

MAKING GOOD CHOICES: AN INTRODUCTION TO PRACTICAL REASONING

CHAPTER 3: FRAMING DECISIONS AND EVALUATING OPTIONS – SINGLE-CRITERION DECISIONS UNDER CERTAINTY

In Chapter 2 we covered two fundamental concepts of practical reasoning: the agent who reasons and the goal the agent desires to achieve. We presented and practiced methods of identifying and analyzing the goal. These methods result in an ordering of objectives according to an agent's subjective values. Now that the goal has been clearly identified and analyzed, the next steps in the process of practical reasoning is to form a set, a "menu," of options, to structure these options into a decision problem, and to evaluate them in such a way that a solution – a rational choice – can be discovered. These steps will be our topic in this chapter. Before turning to them, however, we must introduce several more concepts that are central to practical reasoning: option, state, outcome (the latter already introduced in Chapter 2), and the difference between a certain and a risky decision. We will then be in a position to practice the simplest kind of decision making: an individual decision under certainty, with a simple goal, and with options that are evaluated by a single criterion. This kind of decision problem will give us some of the basics of framing and solving a decision problem, and will provide a foundation upon which we can build in Chapter 4 to more complicated, multi-criteria decision problems.

3.1 Options and outcomes

Options are a set or menu of at least two alternative courses-of-action from which an agent chooses or decides what to do. (The term "choices" is commonly used interchangeably with "options" as in the common question: What are my choices?) Concerning options practical reasoning is, broadly, a process of selection and elimination – selecting what the agent will do from the agent's options, and at the same time eliminating those options that the agent will not do. The courses-of-action that make up an agent's options must, of course, be possible for the agent to do – possible in the practical sense, not just the logical or physical sense. What an agent can't do, given the agent's situation, can't be an option for the agent to choose to do in that situation.

Note that the terms “choices” and “actions” are being used in a special way here. We might speak of a person’s choices and actions (courses-of-action) and be referring to the actual choices the person has made and the real actions the person has done. This is what we normally mean when, for example, we say to a person, “Your choices show that you favor using public transportation over driving your own car.” or, “Your actions speak louder than your words.” But we might speak of a person’s choices and actions and be referring not to past decisions or actions but to the future, to what is available as abstract possibilities for the person to consider doing; that is, not to what a person has actually chosen or done, but rather to what a person should choose or should do. It is this second sense of abstract possibilities that is meant here by the terms “choices” and “actions.” If a choice is made and a course-of-action taken, then the agent’s behavior is said to “realize” or “actualize” one of the abstract possibilities that were available to the agent. Doing an action reveals to others by your behavior what your choice was, assuming – along with many other things – that you did not act contrary to your intentions. (As you might expect, realized choices and actualized courses-of-action are the preferred objects of study in the empirical parts of the theory of rational choice.) This idea of available abstract alternative possibilities is nicely captured by the term “options,” and it is the favored term in the theory of rational choice over the somewhat more ambiguous term “choices.” But for our purposes, either term works. There are several important points to make about the options that are available to an agent.

3.1.1 An option is an action or course-of-action an agent can choose to do. An action in this sense might take little time or effort to complete and be thought of as easy to do, or it might take a long time or a great deal of effort to complete and be thought of as something hard to do. Some options might take years to do and are very complex, involving a huge number of sub-actions, stages, phases, interruptions, and require much coordination, planning, and resources on the agent’s part, if they are to reach completion. Take a person who decides to go for a college degree; this will typically take about 4 years and require much effort to complete. A country that decides to go to war undertakes a course-of-action that could last many years, involving massive efforts and resources. A business that decides to develop a new product is often deciding on a long and costly course-of-action. On the other hand, some actions

are comparatively quick and easy to do, for example, a person deciding which friend to phone first, or whether to have coffee or tea with dessert, or a country deciding to fly its national flag at half-mast for a day. These are not necessarily easy decisions, but once decided on, such actions are relatively easy and quick to carry out compared to the case of a person going for a college degree or a country going to war.

