MAKING GOOD CHOICES: AN INTRODUCTION TO PRACTICAL REASONING

CHAPTER 4: FRAMING DECISIONS AND EVALUATING OPTIONS – MULTI-CRITERIA DECISIONS UNDER CERTAINTY

Let's recall some of the concepts needed to solve multi-criteria decisions. A complex goal is a goal that has at least 2 parts or objectives the agent must achieve in order to achieve the goal. To discover the best option to achieve a complex goal, the goal must first be analyzed into its objectives and the objectives ranked in their order of importance (that is: quantity of the goal achievement each equals). This ranking is then numerically represented on a normalized (sums to 1.0) interval scale. (See Chapter 2 for review in case you are not sure you understand the concepts in this paragraph.)

These ranked objectives will be used to form a set of criteria by which options can be evaluated. We'll do this in stages. First we begin with an example of a decision problem whose goal has only two objectives (deciding on a movie) and work through it without introducing any new concepts or theory. Try to get a rough idea, by way of this example, of what's going on in the case of a 2-objective complex goal. Don't worry at this point if there are ideas or steps that you don't understand. We next move to an example of a decision problem whose goal has three objectives (deciding on a job). As we work through this example we will be adding a new level of practical reasoning concepts and methods for arriving at a rational choice. This is the place to grasp the ideas and rationale of the steps. In the remainder of this chapter we will apply these new principles to a decision problem (deciding on an apartment) involving a goal with five objectives. By this combination of working with increasingly complex examples and being introduced to a new level of rational choice theory and methods of practical reasoning you will gain both a level of skill, and a basic understanding, of making good decisions under certainty when goals are complex.

4.1 Example #1: Deciding on a movie

Suppose one night you want to go out and see a movie. You have two objectives you'd like to achieve, and if you achieve both your goal will have been accomplished. First, you want to see a movie that several of your friends have seen, for you will enjoy discussing it with them, and second, it has to be a movie that's showing at an accessible theater, for you have been studying for a test and don't have much time, only about 45 minutes, to make the last showing. You search the internet to find out what's playing, and realize that you have just three options: (1) a murder mystery at a movie theater that's about 20 miles away by highway, (2) a romantic comedy at a theater 10 miles away by country roads, and (3) an action film at the nearest theater just 5 miles away with typically heavy traffic and many traffic lights. You recall that only one of your friends mentioned seeing the action film, you've heard a few of your friends talking about the romantic comedy, but you know that almost all of your friends have already seen the murder mystery. Which movie should you choose?

1) Forming criteria

To solve this decision problem, we will have to imagine that this agent (let it be you!) has some specific desires and beliefs. These will probably not be true to your actual desires and beliefs, but you'll certainly have others that you could readily draw upon in case you were really in such a decision situation as we have here imagined. So, for the purpose of this example, we'll suppose that you desire to see a movie that your friends have already seen slightly more than you desire to get to the theater in time. Such desires would yield something like the following goal analysis.

Goal: To see a movie

Objectives	Qualitative rank	Ordinal rank	Interval rank(1-10)	Normalized
(a) seen by friends	slightly more valued	2	8	.57
(b) accessible theater	slightly less valued	1	6	.43
			14	1.0

Now that you have made it clear to yourself how you rank these two objectives, you must next relate them to features your options possess that will allow you to compare your options to find the one that best achieves your goal. With respect to the objective “seen by friends” it’s pretty clear from our example that the more friends that have seen a given movie, the more that movie will gain you your goal. Let’s call this feature or attribute “number of friends”. With respect to the other objective, you’ll have to consider two features. It is clear from this example that (i) the closer in miles and (ii) the less time it takes to get to the movie theater, the more your goal will be gained. But which counts more, miles or driving time? You’ll have to determine this before a good decision can be made, and the method is the same as the one used to rank objectives. Here is how it might come out as you consider the two attributes for objective (b). (Remember, these actual numbers are not the point; what matters, rather, is that the numbers overall accurately represent the agent’s values and beliefs.)

Objective:

(a) seen by friends (worth .57)

Attribute	Qualitative	Ordinal	Interval (1-5)	Normalized
(i) number of friends	(only one attribute)	1	5	1.0
			5	1.0

Objective:

(b) accessible movie theater (worth .43)

Attributes	Qualitative rank	Ordinal rank	Interval rank (1 –10)	Normalized
(i) distance in miles	less important	1	4	.30
(ii) driving time	more important	2	9	.70
			13	1.0

Notice something important here that we will specifically be focusing on below when we work through the next decision problem: the attributes for each objective allow you actually to measure or observe something about each option. That is, for each option you can actually count or estimate how many of your friends have seen it and you can actually measure the distances and estimate driving times.

Attributes that are observable or countable in the options are crucial for practical reasoning to result in a good decision.

The next step is to use these attributes to form criteria by which to decide which option is the best choice. How much should each attribute count in coming to a decision? They clearly should not be given equal weight; some mean more to you than others. We adjust the value of the attributes by: multiplying their (normalized) values by the (normalized) value of the objective.

For the first objective (seen by friends), we have value .57 multiplied by its one attribute (number of friends who have seen each movie) having 1.0 value; this gives a product of .57. For the second objective (accessible theater), we have a value of .43 multiplied by its attribute (i) (distance in miles) worth .30; this gives us a result of .13. The same objective value of .43 multiplied by its attribute (ii) (driving time) worth .70 yields a value .30.

The criteria by which you'll choose which movie to see, then, are these three:

1) number of friends – worth .57
2) distance in miles to movie – worth .13
3) estimated driving time to movie – worth .30.

Notice that, in keeping with the base value of the goal set at 1.0, these three criteria sum to 1.0. It only remains now to use these criteria to evaluate the options; that is to discover which movie to choose.

2) Evaluating options

How would you rate each movie under each criterion? It should be pretty clear to you that some of your options do better than others given one criterion, but others do better given a different criterion. Let's rate the options on a (1—10) scale under each criterion. Go back and read the example carefully and you'll see that for criterion (1) almost all of your friends have seen the murder mystery, so it should place high on the scale, let's give it a 9. There is only one friend, as far as you know, who has seen the action movie, so it rates low on the scale; let's rate it a just 1. A few of your friends have seen the romantic comedy, so

it should rate somewhere in the lower-to-middle of the scale; let's say 4. How about the next criterion: distance in miles. On a scale of (1 – 10) how does each movie rate as to distance? Now do the same to each option with respect to the last criterion, estimated driving time. It is convenient to show these results on a grid. If you stick to the information contained in the example, you'll come out with ratings that are the same or very close to these.

