CHAPTER 8: INDIVIDUAL DECISION UNDER IGNORANCE

Sometimes an agent is faced with a decision and cannot form any reasonable degree of confidence that an option will yield the desired outcome. The agent is uncertain whether or not the state will be in place – the state required for an option to result in the desired outcome – and is uncertain to such a degree that the agent can't even estimate the probability that it will. What to do? This is **decision under ignorance**, the topic we will be examining in this chapter.

8.1 Ignorance in decision making

In decisions under ignorance, why is it unreasonable for the agent to try to assign probabilities to states? One possibility is that there is no evidence concerning the probabilities in question, and so the agent can't be expected to have such evidence. Another possibility is that the agent has only incomplete and imperfect evidence. A third possibility is that the agent has good evidence that the existing statistics and probabilities are inaccurate and thus useless. In each of these three categories of ignorance, the agent is aware of the condition of ignorance concerning the probabilities in question. Also, the agent realizes that assigning probabilities to states and forming degrees of confidence on this basis will more likely than not result in false and unreliable assignments. It is important to see that "ignorance" in this connection does not mean that the agent does not know anything. Quite the contrary, the agent knows something that is crucial to the decision situation: the agent knows that probabilities cannot be correctly assigned to states and so the agent should not try to do so. In decisions under ignorance, the agent is aware that she is better off using no information about states rather than using wrong information. Here is an example that we can vary to illustrate these three conditions of decisional ignorance.

Suppose that a college student, Erin, is planning to vote for class president. Let's say that there are two candidates competing for class president: Alice and Bob. Erin has two options: vote for Alice or for Bob. For each option, there are three possible states, each with its outcome. State 1: the majority of student voters vote for Alice (Bob), in which case the outcome is that Alice's (Bob's) winning vote count is increased by one and Erin's vote does not make a difference in the election. State 2: the minority of student voters vote for Alice (Bob), in which case the outcome is that Alice's (Bob's) losing vote count is decreased by one and Erin's vote again does not make a difference in the election. State 3: the student vote is a tie between Alice and Bob, in which case the outcome is that Erin's vote breaks the tie and her candidate wins by one vote. Here is the decision structure.

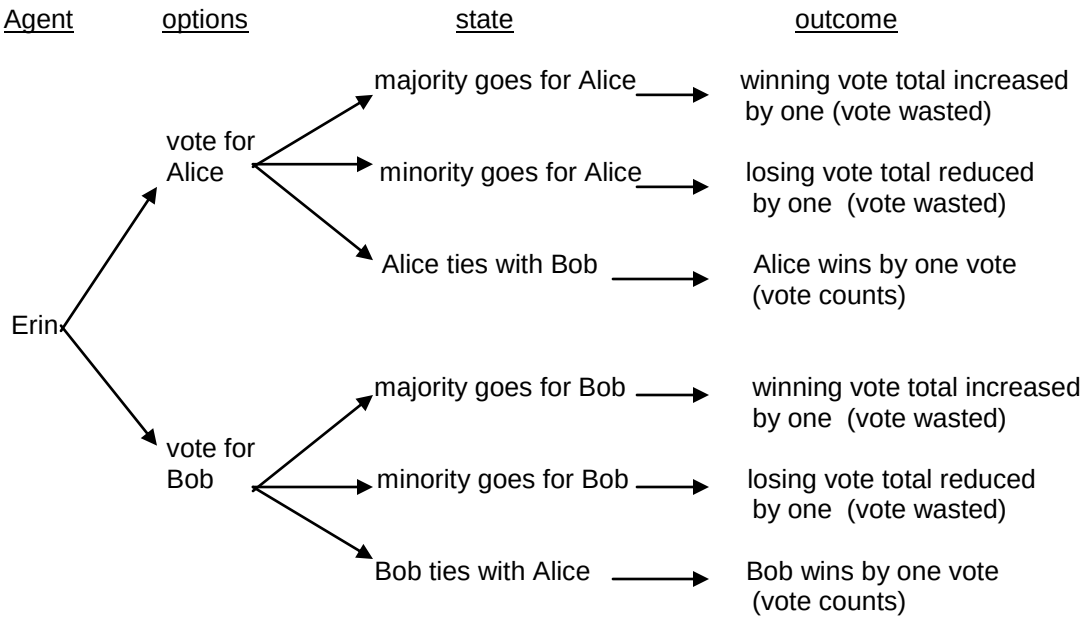

For each option, the states sum to 1.0. Erin is certain that one of the three states will be in place; practically speaking, there are no other possibilities. For this to be a decision under ignorance, Erin must realize that she is not able to assign probabilities to the states in question. The first possibility mentioned above would be a case in which Erin has no evidence at all about how the voting students will vote. Let's suppose that no polls have been taken on campus, that students

have kept quiet about how they will vote, and that there is no sense at Erin's college who might win the election for class president. So, she has nothing on which to base an estimate of the probabilities of these states. Erin realizes that it would be foolish to try to assign probabilities to these states.

The second possibility mentioned above would be a case, let's suppose, in which Erin is aware that a poll was started, but an insufficient number of students responded. As a result, there is incomplete data about how the majority of voting students would vote in the class president election, so incomplete that no reasonable probability estimates can be based on such a poor sample of responses. Erin realizes that she cannot assign probabilities to states on such weak evidence.

Finally, as an example of the third possibility mentioned above, suppose that the campus newspaper did a poll showing that Alice would win by a large margin. However, a scandal has erupted around this poll. A journalism student has blown the whistle and revealed that the so-called poll was deliberately falsified by some of the polltakers who secretly belonged to Alice's election committee. Once Erin gets wind of this, she realizes that the falsified poll results cannot be used as evidence on which to base probability estimates of the states.

These three examples have this central point in common: the agent <u>understands</u> that no reasonable degree of confidence can be formed that an option will result in a given outcome. This is decision under ignorance. To head off a possible confusion, let's contrast two different statements.

(1) The agent has no reasonable degree of confidence that an event will happen.

(2) The agent has a reasonable degree of confidence that an event will not happen.