The relative complexity of an option that would make it hard or take a long time for an agent to do (if chosen) should not be confused with the relative complexity of a goal having many parts. The goal could be simple but the possible courses-of-action that would achieve it might each be very hard, complex, and take a long time for the agent to do. Suppose, for example, an agent has a goal to apologize to a friend for betraying her. This goal has just one part and thus is achieved by accomplishing one objective, namely, saying “I’m sorry.” to the betrayed friend. Yet, this action (once decided on) could be very hard for the agent and take a long time to do, requiring a great deal of inner struggle, planning, overcoming rationalizations, procrastination, efforts to save face, false starts, waiting for the right moment, and so on. Conversely, a complex goal having objectives each of which is quickly and easily accomplished by an agent, typically add up to a complex goal that is achieved by an agent doing a total of easy, quick courses-of-action. For example, suppose an agent returns home from work with the goal of changing into comfortable clothing. Let’s say this requires changing 5 items of work clothing (think of an official uniform) for 5 items of everyday relax-wear. The actions that accomplish each of these five objectives are each relatively quick and easy to do, yet the goal is, by definition, complex. The complexity of the goal and the difficulty of the options that would achieve it, then, are different things. A complex goal is equally complex for any agent trying to achieve it, but the same courses-of-action that would achieve a goal could be very difficult for one agent to do and not as difficult for another agent to do.

3.1.2 The possible actions that make up an agent’s options are always possible **action-consequence pairs**. Actions have results; they change things in some way or other. Some of the changes that result from doing an action are intended consequences (these are outcomes), but other changes are

unintended consequences some of which the agent might know about and expect, while others might come as a surprise to the agent. For example, suppose an agent drives to a movie theater intending to pick up a friend who needs a ride home. One unintended consequence the agent knows about is that a certain amount of automobile pollution is put into the air. But if the agent gets into a car accident bringing the friend home from the theater, this is an unintended consequence the agent would not have expected. But whether intended or not, and whether expected or not, actions inevitably have consequences, and consequences are linked inevitably to actions. One automatic consequence of doing an action is that now that action will have been done, and sometimes just that result – namely, simply to be someone who has done such-and-such an action – was the only goal the agent wished to achieve. It is natural, then, to think of an agent's menu of options as containing possible action-consequence pairs, and to take an agent as deciding on such a *pair* when the agent decides what to do. To choose to do an action is equally to choose to bring about (at least some of) the consequences that result from doing the action. Likewise, to choose to bring about a consequence is equally to choose to do an action within the menu of options whose result is that consequence. An agent cannot choose one without thereby choosing the other, for the members of an action-consequence pair cannot be decided on in separation from each other. To make a rational choice, the rough idea is that, given the agent's goal, the agent should choose from the set of options that action whose consequence either is that goal or most helps to yield it; to do otherwise is to act unreasonably.

Recall from Chapter 2 that **outcome** is the sub-set or the part of the consequences of an option that actually relates to the goal. Of all the possible consequences that could result from a doing a course-of-action, some might gain the agent the whole goal or at least part of it, promote it in some fashion, or brings the agent closer to achieving it (positive outcome). On the other hand, they might threaten the goal, cause the agent to lose part or all of it, or distance the agent from achieving it (negative outcome). The term "**utility**" was briefly introduced in Chapter 1 to mean a measure or degree of good that results from an agent's action that is done to achieve something the agent values. Let's now define (by stipulation) utility more precisely. **Utility** or **positive utility** will be used throughout this text to mean a measure of the degree to which to a positive outcome gains an agent the goal. **Disutility** or **negative**

utility will be used to mean a measure of the degree to which a negative outcome loses an agent the goal or distances an agent from the goal. We will be working with this concept of utility shortly.