Options	(i) number of friends (.57)	(ii) distance in miles (.13)	(iii) driving time (.30)
murder mystery	9 x .57 = 5.13	1 x .13 = .13	10 x .30 = 3
romantic comedy	4 x .57 = 2.28	5 x .13 = .65	7 x .30 = 2.1
action film	1 x .57 = .57	10 x .13 = 1.3	1 x .30 = .30

These ratings (the whole numbers on the left in each cell) do not take account of the important fact that these criteria are not of equal value for the agent. Each rating must be adjusted by the value of the criterion used to do the rating. Multiplying the option's rated value by the value of the criterion gives us this adjustment. This has already been done (number on the right side in each cell after the "=" sign).

Which option has maximum utility, that is, gains the agent more of the goal than the other options? Which option is the rational choice for you to make, given your goal? The answer is now easy to discover: simply add up, for each option, the weighted (adjusted) values and see which option has the biggest total. The murder mystery sums to 8.26, (5.13 + .13 + 3 = 8.26) making it the clear winner. The romantic comedy comes in second at 5.03, meaning it gains you significantly less of your goal than the murder mystery does. Finally, the action film sums to 2.17, a distant third, gaining you relatively very little of your goal.

4.1.1 Addressing a possible objection

Before advancing to the next example of a decision problem with three objectives (and seven attributes), and using it to introduce some of the general principles of rational choice that justify multi-criteria decisions, let's first address a typical reaction you might be having at this point. You might be thinking:

“No way! All this reasoning just to pick a movie! It is ridiculous to do so much work for something not very important.” And I agree with you! The example we just worked through was constructed not for its importance, but (I hope) for its ease of comprehension and familiarity. Not every decision in life demands the same level of care and effort we just went through. Many decisions can be made poorly and carelessly, and goals not achieved, and there are no nasty consequences to face. This pleasant fact of life was acknowledged right up front in Chapter 1. However, I’m sure there have been important decisions involving complex goals you’ve had to confront, and you can be sure that there will be many more of them in your future. Knowing some of the principles behind, and having some practice in, making good decisions are invaluable in such cases. The same ideas and methods of practical reasoning that we just covered for a relatively easy decision problem with a 2-objective goal and 3 options, and which we will be covering much more systematically in the remainder of this chapter, can be applied (or not applied) to decisions as you, the agent, sees fit and believes important enough to do or not.

Another point I would ask you to keep in mind, in case you had the “No way!” reaction I described, is that anyone who actually had a decision to make like the one in the example we just worked through, as relatively unimportant as it might be, will quite naturally mentally go through many of the step we covered. Of course, the person’s practical reasoning will happen very fast, it will not be slowed down as we have done here for teaching and learning purposes. And it will likely not be very systematic, certainly not as carefully worked through as we have done (again, for teaching and learning purposes). And it will probably be subject to several factors that make for poor practical reasoning that our careful steps have avoided. Nevertheless, you can be sure that his or her mind will swiftly consider certain features of the three movies in light of the desired goal and quickly come to a decision about which movie to see. That it will not be a very important decision is beside the point; namely, that a rapid mental effort at calculating and weighing various factors will indeed take place in the average person’s mind. So, if I this is how our minds typically operate when we make decisions, we are already doing a lot of mental work – much more than we ordinarily realize – in our everyday “unimportant” decisions. Then why not take the time to learn to do it right? It seems only smart, and not a ridiculous use of mental effort, to get into good intellectual habits of practical reasoning in making even our everyday kinds of decisions.

4.2 Example #2: deciding on a job

Let's return to the example in Chapter 3 of an agent who desires a good job and has three job offers. We want the goal to be complex for this example. So let's imagine that the agent desires a good job – the goal – and now believes that a good job for him is a job that has a good salary, good working conditions, and good promotion opportunities – these are the objectives. So, we have to distribute the agent's value for the goal among the goal's three parts, according to the agent's beliefs about how important a part of the goal each objective is. This requires a goal analysis. The agent mulls these three objectives over and let's supposes that this is what goes on in his mind.

"I'm really eager to gain a lot of experience from my first job, so promotions are very important to me; a lot of experience in my first job will make me more competitive landing a really nice job when there is an opening in a better corporation. I've just graduated college and have quite a few debts, so a decent salary is important, but right now in my life gaining job experience is more important. I'd like to have good working conditions, but I'm going to be flexible about this factor, for if I'm promoted quickly my working conditions are bound to change for the better."

Thoughts like these would result in the following goal analysis.

Goal: Get a good first job for myself so that I can eventually move into a better corporation.

Objectives	Qualitative rank	Ordinal rank	Interval (1 – 10)	Normalized
(a) salary	pretty big	2	7	.35
(b) working conditions	smallest	1	3	.15
(c) promotion	highest	3	10	.5
			20	1.0

1) Forming criteria

Now a set of criteria must be formed on the basis of these ranked objectives. These criteria will then be used to discover which option has maximum utility. Let's start with salary first. We know from the example we are using that the agent has three job-offers. How would the agent know which job is best under the category of salary? What would the agent have to count or look at and compare? Let's suppose that there are three things that matter for this agent under the category of salary: actual take-home pay, benefits, and transportation costs to-and-from work. How about working conditions? What would tell the agent that one job has better working conditions than another, or that two jobs had roughly equal working conditions? Remember, all of this is up to the agent to determine. No one else can do it, because no one else has the determining desires and beliefs that go into this decision problem. Suppose that the agent cares only about two things under the category of working conditions: a convenient parking space and a flexible lunch hour. Now, finally, let's turn to the important category of promotion opportunities. Here, let's suppose, the agent is concerned about only two issues: how quickly a promotion can be expected after starting a job, and how big a promotion can be expected when one happens.

4.2.1 Concepts and principles for multi-criteria decisions under certainty.

Before continuing to form a set of criteria, this is a good opportunity to introduce some concepts and some theory about what we have just done with the three objectives. Under salary we have listed three things: (a) take-home pay, (b) benefits, and (c) transportation costs. Under working conditions two things were listed: (a) convenient parking, (b) flexible lunchtime. Finally, under promotion two things were listed: (a) speed of promotion, (b) size of promotion. These seven things are called *attributes*. In the theory of rational choice, an **attribute** is a category or property that links a goal's objective to an option's outcome. For example, if a person takes a job, what is it about that job that links it to the objective of a good salary? One attribute provides part of the answer: it is the take-home pay. Namely, that job will have a specific amount of money as its take-home pay, and this specific amount is an instance of the attribute take-home pay. Another attribute provides another answer: it is the transportation costs. Namely, that job will have a specific amount of money as the cost of getting to-and-from the job, and this specific amount is an instance of the category or property "transportation costs". Again, if a person takes a job, what links that

job to the objective of good working conditions? An attribute provides part of the answer: convenient parking is one such attribute. That job will have a designated spot or area for employee parking that will possess some degree of convenience to the person using it, and that particular spot with its degree of convenience will be an instance of this attribute. An attribute, then, is a category or property that has a two-way relation: it specifies something (an objective), and something (in an outcome) is an instance of it. On the one hand, it is a sub-category relative to an objective; it comes under an objective. On the other hand, a specific feature of an option's outcome is selected out by an attribute, and this feature is an instance of the attribute.

