Fundamentals of Artificial Intelligence
Prof. Shyamanta M Hazarika
Department of Mechanical Engineering
Indian Institute of Technology-Guwahati

Lecture-20
Decision Network

Welcome to fundamentals of artificial intelligence, we are discussing reasoning under uncertainty. We have covered probability theory and discuss the syntax and schematics of Bayesian networks. In this lecture we would introduce utility and show how utility is combine with probability theory to yield a decision theoretic agent. An agent that takes rational decisions based on what it believes and what it wants.

Such an agent makes decisions in a context where uncertainty and conflicting goals leave a logical agent with nowhere to go. We introduce decision networks a formalism that extends Bayesian networks by incorporating actions and utilities. But lets quickly review what we mean by uncertainty before proceeding further. So when we say reasoning under uncertainty, uncertainty is about the quality or state of being not clearly known.

(Refer Slide Time: 02:10)

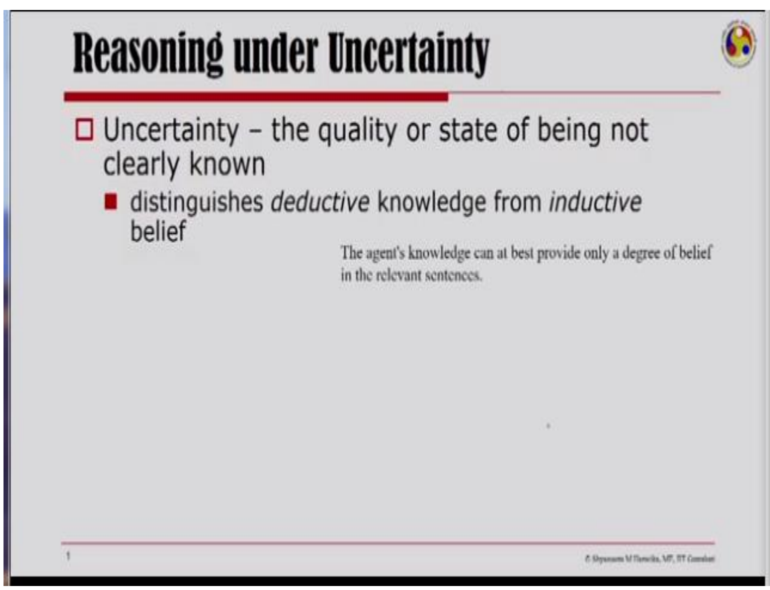

This distinguishes deductive knowledge from inductive belief, the agent's knowledge under uncertainty can at best provide at only degree of belief in the relevance sentences. Now there could be a number of sources of uncertainty.

(Refer Slide Time: 02:36)

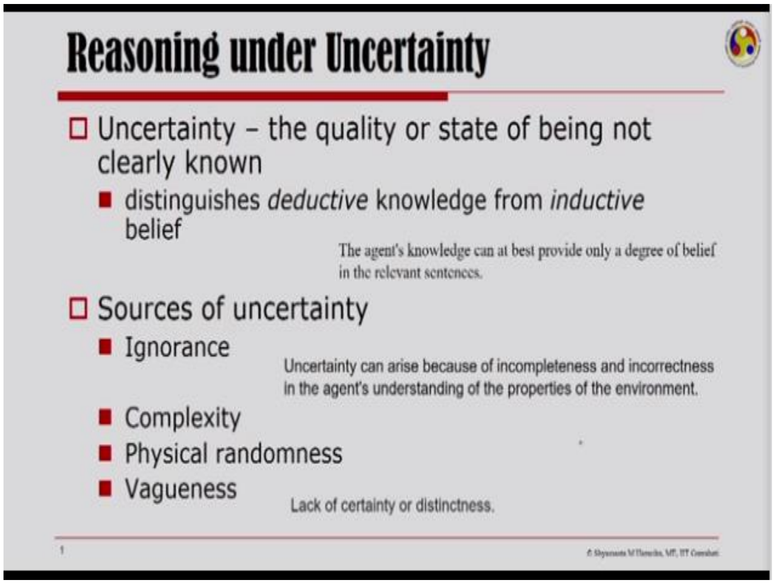

The first and foremost being incompleteness and incorrectness in the agent's understanding of the properties of the environment or what we refer to as ignorance. Thereafter complexity, physical randomness of the system are all sources that act to uncertainty, vagueness or lack of certainty or distinctness is also a source of uncertainty.

(Refer Slide Time: 03:09)

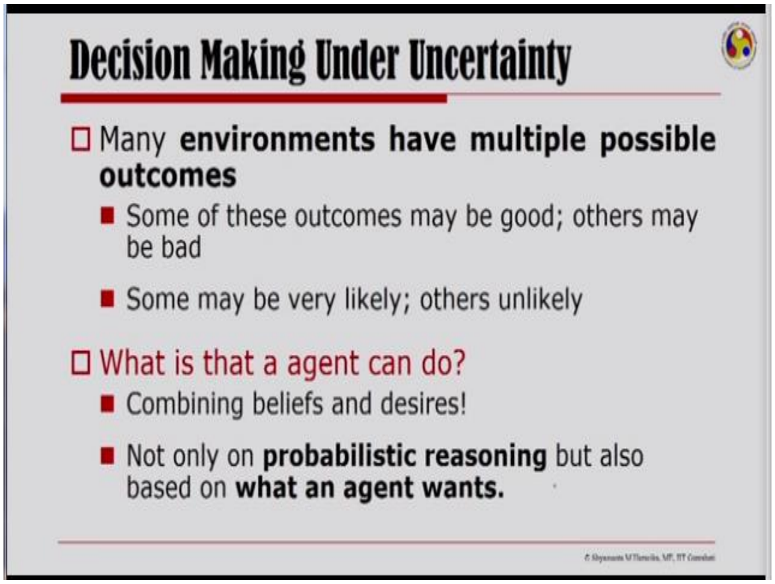

Now decision making under uncertainty is about taking a rational decision by looking at what we call the multiple possible outcomes. Many environments have multiple possible outcomes, some of these outcomes maybe good others maybe bad of these outcomes it maybe that many are very likely others are unlikely. Now under this situation what is that an agent can do. The answer lies in combining the believes and desires of the agent that is it not only uses probabilistic reasoning but also uses some form of it is desire or what it wants to do.

