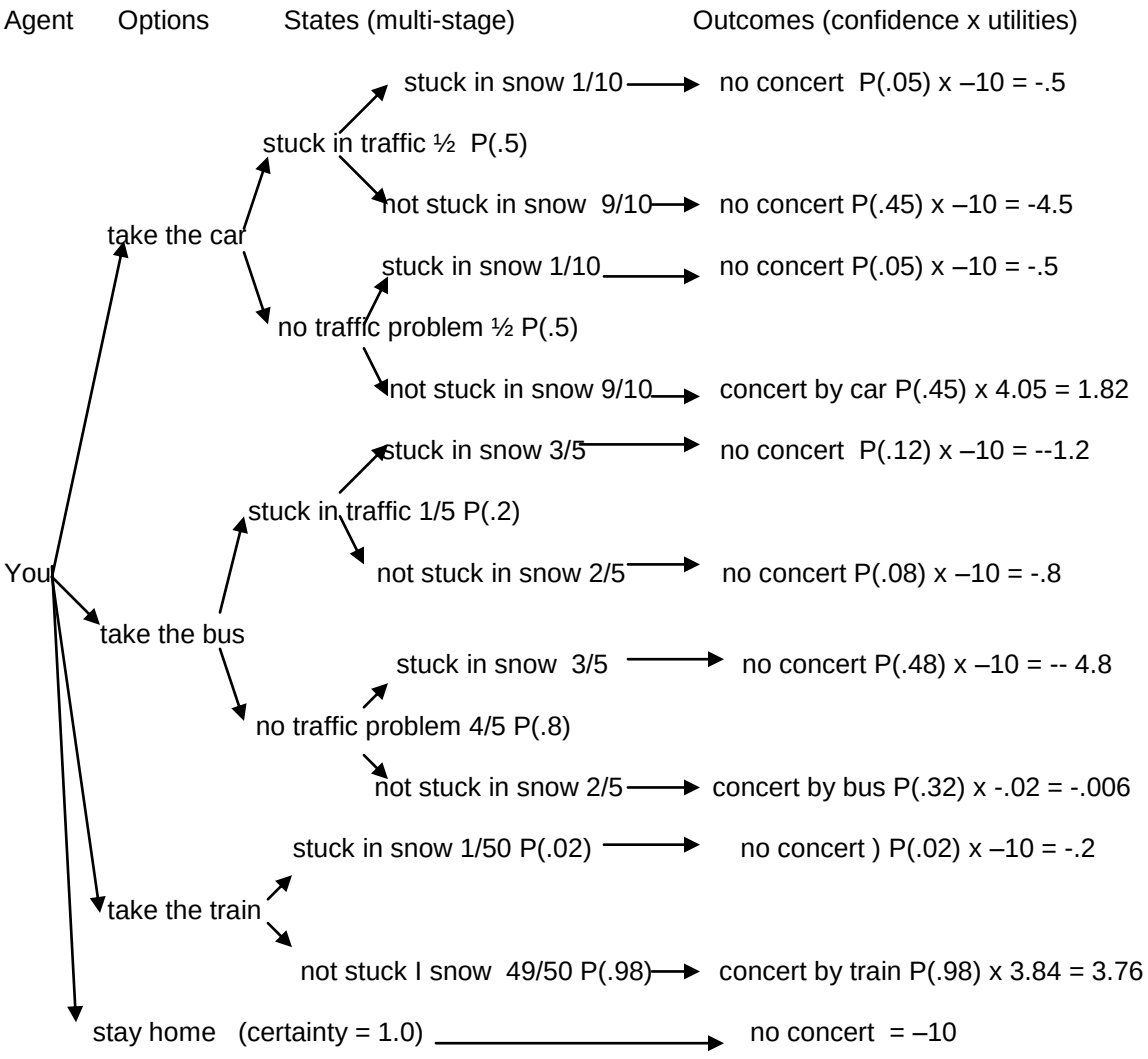

Note that for each option, the best outcome was assigned the final outcome utility; any outcome that achieves less of the goal should be assigned a utility value proportionally less than the final outcome utility.