The event in question in these statements is that an option, if chosen, will result in the intended outcome. The possible confusion is that someone might take these statements to be saying the

same thing. This is not correct. The first statement describes decision under ignorance. However, the second statement does not, and in fact describes decision under risk, for it will yield the appropriate degree of confidence by the formula: P(a) = 1 minus P(not a). (See rule 2, Chapter 5 to confirm this).

Decisions under conditions of ignorance might appear to be hopeless situations, especially when compared to decisions under certainty but also compared to those under risk. If an agent has no idea at all whether or not an outcome will result from doing an action, how can an agent possibly make a rational choice among options? Decisions under ignorance, however, are by no means hopeless "shots in the dark". Practical reasoning can still operate in this kind of decision situation and arrive at a rational choice. Let's look at three different versions of decisions under ignorance: those that are "painful" choices, those that are "can't lose" choices, and those in which the agent has partial probability information.

8.2 Painful choices: decision by maximin

Some decisions under ignorance are bad decision situations, which nevertheless an agent can't avoid; they must be faced. The agent's options, no matter the state, have outcomes that, given the agent's goal, all look in various degrees bad, harmful, or represent damage. The outcomes are undesirable, have disvalue, and yet a decision must be made. Think of a battlefield decision where outcomes are lives lost, or decisions in the midst of a natural disaster about people who are to be left to die. The classic case of the overcrowded lifeboat represents a painful choice well, and is widely used as a vivid illustration. Five people, let's say, must go overboard to die or the entire lifeboat will sink and all 25 will die. Imagine that no one volunteers to sacrifice his or her life. A decision must be made whom to force overboard, a terrible decision that can't be put off. There are very real cases like this, not only involving lifeboats but also situations calling for rationing emergency medical assistance or distributing life saving drugs when there is not enough. In decision situations under ignorance that are painful choices, the rational choice does

not maximize utility so much as it seeks to minimize disutility. The guiding rule is to minimize the harm that will result and try to choose the lesser of all the evils, so that the smallest possible amount of the goal is lost. Here are the steps to structure and solve painful choices under ignorance.

1) Structure the decision into standard option-state-outcome form with clear descriptions of the harmful aspects of the outcomes.
2) Qualitatively rank the all outcomes using negative verbal terms (e.g. very bad, worse case, somewhat bad, etc.) If helpful, ordinally rank outcomes using 1st for the worse case outcome.
3) Select an interval scale having a sufficiently large negative number range and assign disutility values to outcomes in a way that preserves the information contained in the qualitative ranking.
4) Identify for each option the outcome having minimum utility (equivalently, this outcome will have maximum negative utility or maximum disutility). This will be the worse-case, the security limit outcome, for each option. (We are assuming here that no option dominates all the others, for if one does it will be the rational choice. Similarly, if an option is dominated by all the others, it drops out.)
5) Identify the outcome, from the set of outcomes having minimum utility, having maximum utility (equivalently, the one having minimum disutility). This will be, given the goal, the best of the worse outcomes (the least bad outcome): it is the **maximin** outcome.
6) Choose the option containing the maximum of the minimums: the maximin outcome.

The rational choice rule for painful choices under ignorance is called the **maximin rule**, for the agent first sorts out the minimum outcome utilities and then maximizes utility within these minimums – the agent finds the maximum of the minimums. In this way, the lesser of all the evils is discovered and the agent avoids doing any more harm than is absolutely necessary. Most importantly, the agent discovers the option having the worst possible outcome and is especially

careful to avoid choosing it. Here is the abstract structure for painful choices under ignorance, for four options each having two possible states.

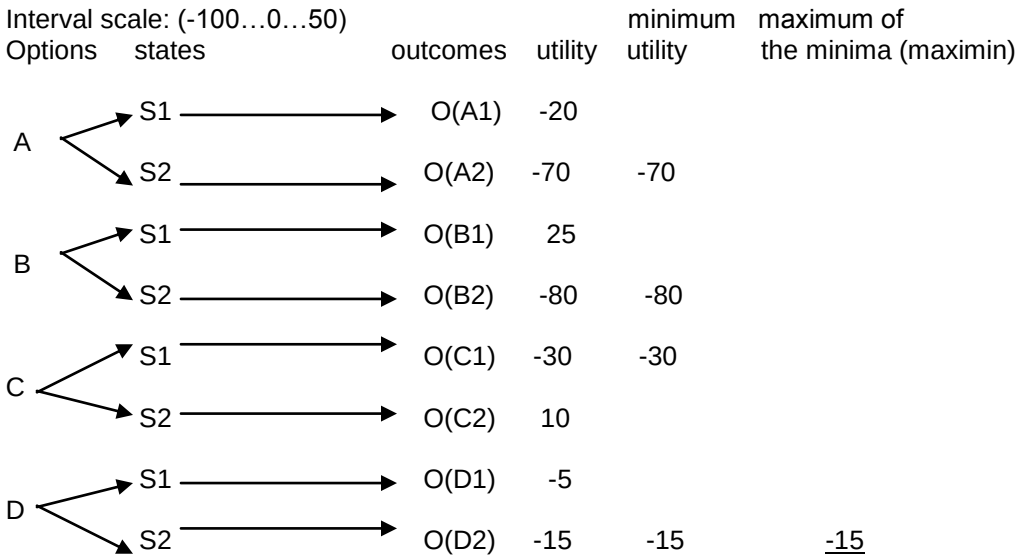

In this abstract structure, I have left out the step of qualitatively ranking the outcomes in their negative aspects. Clearly, outcome O(B2) is the worst case, and O(C1) is pretty bad, and O(B1) is the best outcome. Suppose, for example, we think of these negative utility (disutility) values as numbers of lives lost and the positive utility values as lives saved, and that the goal is to save lives. Without any way of forming degrees of confidence about which outcome each option would result in, it is rational to be gloomy, suspicious, and to choose as if one were a pessimist expecting the worst to happen. It is rational, in other words, to “err on the side of caution” and choose the option whose worst outcome is better than the worst outcomes of the other options, for this guarantees the agent that if the option’s worst outcome results the lesser of all the possible harms will have been done. The rational choice solution is: choose (D) over (C) over (A) over (B).