(Refer Slide Time: 03:59)

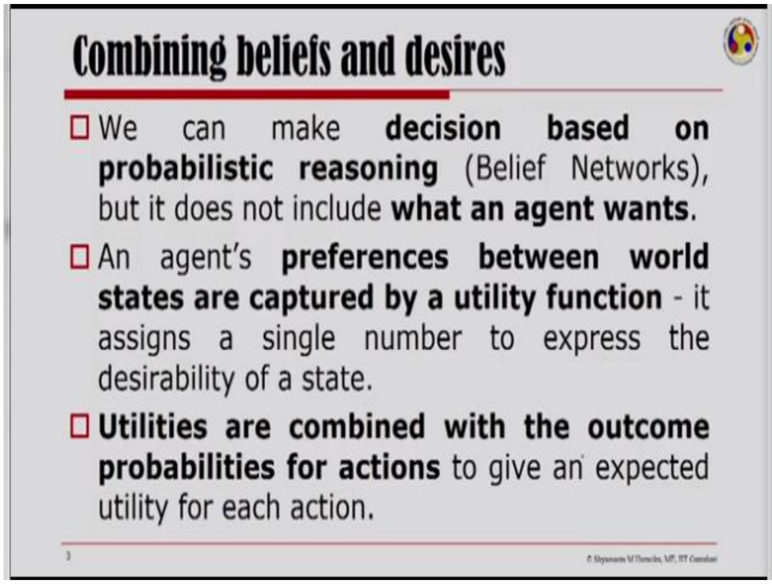

Combining believes and desires one can make decisions based on the belief network and some preferences between the world states which is captured by what is called a utility function. A utility function assigns a single number to express the desirability of a state. And utilities are combine with the outcome probabilities for actions to give an expected utility for each action. And this is what makes decision using just the probabilistic reasoning of belief networks different from decision making using decision networks.

Here we have preferences between states and these preferences are captured by the utility functions and these utilities are then combined with probabilities to arrive at what are called expected utility for each action.

(Refer Slide Time: 05:12)

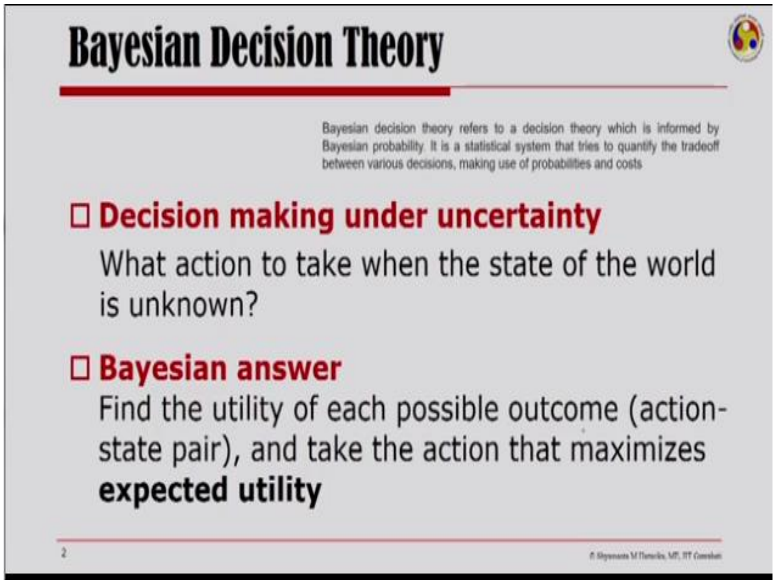

Bayesian decision theory refers to a decision theory which is informed by Bayesian probability, so if this is the statistical system which tries to quantify the tradeoff between various decisions making use of probabilities and costs. And when I say decision making under uncertainty we mean what action to take when the state of the world is unknown. The answer lies in the Bayesian decision theory that is we find the utility of each possible outcomes that is the actions state pair and take those actions that maximizes expected utility.

So even if it is not very clear from the state of the world immediately that what is the action that needs to taken right now unlike a logical world. In an uncertain domain we find out the utility of each possible outcome and take actions that maximizes the expected utility.

(Refer Slide Time: 06:24)

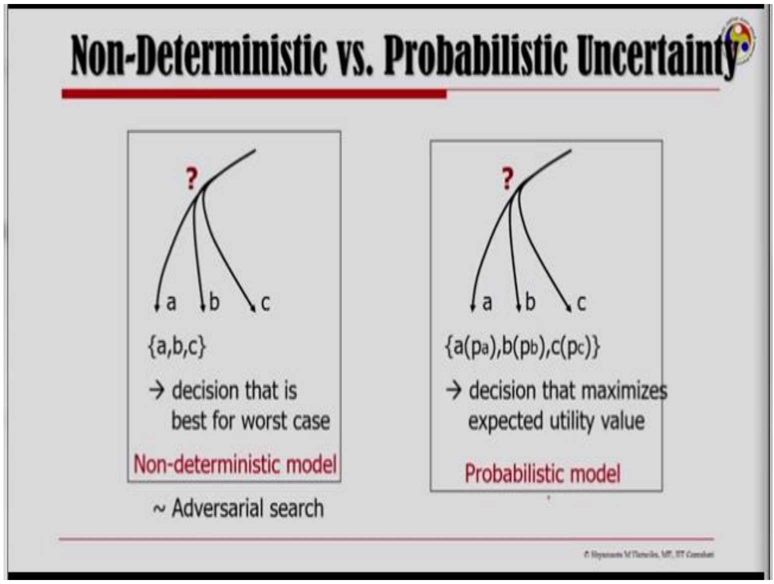

Here it would be nice to review and understand what is the difference between the nondeterministic model that we have been working through while we were talking of adversarial search. And the probabilistic uncertainty model that we are talking of now under the nondeterministic model such as that we have encountered while we were discussing searching game trees, the decision that was to be taken was the best for the worst case.