For each outcome, the agent's degree of confidence was calculated using the conjunction rule for combining probabilities. This degree of confidence (the probability number coming after the description of the outcome) was multiplied by the utility value of the outcome. At this point it is clear that the option to stay home is dominated, so it drops out. But no other remaining option is dominated. Thus, we must add these products for each option to arrive at the option's expected utility.

1) Travel by car: -.5 + -4.5 + -.5 + 1.82 = EU(-3.68)
2) Travel by bus: -1.2 + -.8 + -4.8 + -.006 = EU(-6.8)
3) Travel by train: -.2 + 3.76 = EU(3.56)

The rational choice order is clear: choose (train) over (car) and (car) over (bus)

EXERCISE: Solve the following decision problem by expected utility. It is a multi-criteria decision under risk having multi-stage states.

You are a parent whose daughter is about to graduate from high school. You have saved $2500 for a really good graduation gift, and your objectives are these: the gift should be of educational value for your daughter in college and, if possible, beyond; it should be a source of fun and enjoyment; finally, it should contribute toward her sense of independence and self-esteem. If your gift accomplishes these three, your goal of the perfect graduation will be achieved. Here is what you have been thinking.

Your daughter loves music, she has studied piano for many years, and music is one of four possible majors she is equally considering in college. One option for a gift is a good piano

keyboard. If she majors in music, there is only a slim chance that her love of music will wane, and an overwhelming chance that music will be a life-long interest. In this case, a keyboard will be a gift of lasting value as well as a big help in college. Your are afraid, however, that she might not major in music, in which case there is still a strong chance that music will be a life-long love, but also a chance large enough to be worrisome that her love of music will wane after college. The outcome here is that a piano keyboard will turn out to be the wrong gift, having only some value for your daughter through college, but little lasting value afterwards.

Another gift option you are considering is a trip. Your daughter has hinted how exciting it would be to join two of her friends who are going to travel around Europe for a month before starting college. This would be a very broadening experience for her, and she would be on her own for the first time. She is good at learning languages, and another possible college major is French. Your gift could be a month in Europe with her friends. The worry here is that her friends are quite fun-loving, so there is a good chance – say 30% -- that they will spend the whole month in one spot where there is a lot of “action”. If the spot is in France, the outcome is that your daughter will miss the valuable experience of traveling in Europe, but at least will gain important knowledge of the language and culture of France. But if it is not in France, she will miss out on both the travel experience and the practice she would get in the language that might be her college major. Of course, if the group travels around Europe as they intend to do, the outcome will be both a valuable travel experience and practice in French, for part of the travel will be in France.

Another option is to get your daughter a high-end laptop computer. She spends hours with educational software and doing schoolwork on the computer at home. She enjoys surfing the web and emailing friends. But there are two potential problems with getting her a laptop. First, the college she will be attending is in the middle of a decision to provide all students with dorm-room computers. Money is tight at this college, and you have heard from good sources that there is a 60% chance that college computers will be provided in your daughter’s junior year, and a 40% chance that they will not be provided al all. If they are not provided at all, the outcome is the full

value of this gift for all four years of college, after which the laptop will become outdated technology. The second problem concerns theft. You have discovered that laptops are stolen during the first few weeks of freshman year at the college at a rate of 1 out of 15. A stolen laptop means an outcome of total loss of the value of this gift.

The last option you are considering as a graduation gift is a car. Your neighbors are buying a new car and they have offered to sell you their present car for $2500 – a great deal. Having a car during college is not a necessity for your daughter; friends would always give her a ride to where ever she needed to go. Having a car, however, means that she does not have to rely on others. She would use the car mostly for pleasure trips, for some extra-curricular activities, but rarely for schoolwork. While the car is now in good shape, it is getting old and it will not last much past graduation. Your worry, however, is that it will need frequent repairs during college, for this model has a higher than average frequency of repair history. The data on this year, make, and model indicates a 20% chance that it will need no repairs during your daughter's college years, in which case this gift will have its full value. There is a 50% chance that it will need one repair per semester, in which case the outcome is a gift that she can enjoy for most of her college years. Finally, the data indicates a 30% chance the car will need frequent repairs each semester, and the outcome here is a frustrating gift of little value.