You might want to argue that, if the goal is to save lives why not choose option B for it has as the outcome with the hope limit, the best case of 25 lives saved. True enough. But the agent is deciding under ignorance, and so there is no degree of confidence that O(B1) will result. If it does not result, then O(B2) will be the outcome and it is clearly the one outcome to be avoided at all

costs. Given the goal of saving lives, it would be irrational to choose option B in the optimistic hope that things will turn out well and 25 lives will be saved when, for all the agent knows, B will result in 80 lives lost, the security level of the option set. Instead, it is far more rational, given the goal, to choose the option having the least harmful security level, even if it does not contain the decision's hope limit.

Here is the rational choice rule for painful choices under ignorance:

For any two options x and y: (i) If minU(x) < minU(y), then choose x over y
and (ii) if minU(x) = minU(y), then be indifferent between x and y.

A word on terminology: In the above 6 steps for structuring and solving painful choices under ignorance, we indicated that "minimum utility" and "maximum disutility (negative utility)" are two ways of referring to the same value. In the above abstract structure, -80 is both the minimum utility and the maximum disutility. The maximin rule was stated in terms of minimum utility: first select minima and then select from them the maximum. If, alternatively, we use the term "maximum disutility" instead of "minimum utility", then the maximin rule tells us first to <u>maximize disutility</u> and then select the outcome from these having <u>minimum disutility</u>. The steps are the same, and the values selected are the same, it is just that the terms we use to describe this process of practical reasoning change slightly. Here is the equivalent rational choice rule using "disutility" (negative utility) instead of utility:

For any two options x and y: (i) if max -U(x) < max -U(y), then choose x over y,
and (ii) if max -U(x) = max -U(y), then be indifferent between x and y.

This rule says: (i) if the maximum disutility (-U = disutility or negative utility) of option x is less than the maximum disutility of option y, then x should be chosen over y. Similarly for part (ii) concerning indifference.

EXERCISE: Using the maximin rule, solve the following painful decision problem. This is a single criterion individual decision under ignorance, multi-stage states. The decision is already structured based upon the narrative. Base your qualitative ranking and utility values on your own beliefs and values, keeping to the narrative.

You are a medical social worker who must service a very dysfunctional family containing five children. There have been on-going periods of abuse of the mother by the father and some of the older children, yet the entire family is psychologically dependent on the mother – she holds them together as a family. Several children suffer depression, and some have attempted suicide when the mother had to be hospitalized in the past. There is a history of substance abuse on the part of some of the children, and the father is an alcoholic. In spite of the deep and serious problem this family experiences, your goal is to keep this family together, for breaking this family up will almost certainly result in greater harm than if it remains together. Your single objective right now is to do all you can to stabilize the lives of the family members.

Here is your dilemma: it has just come to your attention that the mother has been diagnosed with a fatal illness and is expected to die in about a year. You have the responsibility to give this sad news to the family or to withhold it from them. If you tell the family, you fear as an outcome that two especially vulnerable children and perhaps even all of the children could go into depression and might attempt suicide. But telling might diminish the amount of abuse suffered by the mother. If you withhold the information concerning the mother's medical condition, the mother might suffer periods of abuse, which could hasten her death; but her fatal illness will not cause an increase in the risk of depression and suicide attempts for the time being. You do not have enough personal experience with this family to form reasonable degrees of confidence about the outcomes, and the social services data on this family is too incomplete to be of any use as evidence to assign risks.

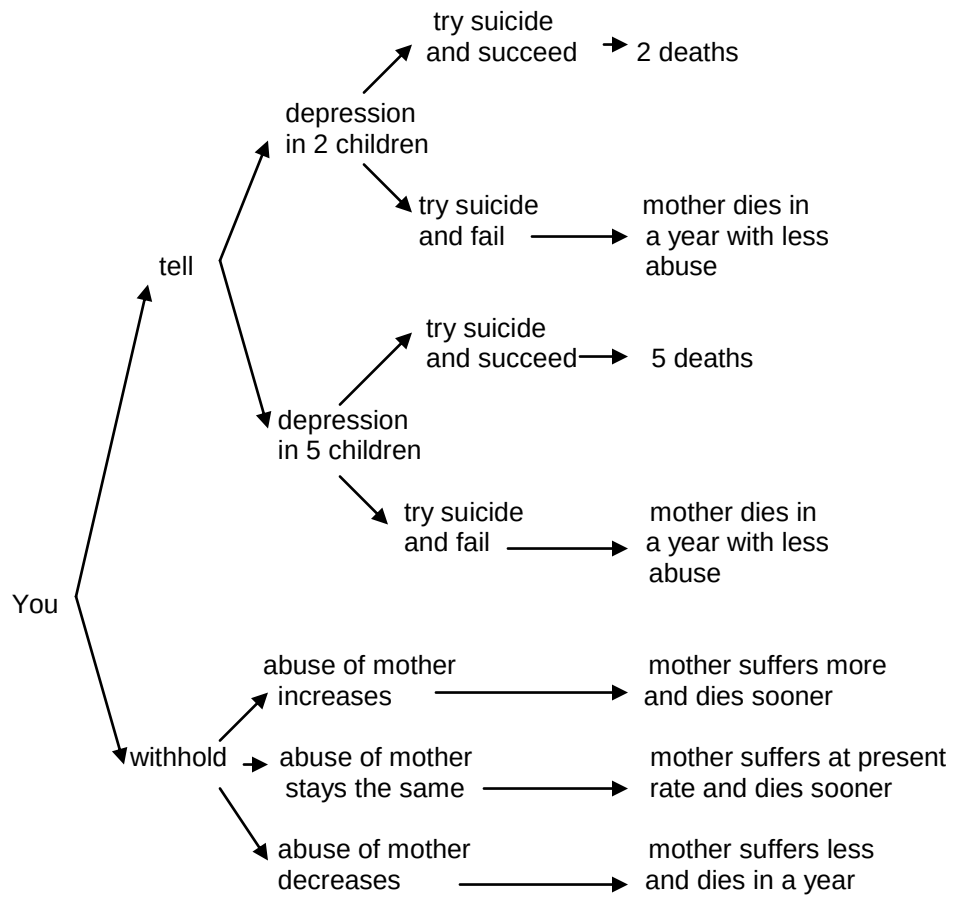

8.3 Can't lose choices: decision by maximax

The second kind of decision under ignorance, in comparison to painful choices, are pleasant decision situations, ones in which the agent can't lose the goal no matter the outcomes. The agent's options and states yield outcomes that, given the goal, all look good. The goal would be largely achieved with even the "worse" outcome. In such a decision situation, the agent can throw caution to the wind and go for the best of the best. It would be irrational to be the cautious pessimist in can't lose decisions, the agent should try for the hope limit and need not worry about avoiding the decision's security level. The guiding rule is to maximize the good, and try to choose

the greater of all the goods. Here are the steps to structure and solve can't lose choices under ignorance.