So I had assumed that there would be 3 optional ways a, b and c and which should I take was based on what would the worst for me given that my opponent tries to achieve the best for him. And that is what we looked at when we were looking for our decisions to be taken under a nondeterministic model in adversarial search. Here under probabilistic uncertainty we want to take decisions that maximizes the expected utility value.

And we would be driven by not only the fact that we have believes about the actions but then we will also be driven by how much of what is available to us, do we want to achieve. So it would be like the believes and the desires.

(Refer Slide Time: 08:13)

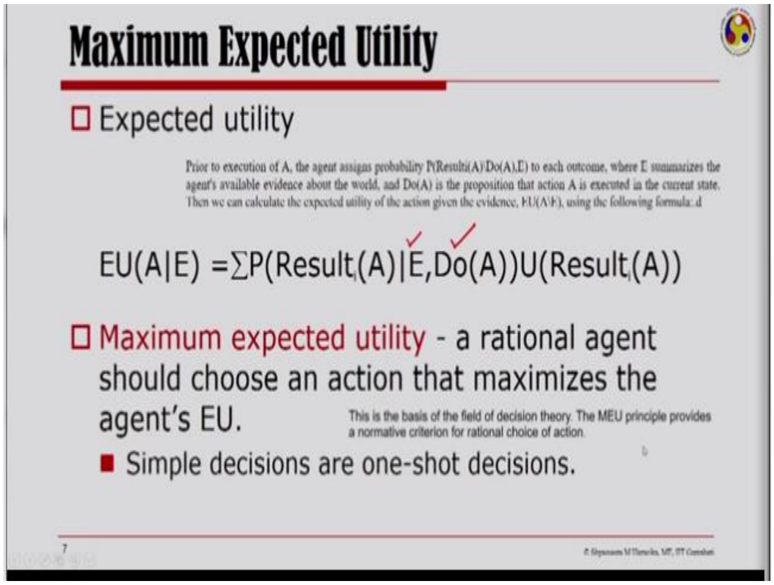

So let us now understand what we have been informally talking of the concept of expected utility if I have an action A prior to execution of the action the agent would assign probability of what would be the probability of the result of A on the action given that I have done the action A and the evidence that I have about the world. Now here the proposition Do A is about action A being executed in the current step and E summarizes the agents available evidence about the world.

This is the proposition that A is executed in the current step, E is the available evidence of the world. So expected utility would be about the probability of the result of A given that I have E and the action is done with the product of it is utility. So what is the utility of the result of A and that would be the expected utility, maximum expected utility is what a rational agent would choose and the action that you would take would be one that would maximize the agents expected utility.

This is the basis of the field of decision theory, the maximum expected utility principle thus provides a normative criteria for rational choice of actions. In our discussion today we will be more concern with just simple decisions are one-shot decisions.

(Refer Slide Time: 10:17)

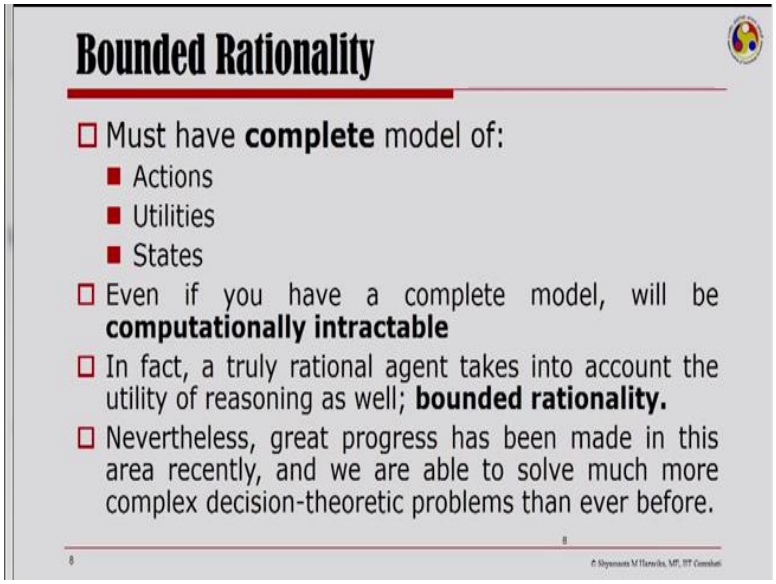

Now once that is in place and we know how a logical agent could derive actions based on maximum utility does that mean that the problem of taking decisions under uncertainty assault it is not so precisely because in order to take this decision the agent must have complete model of the actions utility and states. Now let us assume for a minute that the agent does have a complete model but even if it has finding out the maximum expected utility will be computationally intractable.

In fact a truly rational agent takes into account the utility of reasoning as well when he is trying to do the maximum expected utility and that is bonded rationality. Nevertheless there has been tremendous progress in this area and we are able to solve much more complex decision theoretic problems than ever before, we will here look at just a couple of simple ones.