The important concept of attribute just introduced should not be a source of confusion; try to make it clear. The key is to distinguish carefully these four related concepts: (a) goal, (b) objective, (c) attribute, (d) instance of the attribute in the outcome. Let's focus how these are related to each other. (1) An objective is related to a goal as part to whole. (2) An attribute is related to an objective as a sub-category to a category. (3) An outcome is related to an attribute as a specific instance is to a property. You are by now sufficiently familiar with the relation of part-to-whole from the examples of goal analysis covered in Chapter 2. So let's turn our attention to the relation between objective and attribute (the relation of category to sub-category), and the relation of attribute to outcome (the relation of property to instance of the property). Here is an analogy that might be helpful.

Suppose you wanted to buy a shirt and you wanted it to be colorful. You see a shirt that catches your eye, it is bright red, and you decide to buy it. As you go to pay for it, you notice that someone else waiting in the next cashier's line is buying the very same shirt. It's just like yours, especially the color. Now here we have three things to think about: (a) the category "colorful", (b) the category "bright red", and (c) the visible bright red color that you see in your new shirt and also see in the new shirt of the person who is buying a shirt like yours. How are these three things related? First, what is the relation between "colorful" and "bright red". "Bright red" is a sub-category of the broader category "colorful"; but it is not the only sub-category, for many other colors count as "colorful", bright blue for example. Thus, going from the concept "colorful" to the concept "bright red" you are going from a broad category to a more specific sub-category.

Buying something that satisfies the category “colorful” does not imply that you have satisfied the sub-category “bright red”, for you could have purchased a bright blue shirt. But notice what happens when you go in the other direction. Buying a bright red shirt automatically means that you satisfy the category “colorful” and thus have reached your goal of buying a colorful shirt. In fact, the only way to reach the goal of buying a colorful shirt is to buy a shirt whose color is a sub-category of “colorful”.

Now look at the relation between the concept “bright red” and the observable color that your new shirt and the other person’s new shirt have in common. Here we would not want to say that it is also a relation of broad category to sub-category. For one thing, the color that is in your shirt that you can see is not a sub-category, it’s a visible thing, a physical part of your shirt. So is the bright red color of the other person’s new shirt; you can see it, it is something concrete and specific that you can point to and can count – one time in your shirt, a second time in the shirt of the other person. So, how does “bright red” relate to what you see in these two shirts? The colors in the two shirts that you can see are two instances of the same property “bright red”. One instance of it is in your new shirt (it is probably more correct to say that it is in the dye absorbed by the fabric from which your new shirt is made). Another instance of it is in the new shirt of the other person waiting in the other check-out line. You might even notice, let’s suppose, that the bright red of your new shirt is ever-so-slightly deeper red than the bright red of the other person’s new shirt, the point being that the property “bright red” is instantiated in both shirts, but now you notice that it is instantiated in slightly different degrees.

As this analogy is intended to show, in order to reach the goal of buying a colorful shirt, you had to “travel” from an observable instance to a property, and from a sub-category to your final destination “colorful”. The mid-point “bright red” is a property relative to its instance (the observable color in your new shirt), but it is a specific sub-category relative to the general category that you wanted (colorful). Using this analogy as a model let’s return to the relation between outcome and objective. The mid-point is the attribute. Relative to an objective, an attribute is a specific sub-category (relative to salary, take-home pay is one of three sub-categories; relative to colorful, bright-red is a sub-category). This is a move from something general (the objective) to something specific (the attribute). But relative to outcome, an attribute is a

property an instance of which is observed in the outcome (relative to take-home pay, there will be an observable amount pay-check – a specific amount of money – as a consequence of taking any one of the three jobs; relative to bright-red there is the instance of bright-red you observe in a shirt).

When we return to the example of deciding on a job, you will see these relations clearly illustrated. But now let's continue with theory material. Attributes are key pivot points linking an agent's options to objectives and thus to the goal. There are several conditions that govern attributes. In explaining these conditions, we'll keep to our seven attributes above for use as examples. Under the objective salary we have three attributes: (a) take-home pay, (b) benefits, (c) transportation costs. Under working conditions two things were listed: (a) convenient parking, (b) flexible lunchtime. Finally, under promotion two things were listed: (a) speed of promotion, (b) size of promotion.

4.2.2 Six conditions for forming attributes

(1) Each attribute should be instantiated to some degree by each outcome of each of the agent's options. In other words, each outcome should possess to some degree an instance of each attribute. For example, whatever the job offers that make up the agent's options, each should have some degree of benefits, designated lunchtime, speed of promotion, etc., for each attribute.

This first condition has some flexibility. On the one hand, the agent can think of an attribute as absolutely required with no possibility of substitution or compensation. But on the other hand, the agent is free to allow an attribute that an outcome fails to possess in any degree to be compensated for by substituting something else. For example, suppose the agent discovers that one job that has been in the agent's menu of options up to now has no parking facilities at all. Strictly speaking, the agent should drop this job as an option because its outcome is completely missing an attribute. If the agent treats the attribute as non-compensatory, that is: strictly required, then the option gets dropped. But if the agent notices that this option has special strengths and attractions in some of the other attributes, especially within the same category of objective, she might allow these to compensate or substitute for the missing attribute. So, if

the job with no parking facilities whatsoever was especially strong in, say, lunchtime flexibility, and if this fact was allowed by the agent to compensate for the missing attribute, then this job could remain as an option.

(2) Each attribute must be directly or indirectly observable. The agent must be able to perceive the instances in the outcomes, or estimate their degree of presence, or count them. If an instance of an attribute is not directly observable (a specific take-home pay amount or a specific rate of promotion would be directly observable), then it must be indirectly observable by observing its effects. So, for example, suppose that under the objective "good working conditions" the agent listed as an attribute friendly co-workers. Clearly, the degree of a person's friendliness is not directly observable. How then could someone tell if co-workers were friendly or not? You could indirectly estimate this, say, by observing their behavior toward you – things like pleasant conversation, daily greetings, cards or even gifts on your birthday, help when your work gets too stressful, etc. Literally, these are the things that count, these are the things that you would count up, in estimating how friendly your co-workers were or if one co-worker were friendlier than another.

(3) Each attribute must make a difference. Any attribute that each outcome possesses to exactly the same degree will not help the agent decide which option is the rational choice. Such an attribute equally observed in every option's outcome can't do any decisional work, because it does not discriminate or differentiate between any two outcomes. If there were an attribute such that each outcome had it exactly to the same degree, then all of the agent's option in the menu would have a tie score as far as that attribute goes; it would not matter which option the agent decided on. For example, if under the objective of salary the take-home pay was exactly the same for each job in the agent's menu of options, this attribute could not be used as a deciding factor, for each job would be equally good in this respect. There would be no difference between options as far as take-home pay goes.