1) Structure the decision into standard option-state-outcome form with clear descriptions of the desirable aspects of the outcomes.
2) Qualitatively rank the all outcomes using positive verbal terms (e.g. very good, best case, terrific, pretty outstanding, etc.) If helpful, ordinally rank outcomes using 1st for least good.
3) Select an interval scale having a sufficiently large positive number range and assign utility values to outcomes in a way that preserves the information contained in the qualitative ranking.
4) Identify for each option the outcome having maximum utility. This will be the best outcome for each option. (We are assuming here that no option dominates all the others, for if one does it will be the rational choice. Conversely, if an option is dominated by all the others, it drops out and the agent works with a smaller set of options.)
5) Identify the outcome, from the set of maximum utility outcomes, having maximum utility. This will be, given the goal, the best of the best outcomes, the hope limit: it is the **maximax** outcome.
6) Choose the option containing the maximax outcome.

The rational choice rule for can't lose choices under ignorance is called the **maximax rule**, for the agent first sorts out the maximum outcome utilities and then maximizes utility within these maximums. In this way, the greater of all the goods is discovered. Importantly, the agent does not have to worry that the option containing the maximax outcome, the hope limit, might also contain the decision's security level, for in can't lose choices the security level outcome still achieves the goal, or most if not all of it. Here is the abstract structure for can't lose choices under ignorance, for four options each having two possible states.

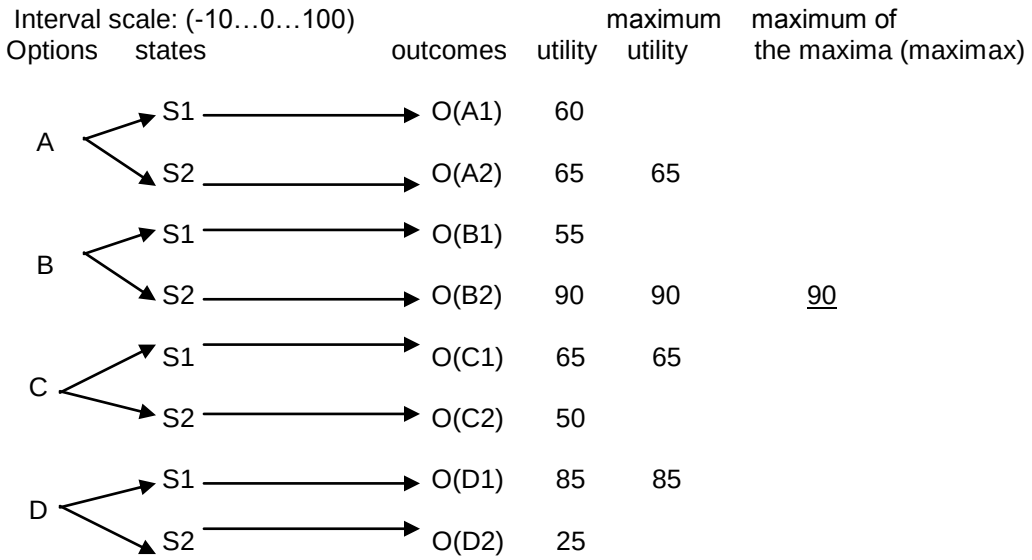

In this abstract structure, I have left out the step of qualitatively ranking the outcomes in their positive aspects. Clearly, outcome O(B2) is the best case, and O(D1) is outstanding, and even O(D2) is a good outcome. Suppose, for example, we think of these options as four outdoor concerts of a band that you love, and the states represent concert conditions: S1 = warm night and S2 = cool night, for the four different concert locations. All outcomes fully satisfy your goal of an enjoyable concert by this band. Which one do you buy tickets for? Without any way of forming degrees of confidence about which outcome each option would result in, it is rational to be a hopeful optimist and choose as if you were expecting the best to happen; you can't lose! It is rational, in other words, to "throw caution to the wind" and choose the option whose best outcome is better than the best outcomes of the other options, for the agent is guaranteed that the "worst" outcome of the option containing the maximax outcome, should it be the result, will still largely achieve the goal. The rational choice solution for this "can't lose" decision problem is: choose (B) over (D) over (A) and be indifferent between (A) and (C).

Here is the rational choice rule for can't lose choices under ignorance:

For any two options x and y: (i) If maxU(x) > maxU(y), then choose x over y

and (ii) if maxU(x) = maxU(y), then be indifferent between x and y.

EXERCISE: Using the maximax rule, solve the following can't lose decision problem. This is a single criterion individual decision under ignorance, single-stage state. The decision is already structured based upon the narrative. Base your qualitative ranking and utility values on your own beliefs and desires, keeping to the narrative.

Your goal is to have a great vacation and your options are these: Hawaii, Florida, or France. The costs turn out to be about the same for each option, so it is not a factor in your decision. You love the outdoors and each option will provide a great time outdoors, especially the beaches, if the weather is good. But even if the weather is bad, there are very interesting things to do: wonderful touring and sightseeing in Hawaii, exciting theme-parks and science centers in Florida, and outstanding museums and cultural events in France. Where should you go? Each option will provide you with a great vacation. It is five months before your vacation and you must decide now, but you can't get any reliable information about the weather in any of these places for the time of your stay that would help your decision; the brochures all say things like "prefect weather all the time for all your activities" which you realize can't be true.