(Refer Slide Time: 11:39)

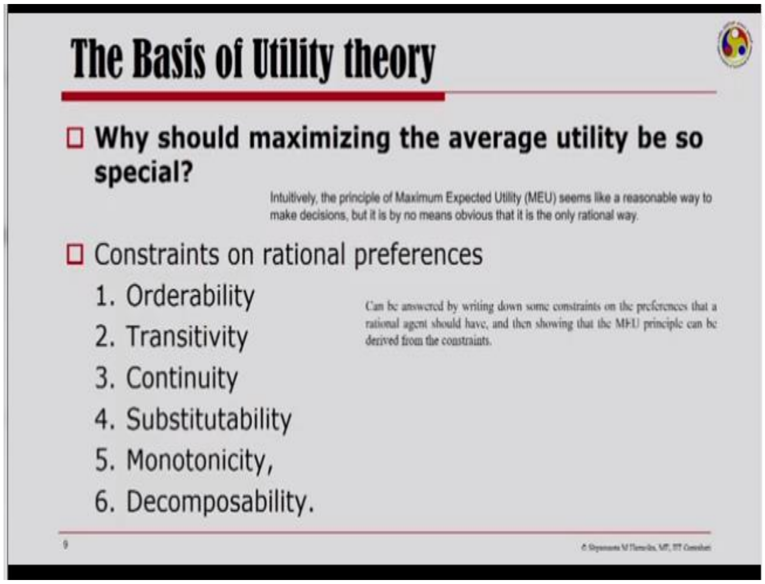

Before that we would try an understand what is the basis of the utility theory, now why should maximizing the average utility be so special. Intuitively, the principle of maximum expected utility seems like a reasonable way to make decisions but it is by no means obvious that it is the only rational way, you could take another metric and somehow look at utility to figure out the best utility for a given action.

This can be answered actually why maximizing the average utility is a reasonable way by writing down some constraints on the preferences that a rational agent should have. And then showing that the maximum expected utility principle can be derived from the constraints. So there are 6 constraints on rational preferences they are orderability, transitivity, continuity, substitutability, monotonicity and decomposability, let us look at each of them one by one.

(Refer Slide Time: 13:10)

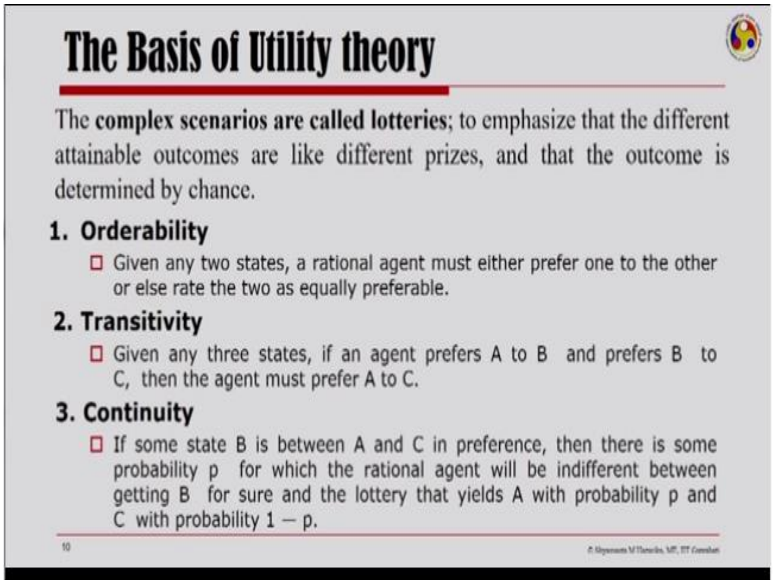

Before that let us introduce a concept called lotteries, in utility theory the complex scenarios are referred to as lotteries. This is to emphasize that the different attainable outcomes are like different prices and that the outcome is determined by chance. So all these constraints the first one the orderability states that if I have 2 states a rational agent must either prefer 1 to the other or else rate the 2 as equally preferable.

Transitivity talks of 3 states if an agent prefer A to B and prefers B to C then the agent must prefer A to C. Whereas on the other hand continuity is about some state B that is between A and C in preference. In that case there is some probability p for which the rational agent will be indifferent between getting B for sure and the lottery that yields A with probability p and C with probability 1 – p, so that is continuity.

So in a sense if you think of these is constraints, this constrain somehow trying to tells us how given certain scenarios in which states come out, a rational agent must make preferences between these states.

(Refer Slide Time: 14:55)

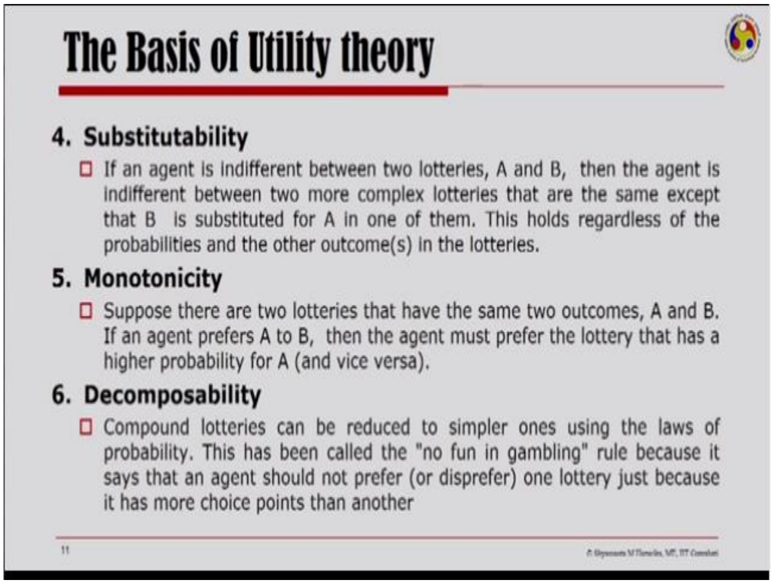

The fourth of the constraints is about substitutability, now if an agent is indifferent between 2 lotteries A and B. Then the agent is indifferent between 2 more complex lotteries that are the same except that B is substituted for A in one of them. Now this holds regardlessly of the probabilities and other outcomes in the lotteries. Then there is monotonicity suppose there are 2 lotteries that have the same 2 outcomes A and B.