You have been overwhelmed by this decision for some time now. Each option has its attractions and its problems. You want your hard-earned $2500 to go for the right gift. Which gift should get for your daughter?

7.4 Risk aversion, risk seeking, and risk neutrality

Before turning to decisions under the condition of ignorance, the topic of the next chapter, let's take a deeper look at the important concept of risk. People commonly use the word "risk" to mean the presence or threat of danger, or doing things that exposes a person to harm. In this

sense, the opposite of risk is the concept “safe”. Thus, we might describe a person’s life-style as risky, and try to get the person, if we care about him, to engage in more safe behaviors. Also, certain emotional states are commonly associated with taking risks. Many people believe that the appropriate emotion to have with respect to risk is fear: the more risk involved in doing something, the more danger there is that harm will come a person, and so we believe that it is natural for a person to be afraid to do it. On the other hand, it is widely recognized that some people react to situations that contain risk and danger with the emotions of thrill and excitement instead of fear. But however people react emotionally to risk, the point is that the word “risk” has a well-established common usage. In light of this fact, it will be useful to review and clarify the special meaning “risk” has in practical reasoning.

In the theory of rational choice, risk is not necessarily associated with harmful or dangerous options. Nor is it associated with any particular emotional state of the agent. An option might be dangerous for an agent to do, of course, but the danger might be a certainty. In other word, in decisions under certainty, there is by definition no risk, yet there could be the certainty of danger or the certainty of harm and even death to the agent. The suicidal terrorist who must choose to blow himself up on a bus or in a store is certain to die either way; there is no risk in this respect. Such a person might well experience fear in such a decision situation, yet there is no risk – in the practical reasoning sense of the term – if the goal is to die for a cause. Likewise, in decisions under risk, risk is present by definition and yet there need be nothing dangerous or harmful or frightening to the agent. Just think of the simple coin flipping decision situation described in Chapter 5.

We have been considering risk in a specialized way: that is, the role risk has within practical reasoning done by an ideally rational agent. A decision under risk does not at all mean that the agent is willing to lose the goal, or is engaging in risky actions, for the thrill of it. Rather, in this category of decision making **risk** refers to something about an agent’s beliefs when the agent is not certain. A rational choice under risk means that there has been: (i) a reasonably accurate

estimation of the agent's degree of confidence that an option will result in an outcome, (ii) this belief strength is formally represented within the decision framework, and (iii) its role is limited by the methods of practical reasoning to just one (important!) job; namely, as a weight to lower the utility value of an outcome, thereby making the outcome's connection to the goal weaker (or less weak in cases of outcome disutility) than it would be if it were certain to result. This lets the agent form a rational expectation of goal achievement for the outcomes of an option, without which the agent couldn't discover if the option is the rational choice or not. We can get a clear look at the role of risk within practical reasoning, as well as possible attitudes toward risk, by combining risk with certainty in a simple little decision problem.

Imagine that you are presented with two options x and y. Option x guarantees you $1000; you choose x and you get $1000 for certain. Option y involves a flip of a fair coin. If it lands heads, you get $2000. But if it lands tails, you get $0. Let's assume in this little decision problem that the utility of money equals the monetary amount of the money, and your goal is to get as much money as you can. Thus, the $0 outcome means that you achieve nothing of your goal, and the $2000 outcome is exactly double in strength over the $1000 outcome in achieving your goal. Your desire over these outcomes is clear: you want $2000 more than $1000 more than $0. However, option y combines the best with the worse outcome in a risky way that should leave you indifferent between options x and y. Here is the decision structure.