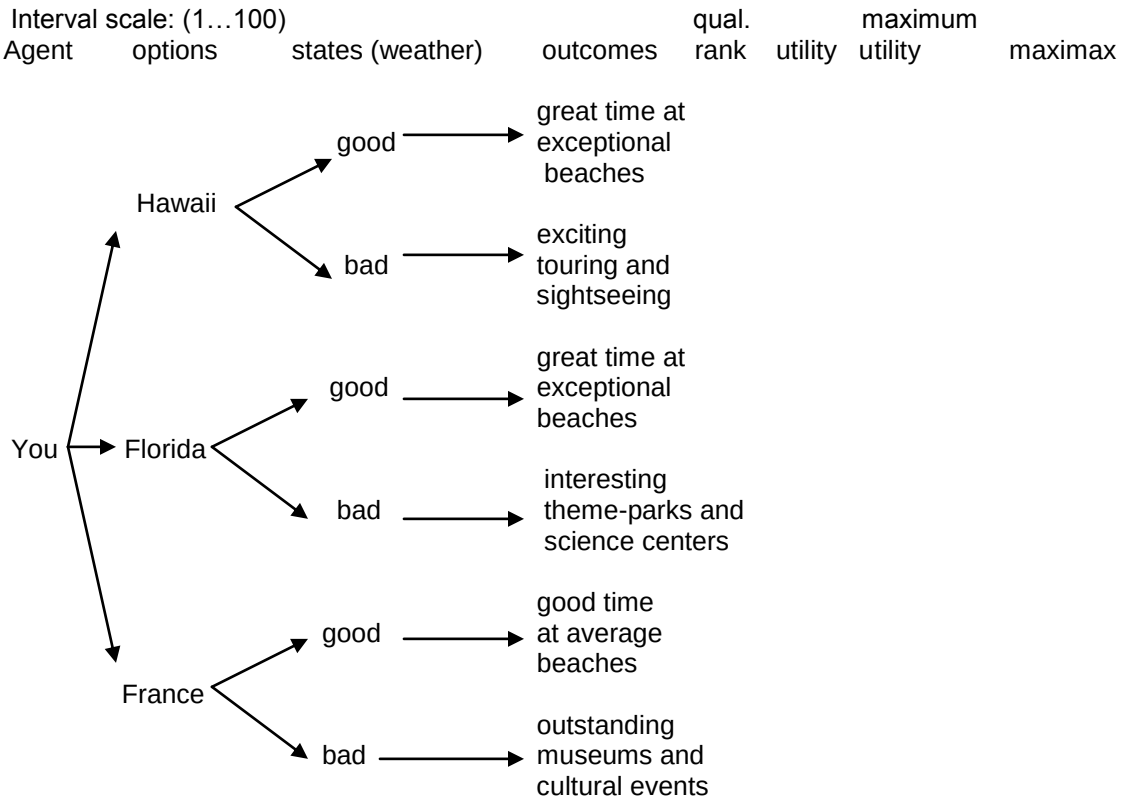

8.4 Decisions with incomplete information: decision by ordinal ranking

In decisions under risk, by definition, the agent has sufficient knowledge of the probabilities of the states to form reasonable degrees of confidence that an option will result in a given outcome. The two kinds of decision under ignorance covered above, painful choices and can't lose choices, were presented as decisions problems in which the agent had no evidence at all on which to base reasonable degrees of confidence. But having evidence about the probabilities of the required state is sometimes not an all-or-nothing affair. An agent could have some useful information about the state's probabilities, some general ideas about their likelihood, and yet still not have sufficient evidence to assign specific probabilities and form specific degrees of confidence. In reality, this could well be the most common decision situation we face. For example, an agent may be able to tell that one state is more (or less) probable than another state, and thus be more

(or less) confident about one than another outcome, without being able to assign specific probabilities to the states in question and so not be able to form degrees of confidence concerning the outcomes in question. Even though the agent's information about the states is partial and imperfect, it can be useful if it is minimally sufficient to form an ordinal rank of the probabilities of the states, and from this to form an ordinal rank of confidence concerning the outcomes. Here are the steps for solving a decision problem under ignorance by ordinal ranking.

1) Structure the decision into standard option-state-outcome form with clear descriptions of the outcomes.
2) Qualitatively rank the all outcomes using verbal terms (e.g. very good, best, worst case, so-so, pretty good, etc.) relative to the goal, and transform this into an ordinal utility ranking with the lowest ordinal number assigned to the worst outcome and the highest ordinal number assigned to the best outcome.
3) Using the partial information concerning the states, qualitatively rank the probability of the outcomes using verbal terms (e.g. least likely, most probable, not very confident, pretty good chance, etc.), and transform this into an ordinal rank with the lowest ordinal number assigned to the outcome in which there is least confident it will happen and the highest ordinal number assigned to the most probable outcome.
4) For each outcome, multiply the two ordinal numbers (ordinal confidence X ordinal utility). This is the outcome's ordinal ranking value (ORV).
5) For each option, sum the ordinal ranking values. This gives for each option a total **ordinal ranking value**.
6) Choose the option having the greatest total ordinal ranking value.

Solution by ordinal ranking is clearly not as refined practical reasoning as solution by expected utility, and thus not as good a method as the latter for discovering the rational choice. Yet it is clearly better than having to make a decision when there is no probability information at all on which to base degrees of confidence, as for example in the pessimist case of solution by

maximin. Here is the abstract structure for decisions with incomplete information, for three options each having two states.