Now if a agent prefers A to B then the agent must prefer the lottery that has a higher probability for A and vice versa. Finally the 6 constraint about decomposability is when compound lotteries are to be reduced the simpler ones using the law of probability. This has also been called the no fun in gambling rule because it says that the agent should not prefer or disprefer 1 lottery just because it has more choice points than another.

(Refer Slide Time: 16:14)

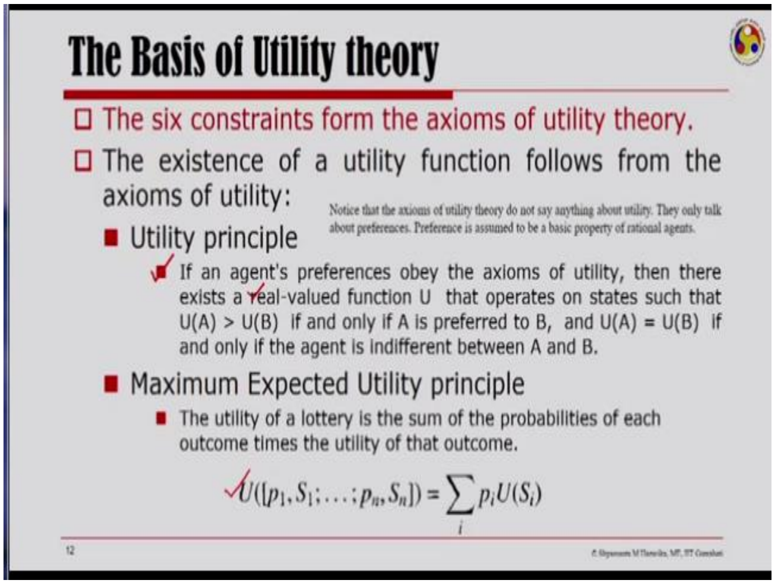

Having this 6 constraints we can then talk of the existence of a utility function, these constraints are called the axioms of utility theory and the existence of the utility function follows from the axioms of utility. The first among them is the utility principle now you should take a moment to realize that even if the utility function follows from the axioms of utility. The axioms of utility do not say anything about utility all they are talking of is about preferences.

So preference is assume to be a basic property of rational agents and that gives rise to the utility principle. Now here is the utility principle which states that if an agents preference says obey the axioms of utility then there exists a real valued function U that operates on states. Such that U of A is greater than U of B if and only if A is preferred to B and they are equal if and only if the agent is indifferent between A and B.

Now this principle of utility clearly brings in the agents preferences and talks of a real valued function utility. We thereafter have the maximum expected utility principle which is the utility of a lottery as the sum of probabilities of each outcome times the utility of that outcome which is one that we have looked at when we have introduced the definition for maximum expected utility, utility functions in a way map states to real numbers.

(Refer Slide Time: 18:35)

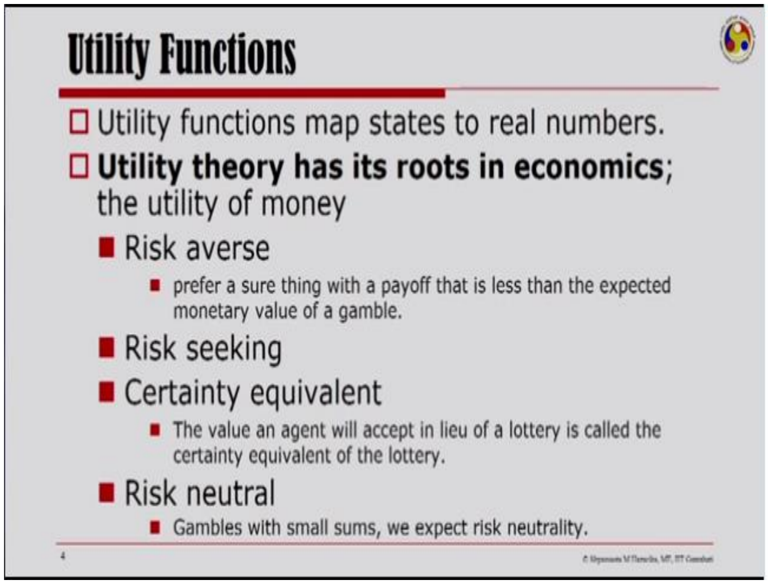

And one is to realize that the concept of utility or utility theory has it is roots in economics, the utility of money. There are a couple of concepts here that one needs to be clear about one is that of risk averse you prefer a sure thing with a payoff that is less than the expected monetary value of a gamble. Then there is risk seeking and certainty equivalent that is the value an agent will accept in lieu of a lottery is called the certainty equivalent of the lottery. Risk neutral, are those gambles which have small sums and we expect under such gambles risk neutrality.

(Refer Slide Time: 19:36)