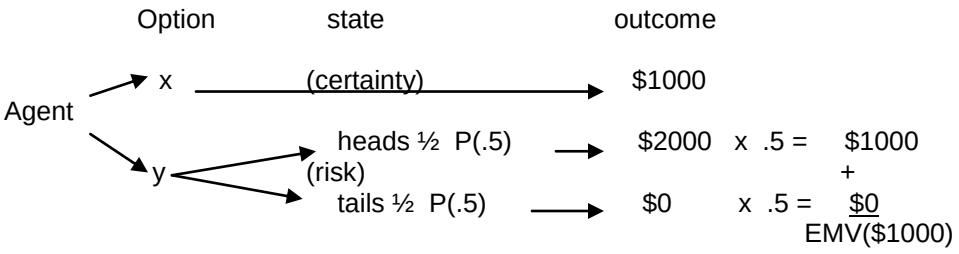

Option y involves risk, so we must discover its EMV. But option x has a certain outcome, so its utility is clear. In this decision problem, $0 outcome is your security level; you can't do worse, given your goal. $2000 is your hope limit; you can't do better, given your goal. The rational choice solution is clear: the ideally rational agent will be indifferent between x and y, for they have

equal outcomes: $1000. The only difference is that one option involves risk while the other is certain. But the risky option has had its outcome value adjusted (de-valued) in proportion to the risk factor, and this rationally adjusted outcome is equal to the certain outcome. In other words, imagine two agents one of which always chooses option x and the other always chooses option y. At the end of the day after choosing their respective options over and over in equal numbers, who come out with more money? Neither, they end up with equally achieved goals. Thus, practical reasoning should result in indifference, in theory!

You may not feel comfortable with this solution. You might protest that $1000 for sure seems better than an option that could result in $0. This is true. But note that there is an equally strong counter-protest: an option that could result in $2000 seems better than one that results in $1000 for sure. From the point of view of practical reasoning, the problem with these protests is that each focuses on only part of the whole decision picture. Comparing only the security level to $1000 makes option x look better than option y, and comparing only the hope limit to $1000 makes option y look better than option x. This is an example of relative value distortion. Once the security level outcome and the hope limit outcome are made to balance each other out, something that the model of practical reasoning for risky decisions requires us to do, it becomes clear that a rational agent should be indifferent between options x and y.

But what if a person is still not indifferent and prefers option x to y, or prefers y to x? This lack of indifference must be due to the person's attitude toward risk and certainty, for this is the only difference between options x and y. There are three possible attitudes to distinguish. (a) **Risk neutrality and certainty neutrality**. This means that an agent neither desires nor avoids risk, and similarly neither desires nor avoids certainty. This agent will be indifferent between options x and y, and this is the rational agent, by definition, in this decision problem. (b) **Risk averse and certainty seeking.** This means that an agent avoids risk, perhaps fearing it, and places special value on options whose outcomes are certain, perhaps finding comfort in certainty. This agent will not be indifferent between options x and y, but will prefer x to y. In this decision problem, this

agent is, by the standards of rational choice, irrational. (c) **Risk seeking and certainty averse.** This means that an agent avoids certainty, perhaps finding it unexciting, and places special value on risk, perhaps experiencing a thrill from risk. This agent will not be indifferent between options x and y, but will prefer y to x. In this decision problem, such an agent is, again by the standards of practical reasoning and rational choice, irrational.

Risk aversion and risk seeking vary with agents. A person might be strongly attracted to risk, or weakly attracted; likewise for risk aversion, it might be strong or weak. The above decision structure provides a way of estimating how much a person is risk averse or a risk seeker. We can vary the $1000 certain outcome, increasing or decreasing the amount. Let's use a blank box in which we can write amounts over or under $1000.

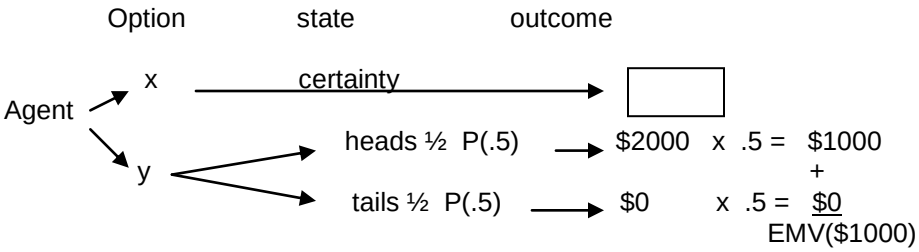

If $1000 is the amount in the box and an agent prefers x to y, then the agent is risk averse, for instead of being indifferent as the rule of rational choice prescribes, the agent avoids the risky option y and is attracted to the certainty of option x. So we decrease the amount in the box that the agent will receive with certainty until the agent reaches the point of indifference. For example, suppose agent Jones reaches indifference when the amount in the box is lowered to $800, but agent Smith reaches indifference when the amount in the box is lowered to $600. This shows Smith to be much more (roughly twice) risk averse than Jones, for the distance from $1000 to $600 is twice that of the distance from $1000 to $800. Smith requires a much greater motive in the form of reduced goal achievement to give up certainty seeking and risk aversion than Jones does.