(4) Attributes must be logically independent. No attribute can imply or equal another, for if this happened, as we saw (Chapter 2) in the case where this condition also applied to the list of objectives,

there would be “double-dipping”, and some attribute would end up over-valued. Also, because they are logically independent, there is no danger that the attributes will form an inconsistent or logically incoherent set of properties.

(5) The set of attributes must be closed. As we shall shortly see, attributes figure centrally in forming the criteria by which a rational choice among the options is made. Thus, all the attributes affecting the decision must be listed. And once listed, they will form a closed set of criteria that will not let anything else sway the decision. So, if the agent has a bias, say, toward jobs in companies that offer stock options as part of the salary package, then this must get entered as an attribute under one of the objectives or it will not be able to affect the decision. If it can’t be entered under any existing objective, then a new objective (which amounts to a re-description of the goal) must be created. If the agent, say, has a desire for a job where some of the benefits included tickets to sports events, then this desire must somehow enter the set of attributes. If it does not get entered, then this desire will no longer be able to affect the agent’s decision about the best job. Once the set of attributes is final, they are designed to exclude the influence of any other factors on the decision. In a way, this is very liberating for the agent because the agent is now free to undergo a wide range of emotional experiences and cognitive states concerning the decision without having to worry that these might negatively affect the decision. There is no way for such emotions or thoughts to gain entry once the set of attributes is closed. This is an important part of the theory and practice of practical reasoning, namely, once a rational decision procedure is in place the decision process is protected from potentially irrational psychological forces (assuming the procedure is followed). The agent is then free to experience all kinds of feelings and moods, thoughts and desires, that otherwise could be sources of trouble by negatively affecting the decision if a system for making rational choices were not in place.

(6) The attributes must be value independent. This condition is important because attributes will be used to evaluate options and the resulting values will be added up. Thus, attributes must obey the principle of additivity. Here are two examples that are not additive. Suppose that you had money but nowhere to spend it. This money would have very little value for you. Suppose, on the other hand, that

you had plenty of places to spend money, but had no money at all to buy anything. These places would have very little value for you. So, you have one situation that has very little value for you and another situation that has very little value for you. These should add up to a situation that has very little value for you. But they don't. What happens when you add these two situations together? Now you have money and you have places to spend it. Suddenly your money is very valuable for you, and the places to spend it are very valuable for you. This new situation is not just the sum of the two previous situations, it is much more. You did not just add up the two previous situations in which the money and the places to shop each had little value for you, for if you had you would get a result of a situation of little total value. Instead, the places to spend your money made the value of your money dramatically go up, and likewise your money made the places to spend it greatly increase in their value for you. These two values affect each other; they are mutually dependent relative values. Here is another example of a non-additive situation. What would the value of a car with no fuel in it be? Very little, right? Maybe it is even worthless. And what would the value of a can of fuel be for you if you had no car to put the fuel into? Also, very little, and perhaps even worthless. If these two situations obeyed the principle of additivity, worthless plus worthless should add up to worthless. But that is not what happens. If you bring together these two situations into one new situation, you now have a car with no fuel plus a can of fuel that you can put into the car. Suddenly, you have a valuable situation, a car that you can drive.

We can generalize these examples and say that when one situation's value for an agent varies dependent upon the degree to which another situation's value varies, then these situations are not value independent (indeed: they have mutually dependent relative value). Another way to state this is this: the sum of the values of the parts does not equal the value of the whole. In the two examples, the value of the whole turned out to be much greater than the sum of the values of the parts (taken separately). What, then, is value independence? It is when the sum of the values of the parts equals the value of the whole. In other words, the values that two situations each have separately for an agent do not change, do not affect each other, when these situations are brought together.

By requiring attributes to be value independent, the operation of addition can be applied to their values. The degree to which an attribute promotes or fulfills an objective will not vary with the degree to which another attribute promotes or fulfills an objective. Here is condition (6) put in symbols: let v be the value of an attribute, and let x,y be any two attributes: $v(x)+v(y) = v(x+y)$. This is the **principle of value additivity**; it applies when the attributes are value independent, but not when they are value dependent.

Example #2 Deciding on a job (forming criteria – continued)

Keeping in mind the above conceptual material, we now return to the example of the agent deciding on a job. Recall that the goal was analyzed into three objectives, and these into seven attributes.

Goal: Get a good first job for myself so that I can eventually move into a better corporation.

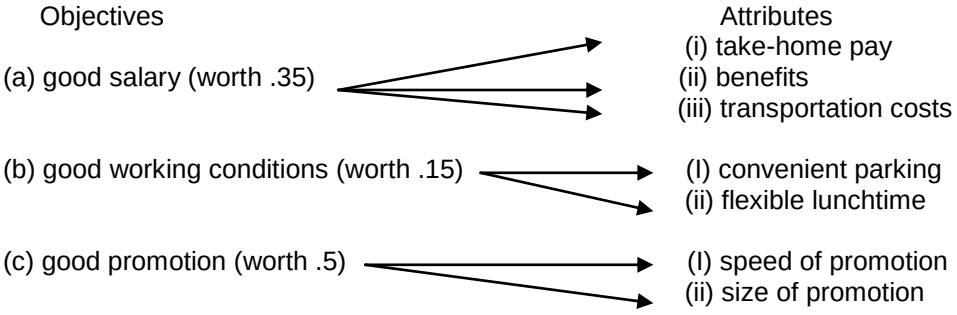

Note that each attribute satisfies all six of the above conditions. (Take a moment to verify this.) The next step in forming a set of criteria is to distribute the value of the objective onto the attributes falling under it. This is done in the same way that we arrived at the normalized value of the objectives. For each objective:

(a) first qualitatively rank its attributes in order of importance.

(b) then ordinally rank its attributes.

(c) then rank its attributes on an interval scale, preserving the ordinal rank but expanding or contracting the intervals to represent the degree of relative importance or value each attribute has for the objective it falls under. As in the case of analyzing objectives, this step is done according to the beliefs and desire of the agent.

(d) Finally, for those attributes under each objective, normalize to 1.0.

In order to carry out this step, we will again enter the mind of the agent and imagine that these are the thought she has.

"I care most about my actual take-home pay under salary, and next are the transportation costs because they directly affect my take-home pay. As far as benefits go, they are not very important at this stage of my career. A flexible lunchtime is very important to me, because I have health problems that require I eat at unusual times. But as to parking, it would be nice to have a spot close to the office, but not a problem if I have to walk a distance. Good promotion opportunities are very important, and I'll value the speed and size of the promotions I can expect equally."