Aside from the outcome (also called “payoff” in some decision situations) of an action, there will typically be other consequences that have no enhancing or threatening connection to the agent’s goal. These can be divided into two groups depending on their affect: (i) those that are costly to the agent, and (ii) those that are irrelevant. Deciding on a course-of-action will result in costs for the agent. Think of **costs** as things that the agent values but must give up if an action is to reach an outcome; time, money, effort, the new condition of tools or instruments (if they are to be used), opportunities to do other things, sometimes health, pride, reputation, moral principles, friendships, perhaps even the agent’s own life are examples. Costs are those things that have to be sacrificed in order for a decision to achieve a goal. Some costs can be recovered, but an important category of costs that can’t be recovered are called **sunk costs**, time being perhaps the best example. In order to make a good decision, an agent must have a way to estimate decision costs and factor them into the reasoning that leads to the decision. If an option’s costs outweigh the utility of the outcome or the value of the goal to be achieved, a rational agent will not choose that option. **Irrelevant consequences**, as the name implies, are any results from a course-of-action that are neither outcome (positive or negative) nor costs. For example, if your goal is an evening entertainment and you choose to go to the movies instead of watch TV, a cost is the price you must pay for the movie ticket but an irrelevant consequence would typically be that the movie theater makes a slight profit from your decision.

It might be worth repeating here an important idea from Chapter 2. If the total consequences of an action are to divide up between outcome and non-outcome consequences, how does this division happen? It is the goal that determines what part of the total consequences counts as outcome. Remember the example: an agent decides to drink a glass of juice, and as a result let’s say there are three consequences: (a) a glass is now dirty, (b) the agent has had an intake of liquid, and (c) the container of juice is now empty. What’s the outcome? There is no way to tell without identifying the goal the agent desired to achieve. If the goal were to quench a thirst, then (b) is the outcome, and (a) and (c) are non-

outcome consequences (leaving aside the question whether either is a cost or irrelevant), for only (b) brings the agent closer to or achieves that goal. But if the goal were to consume the juice, say before it turned sour, then only (c) is the outcome. Because the goal determines, in the total consequence of an option, the division between those consequences that are outcome and those that are not, the agent must make it very clear exactly what the goal of the decision is. We covered in Chapter 2 the practical reasoning techniques for doing this.

Clearly, the outcome is the most important part of an option, but it is not the measure by which a decision is evaluated. A good choice is *not* defined as a choice having a good (positive) outcome. (This would be a version of the false idea that the end somehow justifies the means.) Likewise, a bad decision is *not* defined as a decision having a negative outcome. This must be kept in mind because it is common for people to think of a good decision as one having a positive outcome, and a bad decision as one having a negative outcome. Instead, in the theory of rational choice a **good decision** is defined as a rational decision. "Rational" is a property of the reasoning that leads to the decision, it is not a status conferred on the decision making process retroactively from a positive outcome. It is perfectly possible, given this definition, for a bad decision to have a great outcome. Take the case of a person in serious need of money who scrapes together every last bit of cash he has and, knowing nothing about horse racing, bets on the horse least likely to win the race. The decision to bet was nothing more than an act of desperation, and the horse bet on was selected, let's suppose, by pure whim. This is very poor practical reasoning, very bad decision making (you will see why later, but here it should be intuitively clear to you that it is), and it does not matter if the horse wins and the agents enjoys a great outcome. This would be good luck, perhaps, but certainly not a good decision, not the model of how one should make decisions. So, a bad (that is: an irrational) decision can, by pure accident or by nothing but a stroke of luck, have a good outcome. The converse is also possible; good practical reasoning, good decision making (that is: a rational choice) can have a bad outcome. Well known examples of medical doctors following the most reasonable decision procedures but nevertheless having their patients die as a result are cases in point.

3.1.3 An agent can (always?) decide to do nothing. This means that one of an agent's options is the option not to act. This is to make a decision not to do any of the (other) available courses-of-action. But this itself counts as making a decision to do one of the agent's options. In other words, the decision not to act is thought of as a decision to do *something* – it is a legitimate course-of-action (maybe called by the agent “holding-off” or “buying time” or “procrastinating” or taking a “wait-and-see” approach, and often described as “keeping my/our options open”). Of course, this action too has its consequences, and in many decision situations the choice to do nothing might be the right decision at the time for an agent to make, if the goal is eventually to be achieved. It is important, then, that we distinguish the case of indecision from the decision not to act. If an agent does not make a decision, can't bring himself to choose, wavers and fails to come to a decision about what to do from among the available options, then there is indecision on the agent's part. Indecision is not (typically) one of the agent's options, and being indecisive is not (typically) a good thing for an agent to be. Not being able to decide when we should be making a decision is something agents must work to overcome. But indecision is a far cry from the agent who *decides* to do nothing. Not deciding to act is not the same thing as deciding not to act. Indecision is one thing but taking a wait-and-see stance is something very different.