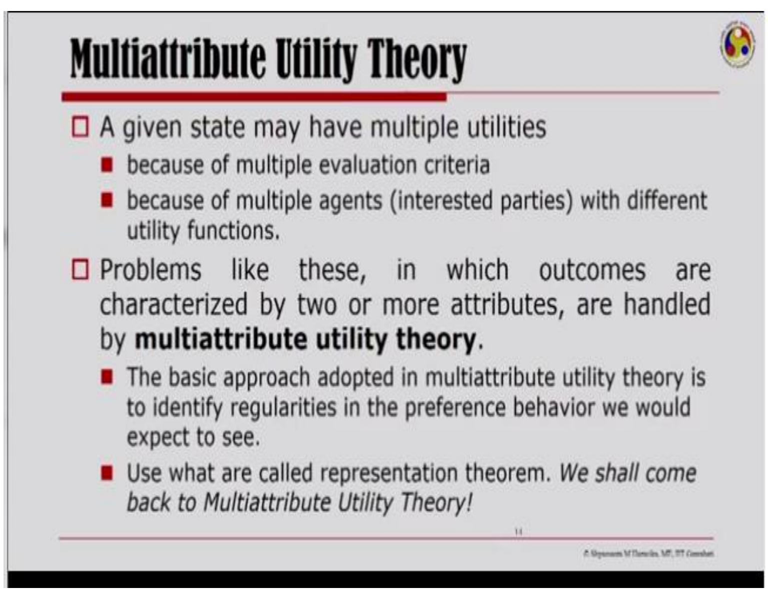

It is not that a state can have only one utility, there are situations in which a given state may have multiple utilities. This is because of multiple evaluation criteria or because of multiple interested parties with different utility functions working. Problems like this in which outcomes are characterized by 2 or more attributes are handled by what are called multiattribute utility theory. The basic approach adopted in multiattribute utility theory is to identify regularities in preference behavior we would expect to see.

And we use what are called representation theorems, now in today's lecture we would not cover multiattribute utility theory. During the course of our discussion on decision making we shall come back to multiattribute utility theory.

(Refer Slide Time: 20:42)

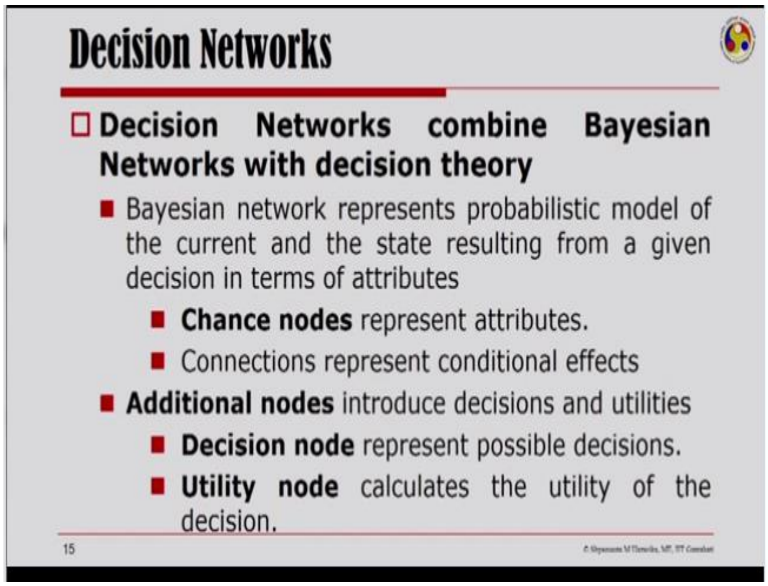

So let us introduce the concept of decision networks, as already pointed out decision networks combined Bayesian networks with decision theory. So recall that a Bayesian network represents probabilistic model of the current state resulting from a given decision in terms of attribute. We have chance nodes that represents the attributes and we have connections that represents conditional effects.

Together with that we now introduce additional nodes on decisions and utilities, decision nodes represent possible decisions and utility node calculates the utility of the decision.

(Refer Slide Time: 21:37)

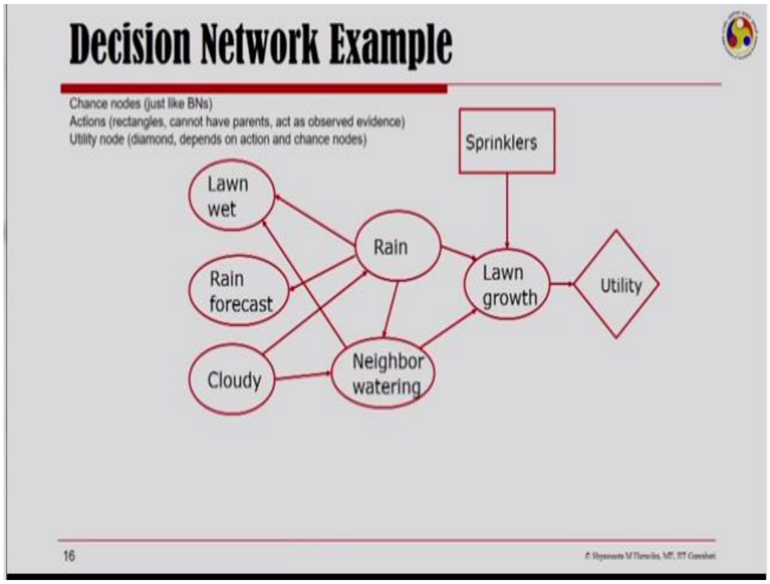

So here is a decision network example this is about whether you should put the sprinkler on to have your lawn water. When you know there is some chance of rain coming or you know that neighbor maybe watering your lawn. Now you want to find out given certain probabilities that they would be rain or certain probabilities that the neighbor would be watering your lawn, you want to take a decision of whether you should put the sprinkler is to watch out a lawn.

Now if you see here we have chance nodes just like the Bayesian network and then we have actions which are the rectangular things here, that is the decision node rectangles. Now these cannot have parents because they act as observed evidence and we have utility nodes which are diamonds. And one should take note that the utility nodes depends on both the action and the chance node.