With this set of beliefs and desires we have imagined the agent to have, the attributes will be analyzed along the following lines.

Attributes:	Qualitative	Ordinal	Interval (1 – 8)	Normalized
for objective (a)				
(i) take-home pay	highest	3	8	.47
(ii) benefits	least	1	3	.18
(iii) transportation costs	pretty high	2	6	.35
			17	1.0
for objective (b)				
(i) convenient parking	least	1	2	.22
(ii) flexible lunchtime	highest	2	7	.78
			9	1.0
for objective (c)				
(i) speed of promotion	high	1	8	.5
(ii) size of promotion	high	1	8	.5
			16	1.0

The last step in forming a set of criteria for multi-criteria decisions is to make the values of the attributes, which now altogether sum to 3.0, sum to the value of the goal which is 1.0. In other words, the value of the goal must now be distributed onto the attributes, using the value of each objective as a sort of control or a "gate," yet keeping to the relative proportions of values the attributes now have. The more important the objective, the wider the gate is opened to let value "flow through", and the less important the objective the more restrictive the gate so that less value "flows through". In effect, the value of each attribute is being proportionally discounted or reduced by the value of the objective it falls under until all attributes sum to 1.0, which is the standard value assigned to the goal. This is done simply by multiplying the value of each attribute by the value of the objective it falls under. This will give us the final attribute weight. These weighted attribute values will sum to 1.0, which is the total value of the goal.

Continuing with our example we arrive at the following.

(i) take-home pay (.35 x .47 = .17)

(ii) benefits (.35 x .18 = .06)

(iii) transportation costs (.35 x .35 = .12)

(iv) convenient parking (.15 x .22 = .03)

(v) flexible lunchtime (.15 x .78 = .12)

(vi) speed of promotion (.5 x .5 = .25)

(vii) size of promotion (.5 x .5 = .25)

For multi-criteria decisions, a **criterion** is defined as (a) an attribute, plus (b) its weighed value. Both factors are important. Without the attribute, the agent cannot tell what to look at in the outcomes to discover which option is the best means to achieve the goal. Without the weighted value, the agent cannot tell how much what she is looking at in the outcome should count in the decision.

2) Evaluating options

These criteria will now be used to evaluate the options. Recall in the example we are using, the example of an agent whose goal is to get a good job, that there were three job offers. Any of the three are hers for the taking; this is a decision under certainty. Here again are the jobs offered to the agent.

1) One job has a great salary, average working conditions, but poor opportunities for promotion.
2) The second job has wonderful working conditions, average opportunities for promotion, but a very poor salary.
3) The third job has excellent promotion opportunities, an average salary, but really bad working conditions.

Using the criteria as guides, the agent now gathers facts about each of these three options. The facts the agent must learn about are, for each job, the instances of each attribute. So, for example, what exactly is the specific take-home pay for each, what are the actual parking conveniences for each, what rate of promotion takes place in each, etc. So, the agent must make phone calls to the human resources or personnel office of each company, talk to employees of each company, ask key questions, etc. Let's suppose the agent has completed the fact-finding mission, and here is what she found out about each job.

1) The starting salary for job #1 is $45,000 with an assured year-end bonus of $2000. No dental benefits, and only 25% medical coverage as health benefits. This job will require expensive commuting by railroad that will cost $100 per week. The lunchtime is somewhat flexible, I can set my own times within certain limits. There are very convenient parking facilities, but I can't use them; the train means I can't use my car. Promotion in this company takes place in small steps and typically does not happen until an employee has worked there for well over a year.

2) Job #2 has a starting salary of $35,000. However, it offers full coverage for both dental and medical benefits. This job has low transportation costs, I get to use a company car, and the company pays the tolls. Fuel will cost approximately $20 per week. My lunchtime is completely flexible. There is a very convenient parking space for me right near my office. I can count on a nice promotion within the first year to a management position at least two levels above my starting position.

3) Job #3 has a starting salary of $40,000. There is a year-end bonus of about $1000. Dental and medical benefits are 50% coverage for each. I can drive to work; fuel and tolls will come to around $50 per week. My lunchtime is not flexible at all, and there are no parking facilities. I'll have to park my car several blocks away and it will be difficult to find a parking space. But the promotion opportunities look really good. I can expect a promotion within the first six months on the job, and because it's a small new company, I'll be promoted to middle management within a year.

Which job should the agent take? What is the rational choice, given this agent's goal? First, notice that no option gives the agent all that the agent desires. It is not a question of achieving the whole goal by taking one perfect job. Rarely does life allow us such options. It is, rather, a question of which option from the menu of options gains the agent as much of the goal as possible compared to the other options. This decision, like most decisions in life, is a matter of compromise and giving up in order to find the right balance of factors that most closely match the goal to be achieved. Without a helpful set of criteria, this decision would be very hard to make. Without criteria, the decision would be subject to all sorts of whims of the moment, last minute feelings, doubts and second-guessing, conflict, and many of the other forces that make for poor practical reasoning.

The second thing to notice is that taking any one of these options will gain the agent some of the goal. In other words, each option has a positive outcome as a sub-set of its total consequences (if one didn't, it would not have been allowed into the set of options to begin with, given the first option-narrowing principle). The problem is to discover which option's outcome has maximum utility; that is, which outcome has the strongest connection to the goal via the objectives. Of course, many other consequences will

result from taking one of these jobs. For example, the agent will meet new people. But such consequences are not part of the outcome. The outcome, recall, is determined by the goal. It is determined by the goal directly in the case of single criterion decisions. In the case of a complex goal, however, the outcome is determined by the criteria which were derived from the goal, first by goal analysis and then by objective analysis.

The first way the criteria were used was as a guide to find the right facts for each option, and describe these facts as the outcome of the option. Now we will use the criteria to discover which outcome has maximum utility. This is done by the following method.

1) Make a grid with the options listed on the left-most column and the criteria listed along the top-most row.
2) Select an interval scale that is large enough relative to the number of options.
3) Rank the specific instance of the outcome attribute on the interval scale under each criterion, going down. The idea is to compare the outcomes to each other only with respect to the instance singled-out by each of the criterion. Do this for each criterion, one by one, until the grid is filled in. The resulting values in each box are called **outcome utilities.**
4) For each column, adjust these outcome utilities according to the criterion's value. Do this by multiplying each outcome utility by the value of the criterion at the top of that column. These new values are called **weighted outcome utilities**. These are the outcome utilities discounted or adjusted by the value of the criteria.
5) For each option, add the weighted outcome utilities, going across each row. These sums are called **final outcome utilities**.