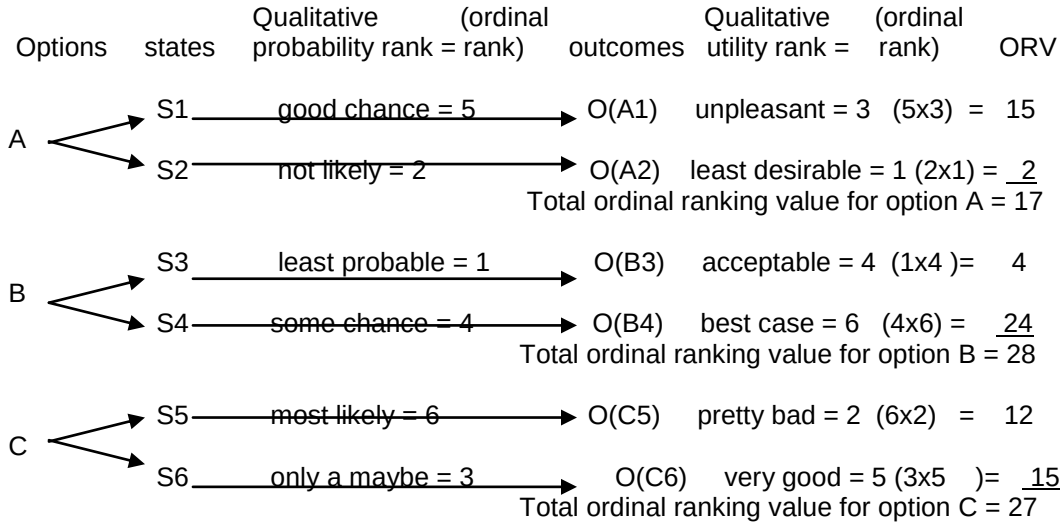

In this abstract structure, the qualitative ranking of the agent's degree of confidence and its ordinal number have been placed between the state and the outcome. The qualitative ranking of outcome utility and its ordinal number have been placed after the outcomes. These two ordinal numbers are multiplied and the products summed for each option to give the total ORV of the option. The rational choice solution for this decision problem is: choose (B) over (C) over (A)

Here is the rational choice rule for decisions with incomplete information, solution by ordinal ranking value:

For any two options x and y: (i) If ORV(x) > ORV(y), then choose x over y

and (ii) if ORV(x) = ORV(y), then be indifferent between x and y.

In the above abstract example, notice that options B and C almost tie. Yet, by examining the outcomes it becomes clear that B beats C, if only by a little. Let's see why this is so. A better chance for a more desirable outcome (O(B4)) certainly beats a lesser chance of a less desirable outcome (O(C6)). And, a low chance for an acceptable outcome (O(B3)) is surely better than the

highest probability for a bad outcome (O(C5)). But also notice that option A contains the worst outcome (O(A2)) and that solution by ordinal ranking value makes it the most irrational choice.

EXERCISE: Using the incomplete information provided in the narrative, solve the following decision problem by ordinal ranking value. (Any equi-probable states or equi-desirable outcomes should be assigned the same ordinal rank.)

Imagine that you are housemates with 3 other people: Ann, Beth, and Charlie. You each have separate bedrooms, but share a common kitchen. One night you wake up at 3:00AM very hungry and realize that won't be able to get back to sleep unless you have something to eat. You go to the common refrigerator and see two containers of leftovers. Suppose that you can snack on one container, but not both. You must choose which one to take food from. The first looks like leftover soup. If Ann made it, it should be very good; she typically makes nice soups. But Ann has been complaining how busy she has been lately, so it's not very likely she made the soup. If Beth made it, it will be okay; soups are not really her cooking strength. You remember that you have been seeing Beth in the kitchen a lot lately, and think to yourself that there is a good chance she made the soup. Charlie is a wiz in the kitchen, and used to work as the soup-chef in several local restaurants. His soups are amazing creations. But you haven't seen Charlie around much in the last few days, so it is quite unlikely the soup is his.

The other item in the refrigerator seems to be leftover stew of some sort. Sometimes Charlie brings home leftover stews from the restaurant where he now works. Their stews are not bad, almost as good as Ann's soups. If Beth made the stew, it will be a real treat, for she loves doing stews and they turn out almost as good as Charlie's soups. But Beth uses only fresh ingredients in her stews, and it is highly unlikely that she has done any shopping for fresh ingredients lately.

Finally, Ann's stews are awful. She keeps trying to cook a good stew but they are all failures. Because Ann has been so busy, it is not very likely she made the stew.

You want a nice late-night snack, and could go for either the soup or the stew. Which one should you pick?

8.4 Practical reason within bounds: decision by satisficing

There is an ongoing debate concerning the relation between the theory of rational choice and the actual decision making behavior of real agents. A hint of this debate was given in Chapter 3 where you were imagined to have a "No way!" reaction to deciding on a movie to see. We will call it the rationalist-behaviorist controversy. In brief, the debate is this. The rationalists argue that the principles and norms of practical reasoning that are contained in the theory of rational choice are the standards by which to judge and justify how rational the actual decisions of real agents are. If we find that the actual decisions of real agents are not in keeping with decisions of the theory's ideally rational agents, then we have good reason to claim that these real agents are making irrational choices; they are irrational in the practical sense when judged by these ideals. The decision behavior of real agents should change to fit the norms of practical rationality.

The behaviorists, on the other hand, argue that when the actual decisions of real agents don't agree with and can't be justified by the norms and standards of rational choice, this actually serves to falsify the theory and shows that something is wrong with these principles of practical rationality. The theory should be changed to fit the facts about how real agents make actual decisions, and how they actually do their practical reasoning. If a theory about rational decision making isn't based in the reality of human decision making, how can it possible apply to human decision makers?

You will know from Chapter 1 that this text is squarely in the rationalist camp. But for the purposes of the topic in this section we will take the point of view of the behaviorist. This will allow us to step back and form a general overview of practical reasoning, as it has been developed thus far, as well as introduce an interesting way of structuring and solving decision problems – satisficing.

What do all the forms of practical reasoning and rules of rational choice that have been presented up to this point have in common? They all require that the agent consider all the options available and evaluate them all in order to discover the best option (or options, in cases where two or more options tie for best place). The "best" option is defined as the one whose outcome is expected to achieve at least as much as, and hopefully more of, the goal as any of the other options. This is the case for decisions under certainty where the agent reasons in a way that will discover the option having maximum utility. This is the case for decisions under risk where the agent tries to discover the dominant option or the option having maximum expected utility. And this is the case for decisions under ignorance that are made by the maximin, the maximax, or the ordinal ranking value methods. So, according to the theory of rational choice, the ideally rational agent always tries to find the option(s) from the menu that gains the agent the greatest amount of the goal – whatever that goal happens to be. Two terms are used in the literature to describe this general feature: "maximize" and "optimize." In the theory of rational choice, a rational agent is a maximizing (or an optimizing) agent.