(Refer Slide Time: 22:55)

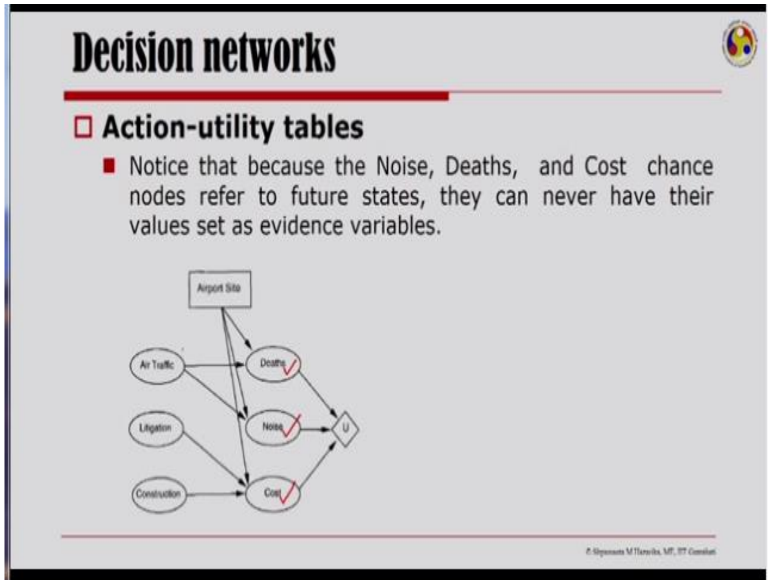

Now here is another decision network where you could see traffic litigation construction going on number of that is noise and cost in an airport side and you want to make a decision. Now here one thing you need to notice is that these nodes the depth, noise and the cost refer to future states they can never have their values at as evidence variables.

(Refer Slide Time: 23:32)

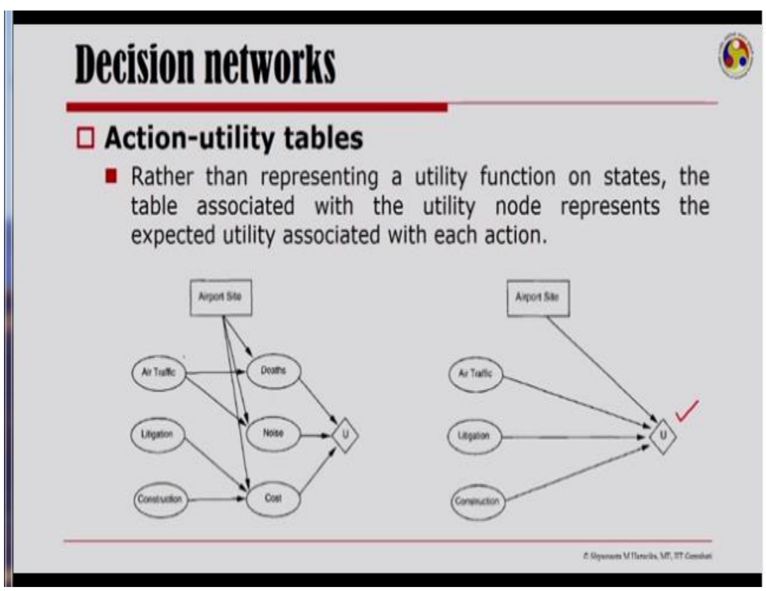

And therefore in the simplified version of decision networks such nodes are omitted, what you have on your right is a simplified version with those nodes omitted there. Now omission of an explicit description of the outcome of whether you would sit there in the airport site or not means that it is less flexible with respect to changes in circumstances. And rather than representing a utility function on states the table associated with the utility node here would represent the expected utility associated with each action and those are then called the action utility tables.

(Refer Slide Time: 24:27)

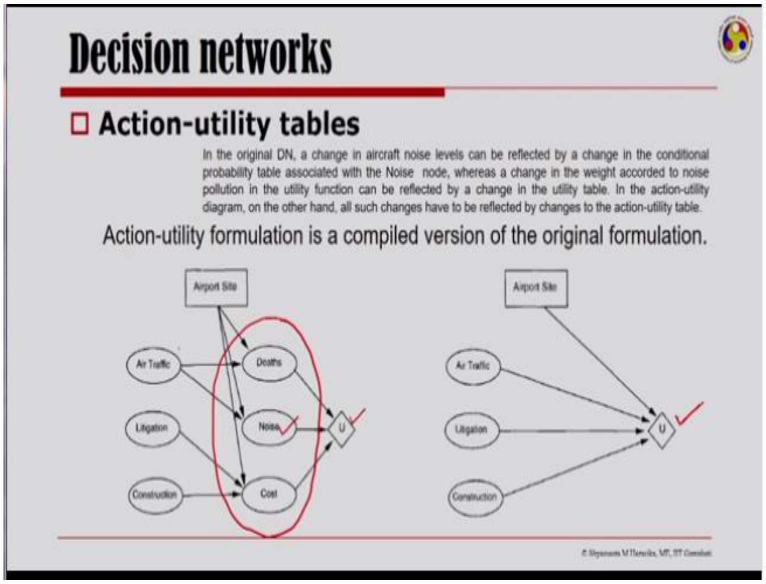

In the original decision network, you could see that any change in aircraft noise level can be reflected in the conditional probability table associated with this node, the noise node. Whereas a change in the weight accorded to noise pollution in the utility function can be reflected in the utility table here. Whereas in the action utility diagram this side all such changes has to be reflected in this action utility table because now I have already taken out this node from my decision network, so action utility formalism is actually a compiled version of the original formulation where you have somehow taken out these utilities reflected where somehow the probabilities associated with a noise node, the death node and the cost node here has been taken from here and compile together with the action utility table on the utility node.