3.1.4 An agent can (always?) **opt-out**. One option an agent typically has is to leave the decision situation altogether. This is to give up the goal, at least for the present decision situation, and perhaps for good. But opting-out is not the same decision as deciding to do nothing. Let's look at some examples that illustrate this difference. Take a college student whose goal is a college degree. One semester, due to many personal problems, the student decides that her goal is best achieved by not taking any classes. So, she “holds-off” enrolling in classes that term, yet maintains her goal of a college degree. Now take another college student who gives up his goal of a college degree, and decides to drop out of college completely. Only the second student has opted-out of the decision situation, the first student has decided to do nothing for a semester because “holding-off” for a semester is the best way, in her judgment, to achieve her goal of a college degree. For another example of the difference between deciding to do nothing and opting-out, take a couple going through marriage difficulties. Their goal had been to remain together. Opting-out would be deciding on divorce; this is to give up on the goal of

staying together. But if the couple decided to do nothing about the marriage difficulties, let's say that for the next six months they decide to take a "wait-and-see" approach, to "keep their options open" concerning their marital problems, they are not opting-out of the marriage but are deciding on an option that affirms the goal of remaining married. These examples should make it clear that among an agent's options, the option to do nothing is not the same as the option to opt-out.

3.2 The State

State means the state-of-the-world in the decision situation that is required in order for an action to produce an outcome. Actions typically do not produce a given consequence unless the world "cooperates." Certain relevant conditions must be in place, certain parts of the agent's environment must be in the right arrangement or be happening the right way, or else the action performed by an agent could end up having nothing but unintended consequences. These relevant conditions might be physical, social, institutional, or interpersonal. The rough idea is that the right choice plus the right state-of-the-world yields the desired outcome: $(C + S \rightarrow O)$. If the right state is not in place, then the action an agent might decide to do could end up going very wrong. For example, suppose an agent decides to drive to the store with the goal of purchasing an item of importance. The state would include such things as a reliable car (no major breakdown on route), the availability of the item (the store did not run out), and no major traffic tie-ups. These are relevant conditions in order for the action to produce an outcome that gains the agent the goal. It may be true that there are other people in the store, trees and houses along the roads taken, and clouds in the sky. However, such things are not part of the state for their presence or absence does not play any role in the connection between the action decided on and the outcome that results. The state, then, includes all and only those conditions that are relevant; that is, that can affect the outcome. Let's look at this idea of "relevant conditions" in more detail.

The situation in which a decision is carried out is often a very complex mix of physical, social, and interpersonal factors. Which factors comprise the state and which are irrelevant to the decision? Most, no doubt, will be irrelevant. But the line between what counts and what doesn't is typically very hard to draw. Here is an example involving your own home. Suppose you smell gas fumes in your house during

a storm that has caused a power outage. You are worried that if the power suddenly returns later and there has been a buildup of gas, it could cause a spark that might explode the gas or cause a fire. You must investigate and try to find and stop the gas leak. But there is no electricity, and it's nighttime. Without some kind of light there is very little chance of achieving your goal. What to do? There is not a strong smell of gas right now; one possibility is to light a candle, another is to try to grope in the dark and hope you find the leak. What is the state that is relevant to an outcome that will achieve your goal of stopping the gas leak? Clearly, the amount of gas fumes in the house is part of the state. But is the size of the rooms relevant? Perhaps, if the sizes are either a very large or a very small areas that would diffuse or concentrate fumes; but perhaps not, if the rooms have small size differences. How about whether or not there are windows? Is this relevant? Perhaps, if they will allow ventilation; but if they won't, then perhaps not. What about how full or empty of furniture the rooms are? Again, this might or might not be part of the state of this decision between using a candle or groping in the dark. The point of this example is that even a relatively familiar place such as a person's house can be very complex, and shows how hard it is to be clear about just what is and what is not part of the state in which a decision is made. Yet this is just what the agent must try to do, for the general formula surely holds: choice plus state yields outcome.