Now in one sense, this seems perfectly correct. If an agent is trying to achieve a goal, isn't it reasonable to assume that the agent is trying to achieve all of it? And if the agent can't achieve the whole goal, isn't it reasonable to assume that the agent is trying to achieve as much of it as possible, given the options? Isn't the agent that chooses the sub-optimal option making a bad decision, making an irrational choice? The theory of rational choice answers "yes" to these questions.

But the behaviorist has a very interesting criticism of this idea that the norm of practical reasoning is to optimize, no matter what the decision problem is. The behaviorist, on the basis of factual studies of how real agents go about making decisions, argues that the theory of rational choice is too abstract and in a way too simple to be of any real use in guiding practical reasoning. It is argued that real decision problems are often far too complex and real agents are often under far too many pressures, to optimize. The theory of rational choice, then, expects too much – it asks too much – of real agents. It sets up an impossible ideal by establishing as rational only optimizing/maximizing forms of practical reasoning. The behaviorist argues that the standards of rational choice should be lowered to something that real agents can meet, given the massive complexity and pressures under which real decisions must typically be made. This idea of practical reasoning within the bounds that are imposed on agents by the complexity and pressures of real decision situations is called the theory of **bounded practical rationality**. The ideal of practical reasoning in this theory is not that of optimizing, it is that of **satisficing**. The term is designed to combine two ideas: (i) being *satisfied* with a sub-optimal option, and (ii) *sacrificing* the optimal option.

What kinds of limits or boundaries do agents typically take into account which would make optimizing unrealistic and satisficing the reasonable way to go? Here are some examples.

(a) Time pressures: an agent might be under time limits and deadlines that are too brief to allow for all the practical reasoning needed to maximize. Suppose an emergency suddenly arose and you had to make an important decision in, say, 7 minutes. The decision problem might involve many factors such that if you tried to maximize you would need several hours to find the rational choice. Behaviorists claim that many of life's real decisions have time limits that make optimizing impossible to achieve.

(b) Low goal value: sometimes an agent has to achieve a goal that is relatively unimportant, given other things going on in the agent's life. For example, suppose that you have to get new tires for

your car so that it will pass inspection, but you are involved in so many projects that you don't especially care what kind of tires you get. Or, it's time to eat but you are in the middle of interesting activities. You won't have the opportunity to eat later, so you desire to eat now but don't care much what or where you eat so long as you can get back to your interests.

(c) Limited resources: optimizing might require more resources than an agent has. Money is an obvious case in point. Imagine an agent sending out job applications to many companies that are located far away. She receives back, say, 10 requests for an interview. In theory, maximizing would require that she goes to all 10 interviews, and (assuming that she receives a job offer from each company) picks the best job from her 10 options. But suppose that she can't afford these 10 trips, they are far away and so each would require an expensive airline flight, hotel and local transportation costs, meals, etc. Maximizing is simply beyond her limited financial resources.

(d) Decision complexity: the decision might be too complex to allow the agent to optimize. Suppose, for example, an agent lives in a large city having thousands of restaurants. The agent desires to eat out, but has no particular desire for one type of food over another that would serve to narrow the options; and let's add that money is no problem for this agent. If this agent were really open to eating out anywhere in the city, optimizing would involve deliberations about such a large number of options that the decision which restaurant to eat at would become too difficult to handle.

The behaviorists point to empirical studies about the way real agents make choices within the bounds imposed by the limits, pressures, and complexities of the decision situation. Bounded practical rationality often settles, these studies show, for an option whose outcome is "good enough," given the goal, and does not try to optimize. Thus, satisficing would seem to be verified by empirical research. According to the theory of bounded rationality, then, an acceptable choice *is* a rational choice in many decision situations, and the optimal choice would be a foolish standard – irrational! – to try to live up to.

8.4.1 Satisficing

Satisicing is a way of solving decision problems that depends far more on the agent than maximizing does. The agent determines what the bounds of the decisions situation will be according to the agent's own sense of time-lines, convenience, and resources. Also, the agent determines what is acceptable or not both with respect to utility and probability. Finally, the evaluation of options tends to remain qualitative and vague, with no more precision or detail then the agents believes necessary. Here are typical steps for decisions that are uncertainty and so require more than one state per option.

1) Given the goal, the agent forms any set of options that respects the limits, pressures, and complexities of the decision situation, as the agent sees these.
2) Utility is simplified into two categories: for example, acceptable (= good enough) or unacceptable (= not good enough), or ok (= I can go with it) or not ok (= I can't go with it). Given the goal, the agent places each outcome of each option in one or the other category.
3) The agent simplifies the probability into two probability ranges, an upper and a lower. These might be very loose such as the pair: likely/unlikely, or the pair: sufficiently probable/sufficiently improbable. Or the agent might feel the need to make the ranges a bit more precise by establishing a probability cut-off point such as: greater than .75/less than .75, or: above 2/3 chance/below 2/3 chance. The agent places each state into one or the other probability range.

The rational choice rule for satisficing: the agent chooses any option having an outcome that is both acceptable (= good enough) and within the upper probability range (= likely enough to happen). The agent is indifferent between options that come out equal by satisficing. Here is an example of a decision made by satisficing.

Suppose that you have had a hard, frustrating day at work and come home to a very stressful home situation. You reach a point where you can no longer cope with things. You realize that need to get out of the house for the evening. You don't much care where you go so long as you can forget your troubles and don't have to deal with a lot of other people. Three possibilities quickly come to mind: you can go to a movie, you can go to your favorite small cozy restaurant, or you can go to the Mall shopping. Three options are enough to consider, for you are not in the mood to do a lot of thinking about where to go, you just want to get going. Here are your quick thoughts about these three options. If the movie is crowded, which is not very likely, it is unacceptable. But if it has either low or medium attendance, it is an acceptable place for you to get away for the evening. It's a pretty popular movie house, so it is unlikely that it will have low attendance. If the small restaurant has low attendance, which is likely given that this is a week night, it is acceptable. But if it has medium attendance or is crowded, it is unacceptable. Finally, if few people are at the Mall, not very likely for there are many sales going on, it is acceptable. But if the mall is crowded, also not very likely due to the fact that this is a week night, or has medium attendance, there will be too many people for you to relax and so will be unacceptable. You are impatient to get going and so don't want to give your decision any more thought than this. Where should you go to get out of the house for the evening and get a break from all the stress? Here is the decision structure.