We now do this for our example of the job decision. Instead of writing the name of each criterion along the top row, we'll use its number and value (it is up to you to refer back to remind yourself what each criterion is). An interval scale of (1 –10) is sufficiently large for three options. Notice that it is the outcomes that are being evaluation, but because this is a decision under certainty, the options (the 3 jobs) can stand in place of the full description of the outcome. We will evaluate the three jobs on a scale

of (1 – 10) going down, for each criterion one at a time. Criterion (i) is take-home pay. So, we are comparing $45,000 plus a $2000 bonus for job #1, $35,000 for job #2, and $40,000 plus $1000 bonus for job #3. Try to use the full range of the scale so that you avoid tie scores. Notice that the salary for job #3 falls roughly midway between those of job #1 and job #2. The interval utility assignments should reflect this fact. Again, the exact numbers do not matter, it is the intervals that encode the important information. After salary, then do the same thing for benefits, then for transportation costs (the less costly the transportation, the higher the utility assignment). We will not walk through this step by step at this point. You should look carefully at the numbers in each box and reread the outcome description for each job to see exactly which facts determined the evaluations, paying particular attention to the distances between the three utilities in each column. Some of these evaluations may not be exactly the same as you would assign utility numbers, if this were a decision you had to make. That's fine. If you are the agent, you can judge the utilities as you see fit. In this case, you are not the agent and we have pictured an imaginary (perfectly rational) agent and have filled in values as we see this agent doing it. But note that if you and I both honestly went by the three job descriptions, our numbers would not differ by much, and the all-important intervals of our two utility assignments would end up very close in proportion to one another.

The top number in each box is the outcome utility and the bottom number in each box is the weighted outcome utility (the outcome utility adjusted by the worth of the criterion by multiplying: utility x worth). The last column on the right contains the sums of the weighted outcome utilities for each option: each option's final outcome utility.

Interval scalel (1-10)	(i) .17	(ii) .06	(iii) .12	(iv) .03	(v) .12	(vi) .25	(vii) .25	Final outcome utility
Job #1	10 1.7	2 .12	1 .12	1 .03	4 .48	2 .5	2 .5	3.45
Job #2	1 .17	10 .6	9 1.08	10 .3	10 1.2	5 1.25	5 1.25	5.85
Job#3	5 .85	5 .3	5 .6	1 .03	1 .12	10 2.5	10 2.5	6.9

Note that sometimes an outcome has a very high utility, but it does so under a criterion that has low value for the agent. On the other hand, an outcome can receive a low utility for a criterion that is highly valued by the agent. Without a system that allows the agent to reason clearly in bringing such factors together into an overall balance, making a multi-criteria decision would quickly break down in confusion and overwhelm an agent.

As you can see, job #3 has maximum final outcome utility. Thus, it is the rational choice for the agent, given the agent's goal. The rational choice rule for individual decisions under certainty is:

the agent should prefer the option whose outcome has maximum final utility.

If the agent in this example were to choose job #2, that would be a bad decision – an irrational choice – and to take job#1 would be an even bigger mistake. We should appreciate, however, in what sense jobs 1 and 2 are bad decisions for this agent. It is not as if someone were imposing a standard of practical rationality on the agent from the outside. And it is certainly not the case that we are telling the agent what to do, and then judging the agent irrational if she does not do what we say. The theory of rational choice respects the agent's own beliefs and desires. It sets up a system of practical reasoning within which the agent's goal, something that the agent freely desires and something that is respected as a given, becomes the basis for the agent to discover whether or not a certain course of action is the best means to that goal. "Rational" in the sense of practical rationality, then, refers to an ordering of the agent's beliefs and desires into a system of internal support and consistency. "Irrational" in the practical sense – the sense we would use if the agent had decided on, say, job #1 – means that the agent is fighting against or undermining herself, being her own worst enemy. It means that the agent is freely deciding things in a way that will *not* most achieve what the agent most wants to achieve.

4.3 Example #3: Deciding on an apartment

Here is another example of an individual decision under certainty, with a complex goal, and thus requiring multi-criteria evaluation. We will work through this example at a much faster clip than we did with example #2. To follow the steps, then, you should be prepared to refer back to the relevant sections of

explanation and theory material that has been covered in this chapter thus far. Here is a narrative describing the decision problem.

Your goal is to rent a nice apartment for yourself. You work, go to school, and socialize in Boston where most of your friends live. You don't make a whole lot of money, so the apartment needs to be affordable, reasonably clean, and it must be safe for both you and your car. A good location is also a factor, and so is noise, for you need your sleep and you are a relatively light sleeper. You have been looking at apartments and have narrowed your options down to three. Each of these three is large enough, has air conditioning, good heat and hot water, and you are equally pleased with the layout of each. The landlord of each apartment is waiting for your call; you can have any one you decide on. Here are your three options.

Apt. #1: In Boston, MA, no garage for your car so you must park on the street wherever you can find a spot. It is in a so-so neighborhood, in an old apartment building with thick walls, near public transportation making it easy to get around town, but also making for a fair amount of traffic noise. It is cleaned 2 times per month, halls and stairs only. The monthly rent is $1550, due on the first of the month, one month rent security returned when you leave. Police cruise by fairly often, especially during the night.

Apt. #2: In Nashua, NH, approximately 35 miles from Boston. This apartment is on a very safe and quiet side street, garage for your car. It is completely cleaned weekly very thoroughly. $1350 per month rent, no security required, rent due any time during the month. Police come by only when called, typically within a half hour.

Apt. #3: In Stoneham, MA, roughly 20 miles from Boston. No garage, but a driveway spot for your car near the building. The area is very safe. This apartment, however, is near a major highway and so is quite noisy. It is not cleaned at all, you have to do all the cleaning. $1400 rent per month, due on the first. $1000 security required, not returned. Pretty thin walls, so you can hear the neighbors. The police station is just down the street.

You must make up your mind which apartment to rent by the end of the day. Which one is the rational choice for you, given your goal?

Step #1. Goal identification and analysis. From the above narrative, here is a reasonable goal statement and set of objectives with their values.

Goal: I want to rent the best of these three apartments that fits my needs.