Because the state links an action the agent might decide to do to an outcome the agent intends to bring about, states should be thought of as part of an agent's options for each option the agent has. So, instead of defining an agent's choices or option as alternative action-consequence *pairs*, as we did above, we will now improved the definition of an option to include the state. An agent's **options** are at least two available alternative action-state-consequence *trios* that make up a set or menu of possible actions from which an agent must choose which to do in order to achieve a goal.

3.3 Certainty


In making a decision, what might an agent believe about the state? One possibility is that the agent is sure that the required state exists. This is called making a decision under certainty. Another possibility is that the agent is uncertain, has a degree of doubt, about what state will exist when an option is acted on.

This is decision making under risk, and will be our topic in latter chapters. **Decision under certainty** means that for each of the agent's options the agent knows with practical certainty that doing a given action will result in a given outcome. In other words, the agent is as sure as can reasonably be expected that the state required for an option to have its outcome is in place.

The basic structure for decisions under certainty is:



If we expand this basic structure to, say, four options we have a diagram like this:



Certainty, in this structure, is a maximum strength of belief an agent has about the state-of-the-world required for a possible course-of-action to have its intended outcome. In the field of practical reasoning, this certainty is practical not logical. It is always logically possible for something to go wrong between the time an agent starts to act, having decided on an option under certainty, and the time the outcome happens, even if this time span is very brief. The universe might suddenly end, or aliens might interrupt the process, or God might perform a miracle – such things are always logically possible. Decision under certainty, however, means that the agent knows that there is no practical possibility that the required state is not in place or will change for the worse. In other words, the agent knows that the probability of an action not reaching its intended outcome is so small as to be negligible, and for all practical intents and purposes can be considered zero.

How might an agent acquire such practical certainty? One way an agent would be in a decision situation under certainty is if the agent has **full control** over the required state. For example, once an agent decides to take a trip by car rather than by plane or train, having a reliable car becomes part of the required state. The agent can have the car checked over, or perhaps even buy or rent a new highly

reliable model for the trip. If, to take another example, a course-of-action requires the cooperation of other people in order for the desired outcome to happen, the agent can work to gain and maintain the necessary cooperation. In many decision situations, the agent can prepare the way by exerting control over the state, keeping conditions in place for the outcome to happen.

Another way an agent might be completely confident about the state is **by default**. This means that the agent has no good reason to doubt that the state is fixed and exists as needed. In this sense, the agent relies on the fact that certain things have always existed, or events have always occurred, and the agent knows nothing that disturbs the way the world has always been. An agent going on a hike is practically certain that the mountains are still there; the agent has no reason to believe that, say, a large earthquake has destroyed the mountains. An agent going to the beach is completely confident that the ocean and beach are still there; the agent has no good reason to think that, say, a massive tsunami has altered the shore line and destroyed the beaches. Someone going shopping is maximally sure that the stores are still there; there is nothing the agent knows that would cause her to doubt the stores still exist, such as news of a fire or explosion destroying them, or word of a bankruptcy closing. In all such cases, it is not that the agent has actual up to the minute evidence about these things. Rather, that's the way the world has always been, and the agent has no reason to doubt that the world is the same in the relevant ways. There is no reason, to the best of the agent's knowledge, to believe that the required state is not in place.

The third way the agent might gain a maximum degree of certainty concerning the state is **by direct evidence**. This means that the agent actually possesses good evidence that the required state exists. Suppose that you need your car that your friend has borrowed. You are not sure if it has been returned. You take a look and there it is in the driveway. Now you are certain (in the practical sense of the term) that your car is there, for you have seen it. You are not certain by control or by default but by direct evidence, in this case by direct observation. Let's say that a hiker hears that a massive earthquake has destroyed the mountains where she loves to hike. But the reports do not sound reliable. She calls the authorities who assure her that the earthquake left her favorite mountains intact. Now, by direct

evidence, in this case a source she has every good reason to trust, she is completely confident that going on a hike to her favorite mountains will yield the intended outcome.