Goal: to get out of the house for the evening to a place that's not crowded, as a stress-break for myself.

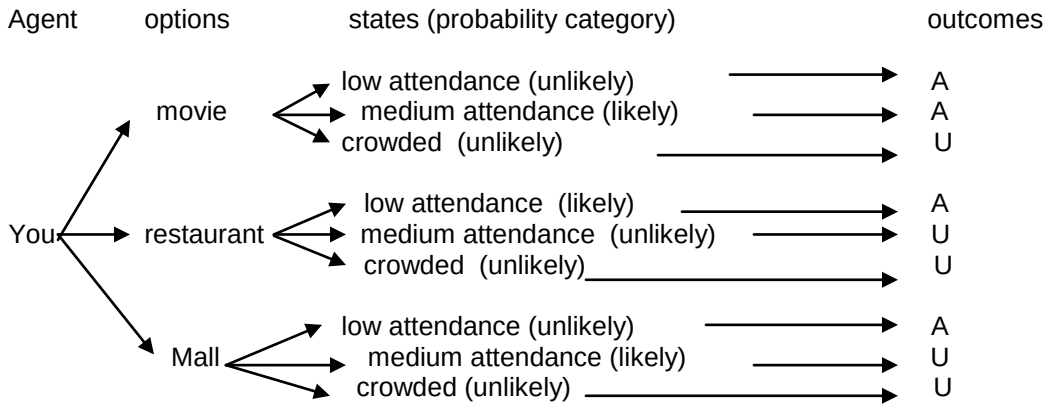

Both the movie and the restaurant have outcomes that are both likely and acceptable. The Mall's only acceptable outcome is unlikely. Thus, the satisficing solution to this decision problem is: be indifferent between (movie) and (restaurant) but choose either (movie) or (restaurant) over (Mall). At this point, you only have to flip a coin to settle between movie or restaurant.

8.4.2 Strengths and weaknesses of satisficing

The chief virtue of bounded practical reasoning, as the above stress-break decision problem is meant to highlight, is simplicity. This makes it an easy, efficient, and fast form of practical reasoning. And it yields a solution that is good enough to the agent. The agent's thoughts as described in the narrative are enough for the agent to make her mind up, and making a full-blown decision analysis/evaluation with all required details filled in is entirely unnecessary. Also, if it is true, as the behaviorist position claims, that real decision makers actually are satisficers and not maximizers, then the theory of bounded practical reasoning is not only verified by the facts, it's standards allow us to judge the actual decisions of real agents to be by-and-large rational choices.

The chief weakness of bounded practical reasoning, again as the above decision problem is meant to highlight, is its simplicity. This makes it prone to error – "error" of course, relative to the standards set by the theory of rational choice for practical reasoning that maximizes. In solutions by satisficing, for example, the agent is inevitably going to overlook options with outcomes that will better attain the goal than those being considered. In the decision problem described above, let's suppose that if you had thought about your options more systematically you would have realized that a forth option would have been to visit a friend. And suppose that, had you been more refined in analyzing your goal, assigning utilities to outcomes and probabilities to states, you would have discovered that visiting with this friend would have made for a much more satisfying evening, given your goal, than either the movie or the restaurant. Suppose that a visit with this

friend would have done wonders for your mood, whereas the movie or the restaurant (whichever you ended up doing) didn't help you all that much as you eventually came to realize. However, you missed out on this optimal option of visiting your friend because bounded practical reasoning does not require the agent to try to live up to the highest standards of practical reasoning.

EXERCISE: Solve the following decision problem by satisficing. Note that there are multi-stage states for the first and the last options.

You have two part-time jobs. You have been delayed at the first one and are rushing over to your second job. But first, you need to grab a bite to eat. You look only on the side of the street where you are now walking and see three possibilities in front of you: a sandwich shop, a pizza place, and a coffee shop. You are not sure which one to go to and have no time to go back and forth checking each one out. Here are your thoughts: "If the sandwich shop has ready-made sandwiches and not a lot of customers, it will be acceptable. It probably has ready-made sandwiches, but it also probably has quite a few customers at this time. If the pizza place has hot pizza slices ready to go, not very likely because I don't see a sign saying "slices to go", that would be great, for I could eat one on my way to my job. It is unacceptable if I have to order a small pizza and wait for it. This is probably the case. The coffee shop will likely have a take out service. If not, I can't wait and order at the counter. If there are many customers at the counter, the take out service will be slow. This is unacceptable. Otherwise it should be pretty quick, and this is acceptable. From the look of things, I'd say that there is equal chance that there are many or that there are few customers at the counter." You have to make a decision right now where to grab a bite to eat and then hurry to your other job. Which place should you go to based on your thoughts and observations?

Sources and suggested readings:

The material in this Chapter on decisions under ignorance draws largely from Mullen and Roth (2002) Chapter 6, Resnik (1987) Chapter 2, and to a lesser extent Luce and Raiffa (1957) Chapter 13. The presentation of bounded practical reasoning and methods of satisficing relies on Byron (ed.) (2004) Introduction, and Mullen and Roth Chapter 6. The study of bounded decision making and satisficing originates with the seminal work of H. Simon. His “Alternative visions of rationality” in Moser (ed.) (1990) Chapter 8 is a clear non-technical argument on behalf of the behaviorist philosophy of bounded practical reasoning. Plous (1993) Section III is a good source for empirical studies of bounded rationality. There are several articles containing clear descriptions of the difference between satisficing and maximizing, and contrasting their relative merits especially in moral decision problems, in Byron (ed.) (2004).