Objectives	Qualitative rank	Ordinal rank	Interval rank (1 – 10)	Normalized
1) Affordable	most important	5	10	.37
2) Clean	least important	1	1	.04
3) Quiet	not very important	2	3	.12
4) Location	quite important	3	5	.18
5) Safe	very important	4	8	.29

Step #2. Forming criteria. Here is a reasonable set of criteria derived from these five objectives. You should verify that each attribute satisfies all six conditions stated above in section 3.3

Attribute	Qualitative	Ordinal	Interval (1 – 10)	Norm.	Criteria value
Under affordable (.37):					
(i) rent amount	highest	3	9	.6	.37x.6=.22
(ii) rent due	middle	2	5	.3	.37x.3=.11
(iii) security	lowest	1	2	.1	.37x.1=.04
Under clean (.04):					
(iv) cleaning rate	least	1	4	.3	.04x.3=.02
(v) thoroughness	most	2	8	.7	.04x.7=.03

Attribute	Qualitative	Ordinal	Interval (1 – 10)	Norm.	Criteria value
Under quiet (.12)					
(vi) traffic noise	important	1	10	.5	.12x.5=.06
(vii) people noise	important	1	10	.5	.12x.5=.06
Under good location (.18)					
(viii) distance to work	highest	3	10	.5	.18x.5=.09
(ix) distance to school	big	2	8	.4	.18x.4=.07
(x) distance to friends	least	1	2	.1	.18x.1=.02
Under safe (.29)					
(xi) during day	not very	1	2	.1	.29x.1=.03
(xii) during night	very	2	8	.4	.29x.4=.11
(xiii) police availability	not very	1	2	.1	.29x.1=.03
(xiv) car safety	very	2	8	.4	.29x.4=.11

Attributes plus their values = criteria. Note that the interval, ordinal, and qualitative rankings match one another. Also, where there are equal ranks qualitatively assigned to attributes, the quantitative values appropriately preserve this information. Finally, note that both the number and values of the attributes are somewhat flexible. They depend on the agent's beliefs and desires, as interpreted from this example's narrative. But given the constraints contained in this narrative, your set of criteria would not vary greatly from the above set of criteria. Of course, if you were the agent making this decision, you would no doubt have your own set of beliefs and desires concerning an apartment. Your attributes and values, and thus your set of criteria, could look very different from the above, and the options you had to choose from would not be the above three imaginary apartments. This is to be expected. The theory of rational choice provides the structures and methods of practical reasoning for making rational decisions, but the agent must provide the facts concerning a specific decision problem. These facts supply the content of the decision problem; structures and methods give us its form.

Step #3: Option evaluation. Now we will use these fourteen criteria to evaluate the three apartments the agent has in the set of options. This will be done criterion by criterion. Because there are a relatively large number of criteria, we will rotate the utility grid so that the three options are along the top row and the criteria are listed along the left most column. The weighted utilities for each option, then, will be added down to find the final outcome utilities, instead of across as we did above in our first and second examples. You should re-read the description of each apartment above as a basis for understanding the outcome utilities that are assigned. These might not be your exact utility assignments, but yours and those provided would turn out to be pretty close if we were solving this decision problem separately. Again, the exact numbers do not matter, it is the comparative, proportional intervals that contain the relevant information.

Interval scale: (1 - 10) Criteria	Options: Apt#1	Apt#2	Apt#3
(i) rent amount (.22)	4 (.88)	10 (2.2)	5 (1.1)
(ii) rent due (.11)	1 (.11)	10 (1.1)	1 (.11)
(iii) security (.04)	2 (.08)	10 (.4)	3 (.12)
(iv) cleaning rate (.02)	5 (.1)	10 (.2)	1 (.02)
(v) thoroughness (.03)	5 (.15)	10 (.3)	1 (.03)
(vi) traffic noise (.06)	3 (.18)	10 (.6)	4 (.24)
(vii) people noise (.06)	10 (.6)	10 (.6)	1 (.06)
(viii) distance to work (.09)	10 (.9)	1 (.09)	5 (.45)
ix) distance to school (.07)	10 (.7)	1 (.07)	5 (.35)
(x) distance to friends (.02)	10 (.2)	1 (.02)	5 (.1)
(xi) during day (.03)	5 (.15)	10 (.3)	10 (.3)
(xii) during night (.11)	3 (.33)	10 (1.1)	10 (1.1)
(xiii) police availability (.03)	5 (.15)	1 (.03)	10 (.3)
(xiv) car safety (.11)	3 (.33)	10 (1.1)	7 (.77)
Final outcome utility:	4.86	8.11	5.05

Step #4: The rational choice rule for individual decisions under certainty is now applied. Given the agent's goal, it is rational for the agent to prefer the outcome having maximum utility. Thus, in this example the agent should strongly prefer apt. #2 over apt. #3, and weakly prefer apt. #3 over apt. #1.

4.4 Summary and review

Multi-criteria individual decisions under certainty are harder decision problems to solve than are single-criterion decisions. They are harder in one sense because we must apply the same methods of practical reasoning and principles of rational choice repeatedly within the same decision problem that are applied just once in the single-criterion case. They are harder in another sense because there are new steps we must do that have been added to the steps to solve single-criterion decisions. Here is a summary of the steps to solve **multi-criteria decision problems**.

1) Agent, clearly identify your goal and form a menu of options.
2) Form a set of criteria by which to evaluate your options.
    a) Analyze the goal into objectives
    b) Rank objectives on an interval scale and normalize to 1.0
    c) For each objective, form a set of attributes.
    d) Rank the attributes for each objective on an interval scale and normalize to 1.0.
    e) Multiply the objective values times the attribute values, for each objective.
    f) An attribute plus its final value (which for all attributes should sum to 1.0) = a criterion.
3) Evaluate the options.
    a) Form a grid with the criteria along the top row and the options along the left column.
    b) Evaluate the outcomes of each option on an interval scale, criterion by criterion,

one at a time (i.e. going down on the grid, column by column). These values are the outcome utilities.

c) In each column, multiply each outcome utility times the criterion's value. This value is the weighted outcome utility.

d) For each option, add the weighted outcome utilities in each option's row. This sum is the option's final outcome utility.

4) The rational choice for the agent is the option from the menu whose outcome has the greatest final outcome utility. This is the option the agent should prefer.

EXERCISE:

1) Here are three multi-criteria decision problems. Select one and using the information contained in the narrative, solve it. From the goal, derive at least three objectives, and for each objective, form at least two attributes, so that you are working with a minimum of six criteria. Depending on your objectives, some of the data in the descriptions of the options will form outcomes and other data in the description of options will be non-outcome (= irrelevant) consequences.

Decision problem #1: Deciding on a car.

Your old car has broken down beyond repair, and now you need a new one. Your transportation needs are primarily to and from your place of work, about 30 miles each way, mostly highway driving. Thus, your primary goal is a vehicle for good transportation, all seasons, for yourself, with the small possibility of using your car for carpooling a few times per month. Any payment above $200 per month will really stretch your budget. You live in New England where there are frequent ice- and snow-storms during winter. After looking at many used cars, you have narrowed your options to these three. Whichever one you decide on will be yours.

Car #1: 4 wheel drive, 17 mpg, large, heavy, safe station-wagon. $6500 ($150 per mo. payments at 5% interest rate). 35,000 miles, 5 years old. Well-maintained, but this model has an above average frequency of needing repairs. It has somewhat poor handling at highway speeds, but excellent handling in slippery conditions. It has a great stereo system (radio, tape, CD). This car is made in the USA. It comes with 1 year warranty parts and labor. Service is widely available.