As you can see, practical certainty in decisions under certainty is not the same as absolute certainty. (Is anyone absolutely certain about anything?) It is a bounded kind of certainty. It is certainty relative to what the agent knows or trusts with good reason to be the state, either by having control or by default or by direct evidence. Beyond this, there is no practical certainty.

With the concepts of option, state, and certainty now in place, we are ready to discuss the practical reasoning steps of forming a menu of options, framing the decision problem, and discovering the rational choice for the case of decisions under certainty when the goal is simple.

3.4 Forming options: two option-narrowing principles.

There are many actions a person could do at any given moment. The possible actions you could do at this very moment, for example, forms a limited but nevertheless large set of possibilities. However, if a person wanted to act in a way that achieved a goal, then from all the actions a person could do there would typically be a narrow sub-set of possible actions selected, namely, only those that the agent believed were related in a special way to the goal. These possible actions would be the agent's options. Let's see this by way of an example.

Suppose an agent has a goal to get a good job for himself. This goal automatically set limits on what can count as an option. No possible action that has nothing to do with gaining employment can count as an option for the agent, given this goal. For example, staying in bed can't be an option. Why not? The reason is that the total consequences resulting from the activity of staying in bed do not contain any expected consequences that bring the agent closer to the goal. In other words, this action has no intended outcome, given the goal. To be an option, an action that is possible for the agent to do must have among its consequences an intended outcome (in our specialized sense of "outcome"). But what if the agent believes that staying in bed will have an outcome, given the goal of getting a good job?

Well, such an agent would probably be very wrong to believe this, and will discover the error soon enough; but so long as this is sincerely believed, then staying in bed could count as an option. We'll put this general idea as a principle of forming the menu of options:

Option-narrowing principle 1: To be an option, the agent must believe that an action possible for the agent to do has an outcome, given the goal.

This principle makes it important for the agent carefully to consider the expected consequences of possible actions to verify that some of these consequences can qualify as outcome. It is only those expected consequences that count as outcome that qualifies the action having those consequences to enter the menu of options, and it is the goal (as described in the goal statement) that determines which of the expected consequences count as intended outcome.

There is another important principle that serves to narrow an agent's options. An agent might have in his or her belief system norms and rules that would disqualify actions as option, even though the actions have outcomes. Indeed, it would be very surprising if an agent didn't have principles or a code of conduct that make some actions impermissible. For example, suppose that telling a lie would gain an agent a good job. Upon reflection, however, the agent clearly sees that lying in such cases is wrong, and so does not permit herself to lie in order to land a good job. Lying is now disqualified as an option in this agent's menu. Most agents have such rules and try to live by them. They could be a moral code, a legal system, religious considerations, social norms, a professional code of ethics connected to an agent's career, or even psychological inhibitions of a personal nature, for example actions an agent finds too embarrassing to do. An agent's moral principles, respect (or fear) of the law, religious commandments, etc., serve to rule out certain actions as wrong, or illegal, or sinful, or ill-mannered, or personally immodest. These actions are dropped from, or not let into, the menu of options. They disqualify actions as options, even though the actions have outcomes. Let's form this as a principle:

Option-narrowing principle 2: To be an option, an action that is possible for the agent to do must pass the agent's disqualifying rules.

It is important to note that the disqualifying rules should be brought into play at the earliest possible stage of forming a menu of options. They help narrow the agent's choices before a decision is made. One benefit to this is that a decision is often easier to make the fewer options the agent must evaluate. The more important reason, however, is that if the disqualifying rules are brought into play as an "afterthought", after the options have been evaluated, the agent could end up in an uncomfortable situation. The agent might find that the best decision has to be eliminated because it turns out not to pass the disqualifying rules. The agent could experience the temptation to override moral, legal, or religious beliefs in favor of the best option, or else to give up the best option for achieving the goal in favor of moral, legal, or religious principles. In such a case the agent has a dilemma that could and should have been avoided. Practical reasoning need not clash with one's moral, legal, religious, or social values, or one's personal code of behavior when it comes to possible options, if the agent brings disqualifying rules into play early in forming the set of options.