If $1000 is the amount in the box and an agent prefers y to x, then the agent seeks risk and avoids certainty. Now we increase the amount in the box until the agent becomes indifferent. The amount of increase is a measure of how much the agent requires in order to balance the attractiveness of risk. For example, if agent Jones reaches the indifference point at $1200 and agent Smith reaches it at $1600, then agent Smith is roughly three times as risk seeking as Jones. Smith, in this example, requires a lot more incentive in the form of increased goal achievement to give up risk for certainty than Jones requires.

It would be impossible to evaluate how strongly an agent is a risk seeker or averse to risk if there were no standards to use as criteria. We have such standards; they come from the assumptions, the practical reasoning models, and the rational choice rules that are contained in the theory of rational choice. Given the goal of maximum money, the rational choice is (x indifferent to y) if the box contains $1000, (x in chosen over y) if the box contains any amount over $1000, and (y is chosen over x) if the box contains any amount under $1000. Agents who seek risk or who are risk averse fail all three of these rational choices, some more and some less. The theory of rational choice does not explain why some agents are risk seekers or averse to risk, nor does it tell us how to correct such a trait if an agent wished to. Also, it does not tell us the general cause of risk aversion (which seem factually to be far more prevalent than risk seeking and may be based deep in human evolution), as interesting as this topic is. But it does provide the conceptual framework for (i) defining the practical reasoning concept of risk, (ii) relating risk to rational choice, (iii) identifying agents who are risk neutral, risk averse, or risk seekers, and (iv) estimating the degree to which an agent is irrational with respect to risk. These appear to be interesting and important advances in the understanding of human rationality.

Sources and Suggested Readings:

The sources for this chapter are Keeney and Raiffa (1993) Chapters 5 and 6, primarily Mullen and Roth (2002) Chapter 7, to a lesser extent Raiffa (1997) Chapter 4, and Resnik (1987) Chapter 4. There is a large body of research, both formal and behavioral, on risk. Keeney and Raiffa devote much of Chapter 4 to relatively formal results. The behavioral field is dominated by the psychological studies of Kahneman and Tversky. See their "Prospect Theory: An Analysis of Decisions under Risk," Chapter 22 in Tversky (2004) or Chapter 6 in Moser (ed.) (1990), and "The Psychology of Preferences" in *Scientific American* 246 (1982); both are highly recommended to anyone with an interest in social science. Expected utility theory (at least in some of its versions), primarily the aspect of risk, is notoriously subject to several problems or "paradoxes": Allais's, Ellsberg's, The St. Petersburg's, and Newcomb's. That expected utility has attracted such scrutiny is a testament to its power and importance as a method of practical reasoning. See Resnik, Chapter 4 for a clear presentation of all four paradoxes. Hacking (2001) is especially recommended for the St. Petersburg's paradox (Chapter 8) and Allais's paradox (Chapter 9). Moser (ed.) contains the original articles of Ellsberg (Chapter 4) and Allais (Chapter 5), as well as Nozick's original presentation of Newcomb's problem (Chapter 9). By general agreement, Newcomb's problem is the most fascinating and philosophically interesting of the four. Nozick's presentation (in section 1) is very accessible and highly recommended. Straffin (1993) Chapter 6 offers a clear presentation. Shafir and Tversky's "Thinking through Uncertainty" (Chapter 28 in Tversky) also offers a clear description of this problem as well as a summary of the choices of people who have been exposed to it. A report on Newcomb's problem, based for example on Nozick's "Reflections on Newcomb's Problem: A Prediction and Free Will Dilemma," *Scientific American* 230 (1974), would make a wonderful term paper topic.