Car #2: Front wheel drive, 30 mpg, mid-size passenger car. $8500 ($210 per mo. payments at 7% interest rate). 65,000 miles, 7 years old. It has a bit of superficial rust, but is mechanically in excellent condition. It has been well-maintained. This model has an excellent history of frequency of needing repairs. Excellent handling in all weather, but average handling at highway speeds. Radio only. This car is made in Japan. It comes with 10,000 warranty, parts only. Service is risky except at dealers.

Car#3: Rear wheel drive, 40 mpg, small, sporty-type compact car. $7500 ($150 per mo. payments at 0% interest rate). 55,000 miles, 6 years old. Excellent condition, body and engine. This model has an average history of frequency of needing repairs. This car is made in Europe. Repairs are costly and it frequently takes a long time to get parts. Radio and tape player. Excellent handling at highway speeds, but poor handling in slippery conditions. 3 years warranty parts and labor, but service only at place of purchase or warranty voided.

Which vehicle is the rational choice for you?

Decision problem #2: Deciding on real estate.

You are a small successful high-tech company that has outgrown its current facilities. Your goal is to move to better facilities, where you plan to stay for at least the next ten years. You will need larger facilities that will accommodate your continued rapid growth, at least double your existing space. The new facilities should not be too far from your existing location, the upper limit is 40 miles, for you want to keep all of your very valuable current employees and a move too far away would mean that many would resign. Because you are still a small company, cost is an issue, especially during times when the high-tech

industry experiences temporary down-turns. You have narrowed your options down to the following four pieces of available real estate. In all other relevant respects, these four equally satisfy you, but they differ in these relevant ways. Which one should you rent?

Facilities #1: A section of a very famous mid-town skyscraper 25 miles from your present location. The other sections of the building are rented by very successful high-tech and professional businesses. Easy access to this location by public transportation, but access is difficult by car. Three times as much space as your present facilities. The rent is $250,000 per month for a two-year lease, and after two years the rent increases by 5% per year. Many restaurants and clubs nearby for employees and to entertain clients. Entire building is professionally cleaned daily and is maintained beautifully.

Facilities #2: A former business headquarters in a country setting 40 miles from your present location. Large beautiful grounds, and a large parking lot near the building. Easy access to this location by way of a major highway, but no access by public transportation. Four times as much space as your current facilities. No cleaning services are included in the rent, but full buildings and grounds maintained are included. Rent is $125,000 per month, three-year lease, and after three years the rent increases by 3% per year.

Facilities #3: A section of a mini-mall just 10 miles from your present location. Somewhat difficult to get to due to local traffic during rush hours, many traffic lights, and school crossings nearby. Twice as much space as your current facilities. Rent includes full maintenance, and full cleaning once a week. The rent is $100,000 per month, three-year lease, after which rent increases by 5% per year.

Facilities #4: A large abandoned warehouse15 miles from your present location. This is an unusual situation, part of the city's urban development program. You must purchase the building from the city for $100,000, but rent the land from the city for $15,000 per month. You would have to renovate the building to make it suitable for your business needs. This will cost $10,000,000 that the city will loan to you. You must repay at the rate of $1,000,000 per year for ten years, with no interest. You will end up with four

times as much space as your current facilities. This location includes use of the city's parking lot, which is about a ten minutes walk from the building. Access to this location is average, neither easy nor difficult, by car or by public transportation.

Which piece of real estate is the rational choice for you?

Decision problem #3: Deciding on a night out.

The weekend is coming and you are looking forward to having fun. Your goal is to enjoy a Saturday night out. For you this means that you want to be with friends, and that you want to do something that you enjoy. But you don't want to stay out too late because you must be up very early the next day to take care of an important personal matter. 1AM is your limit. You are considering four possibilities: you have been invited to a party, there is a new movie you'd really like to see, there is a concert with one of your favorite band's, and there is an evening baseball game with your favorite team playing. All of these options look good, but for Saturday night you can only do one. Here is what you know, as you think these choices over.

The party: Several of your friends will be there, and three of your really close friends. You'll have plenty chance to interact with each other. It starts around 9, but will become much more fun around 11PM and will end around 2AM. You realize that you can't stay until it's over, that would be too late for you. If you go, you'll have to leave around 12:30AM, missing some of the fun. You will enjoy the music, the conversation, and the food, but what you will most enjoy about the party is meeting new people. If you miss the party, it won't be repeated. You'll bring a bottle of wine if you go, and this will cost around $15.

The movie: You have heard that it's excellent, and you are confident that you will enjoy it slightly more than meeting new people at the party. Two close friends want to see it with you. You enjoy their company very much, but at a movie there will not be much opportunity to interact with each other. It starts at 9 and you'll be home by 11:30PM. The movie will be playing for several more days, so if you don't go Saturday night, you'll be able to see it during the week but you'll have to go alone. The movie will cost under $10.

The concert: You know that it will draw a big crowd, and three good friends will be there, with plenty of opportunity to interact with each other. It starts at 9 and will end around midnight, and you'll be home by 1AM. You enjoy this band a lot, almost as much as you would enjoy meeting new people at the party. Tickets are $35. This particular concert will not be repeated.

The game: Your favorite team is playing and they are playing a team that is pretty evenly matched, so you are sure it will be an exciting game. You really enjoy baseball, and going to this game will be at least as much fun as the movie, if not a bit more. Trouble is that only one good friend will go to the game with you. It starts at 7:30 and will end around 10:30, and you'll be home by 11PM. This game is, of course, a one-time event, even though these two teams will play each other several times during the season. Tickets are $30.

Given your goal, what is the rational choice for you come Saturday night?

Sources and Suggested Readings:

This Chapter as well as Chapter 3 is primarily based on the presentations in Keeney (1992) Chapter 4, and to a lesser degree Chapters 5 and 6, and in Mullen and Roth (2002) Chapter 3. Along with the terminology alert in this section of Chapter 2, it should be pointed out that "criterion" in this chapter (and throughout this text) is "attribute" in the sources and suggested readings. What I have called "multi-criteria" decisions (both to highlight the normative/justifying orientation of the text and to build on the reader's intuitive understanding of applying criteria to evaluate things) is called "multi-attribute utility

theory (MAUT)" in these sources. The beginning sections of Chapters 2, 3, 4, 5, and 6 in Keeney and Raiffa (1993) are recommended for definitions and examples of multi-criteria decisions, but the latter sections of these Chapters tend to become more mathematical and technical. For a subscription online source that presents an accessible summary of some of the ideas in this chapter, see McGraw-Hill's Science and Technology Encyclopedia under "decision theory" if your library has access to this resource. For important and influential material on difficulties in framing and evaluating decisions, both multi-criteria decisions under certainty as well as under risk (coming in Chapters 5, 6, and 7) see both Tversky and Kahneman "Rational Choice and the Framing of Decisions" Chapter 24, and Shafir, Simonson, and Tversky "Reason-Based Choice" Chapter 39 in Tversky (2004).