Let's imagine that an agent has narrowed his options down by applying the two option-narrowing principles to the large set of all actions possible for the agent to do in a given decision situation. This yields the agent's menu of options relative to the goal to be achieved. We now turn to the reasoning that will discover which option makes the best choice.

3.5 Individual decisions under certainty, simple goal, and single criterion.

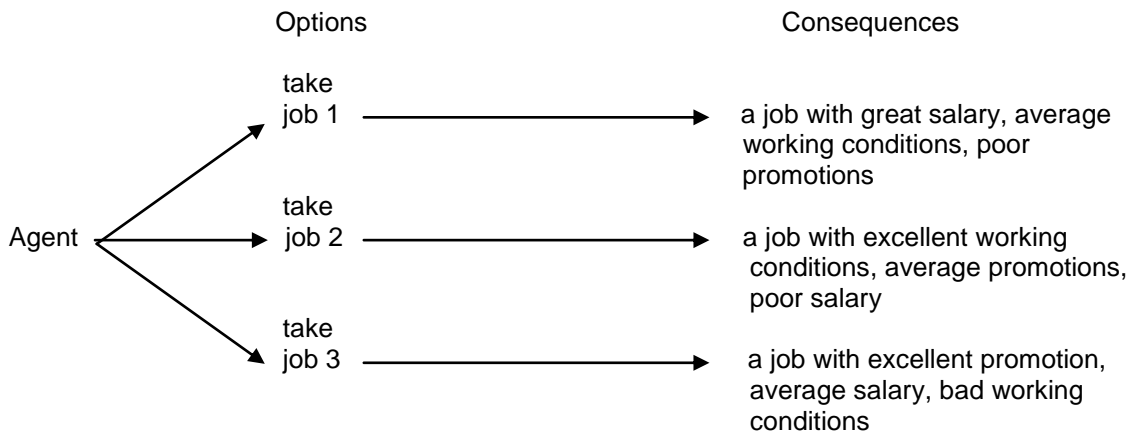
Recall that decisions under certainty means that the agent is completely confident about the state of the world that must exist if the action the agent decides to do is to have its intended outcome. The agent has no worries at all about the state. Also recall that a simple goal is a goal that is completely achieved by achieving one objective; the description of that objective is the description of that goal, they are one-and-the-same. Here is an example of the most elementary kind of decision problem, the easiest to solve.

Imagine that the agent whose goal is to get a good job has narrowed the options down to three job-offers. Any of the three are his for the taking, and the agent is certain (in the practical sense) of this. Here 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.

Let's structure these three options.

Goal: To get a good job for myself.



3.5.1 Option evaluation.


What is the rational choice? Which is the best job for achieving the agent's goal? This cannot be decided until the goal is identified and analyzed. Identifying the goal, you will recall, must address 3 things: stakeholders, outcome, and kind of value. Analyzing the goal, in this case a simple goal, means clearly describing its single objective (whose value is 1.0). Describing this objective will do several things.

- 1) It makes it clear to the agent that by achieving it, the agent will achieve the goal (by definition).

- 2) It will divide the consequences of each option into outcome and non-outcome consequences. One objective will select just one property in the set of consequences, and this one property will be the option's outcome.
- 3) It will provide the criterion by which to evaluate which outcome has maximum utility; that is, the criterion by which the agent can compare outcomes to discover which option has the outcome that will bring the agent the biggest part (or all!) of the goal.

The rational choice rule for decisions under certainty is: choose the option whose outcome has maximum utility. Recall that utility is defined as the strength with which an outcome achieves a goal. So, now it is purely up to the agent as to what defines a “good” job. Let’s suppose that for this agent the single objective that equals the goal is: promotion opportunities that provide job experience. By this objective, the promotion opportunities of the 3 options become outcomes; salary and working conditions are other consequences that are now not able to influence the decision. The original diagram is changed into:

Goal: To get a job offering good promotion opportunities in order for me to gain work experience.



This one objective – promotion opportunities – becomes the criterion by which to evaluate the options. Each outcome can now be ranked as to utility. The agent can clearly observe that job 3 has highest utility, job 2 comes next, and job 1 has lowest utility. Thus, the agent should choose job 3 over job 2, and job 2 over job 1. Job 3 is the rational choice and is the solution to this decision problem. Any other choice is irrational, for given the goal and the available options, another choice would lose the agent part, if not all, of the goal. Of course, had the agent defined “good job” by the single objective of, say,

salary, then only salaries count as outcome for the three options and job 1 is the rational choice, job 3 having least utility relative to this very different objective.

Single criterion individual decisions under certainty are the easiest decisions problems to solve. Here is a summary of how these decision problems are solved.

- (a) The goal's single objective supplies the agent with a **single criterion**.
- (b) This criterion selects from the consequences of each option a single property or feature that becomes the option's outcome.
- (c) Option evaluation – by comparing the outcomes with each other according to the criterion, the agent discovers the outcome that most satisfies the criterion; it will be the outcome having maximum utility.
- (d) The option with this outcome is the solution to the decision problem – it is the rational choice.
- (e) The other options represent various degrees of irrational choices, the one with lowest utility being the most irrational choice.

Although simple to solve, these kinds of decision problems contain many of the concepts and methods central to practical reasoning. For the next level of decision problems we will (i) retain the condition of certainty but (ii) move from simple to complex goals, and accordingly move from single to multi-criteria option evaluation. These more complex kinds of decisions will give us the opportunity to introduce several more concepts and methods for making good (that is: rational) choices.

EXERCISE: For each of the following three single-criterion decision problems:

- (a) identify a simple goal,
- (b) form a single criterion,
- (c) evaluate the options, and
- (d) identify the rational choice.

Decision problem #1: Deciding on a car.

Your goal is to own a good used car, and you have narrowed your option down to these three. All are in equally excellent condition, body and engine:

Car #1: A large station wagon. 4-wheel drive. 17 mpg. \$6500. 35,000 miles, 5 years old. It has a great aftermarket stereo system (radio, CD, MP3 audio input). This car is made in the USA. It comes with 1 year warranty parts and labor.

Car #2: Front wheel drive, 30 mpg, mid-size passenger car. \$8500. 65,000 miles. 7 years old. Radio only. This car is made in Japan. It comes with 10,000 warranty, parts only.

Car#3: Rear wheel drive, 40 mpg, small, sporty-type compact car. \$7500. 55,000 miles, 6 years old. Standard radio with CD player. This car is made in Europe. 3 years warranty parts and labor.

Which vehicle is the rational choice for you, relative to your single criterion?

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. You have narrowed your options down to the following three pieces of available real estate:

Facilities #1: A section of a very famous mid-town skyscraper 25 miles from your present location. 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. 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. 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 maintenance is included. Rent is \$125,000 per month.

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.

Which piece of real estate is the rational choice for you, relative to your single criterion?

Decision problem #3: Deciding on a night out.

Your goal is to enjoy a Saturday night out. 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.

The party: Several of your friends will be there, and three of your really close friends. It starts around 9, but will become much more fun around 11PM and will end around 2AM. You will enjoy the music, the conversation, the food, and meeting new people. 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. It starts at 9 and you'll be home by 11:30PM. The movie will cost \$10.

The concert: You know that it will draw a big crowd, and three good friends will be there. It starts at 9 and will end around midnight. You enjoy this band a lot. Tickets are \$55.

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. Only one good friend will go to the game with you. It starts at 7:30 and will end around 10:30PM. Tickets are \$30.

What is the rational choice for you do Saturday night, relative to your single criterion?

Sources and Suggested Readings: Chapters 3 and 4 develop the same topic: individual decision making under conditions of certainty. See the Sources and Suggested Readings section in Chapter 4 for important material covering both single and multi-criteria decisions.