(Refer Slide Time: 25:50)

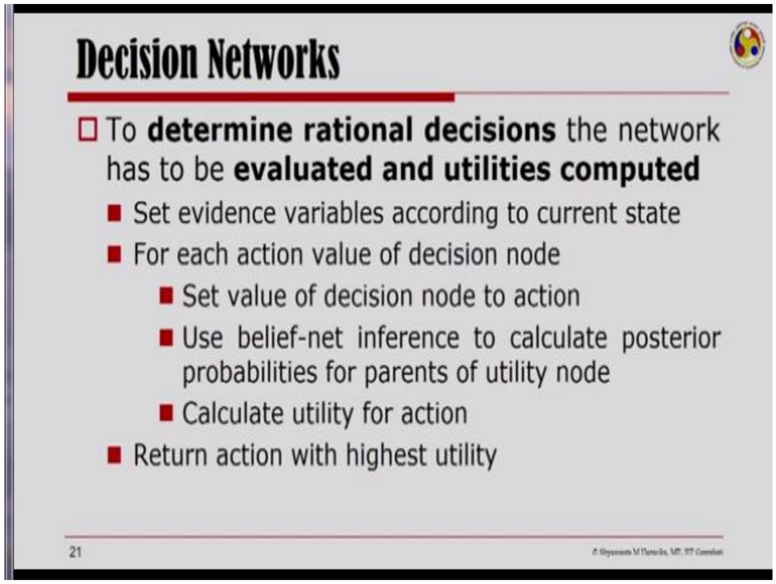

So decision networks determine rational decisions, to determine rational decisions in the decision network, the network has to be evaluated and the utilities computed. This is done by setting the evidence variable according to the current state. For each action value of the decision node you said value of decision node to action use belief net inference to calculate posterior probabilities for parents of utility node, calculate the utility for action, return action with the highest utility.

(Refer Slide Time: 26:35)

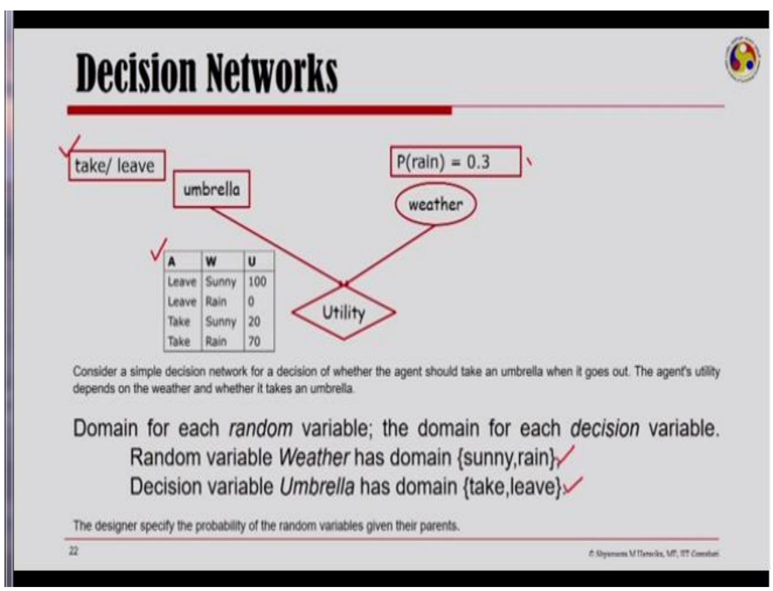

Now let us take an example and try to understand decision networks, here is an example of you having an umbrella and the weather being either sunny or rainy. And you want to take a decision on whether you would love to take or leave the umbrella. So you have a simple decision network here whether you should take an umbrella when you go out. Now utility depends on the weather and the utility here the table shows that for different actions of living it or taking it.

Depending on whether you have a sunny or a rainy weather outside you can have utilities assigned to that action. So for example you decide to leave your umbrella and it is sunny there is a utility of 100 and you leave your umbrella when you go out and it rains then the utility of the umbrella is actually 0. And similarly you can think of other situations where you take it and it is sunny it does have a utility when you take it and it is raining it does have a utility again.

Now if you look at this the domain for each random variable or the domain for each decision variable that we are talking of the random variable weather has in it is domain sunny and rain. The decision variable umbrella has take or leave and the designer of the decision network need to specify the probability of the random variables given their parents.

(Refer Slide Time: 28:36)

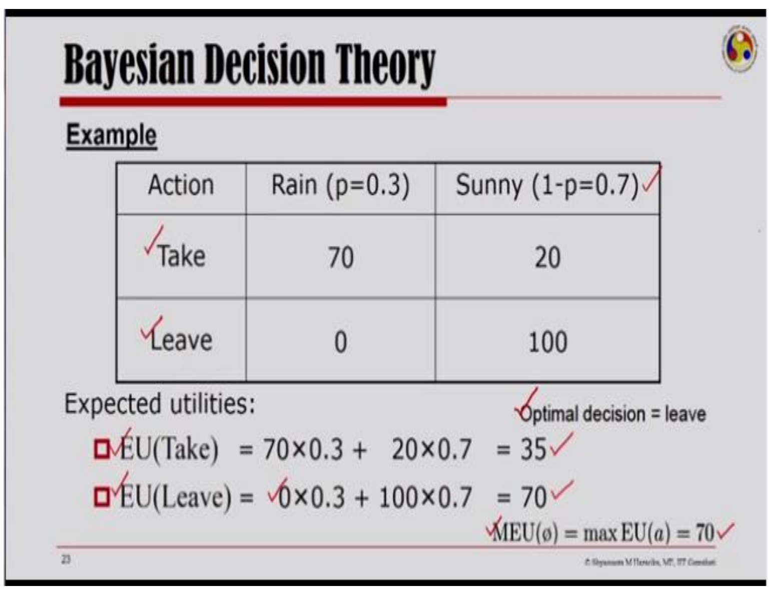

Under this scenario here is the action about whether you would take it or leave it and the different utilities originally specified, you would love to compute the expected utilities before taking a decision. So the expected utility of taking the umbrella is about the expected utility that you get into the probability that it may rain plus the expected utility that you get if it is sunny into the probability that it is sunny.

So you have an expected utility of 35, if you leave it while going out then you know if you leave it and it rains you have an utility of 0, so it is 0 into 0.3 plus you know if you leave it and it is sunny you have an utility of 100. And there is a probability that it would be sunny 0.7 so you have 0.7 into 100 which is 70. Now the maximum expected utility under no information other than the utility of taking it out or leaving it back has come out to be 70.

And this is for leaving it while going out, so the optimal decision is to leave the umbrella while going out.

(Refer Slide Time: 30:16)

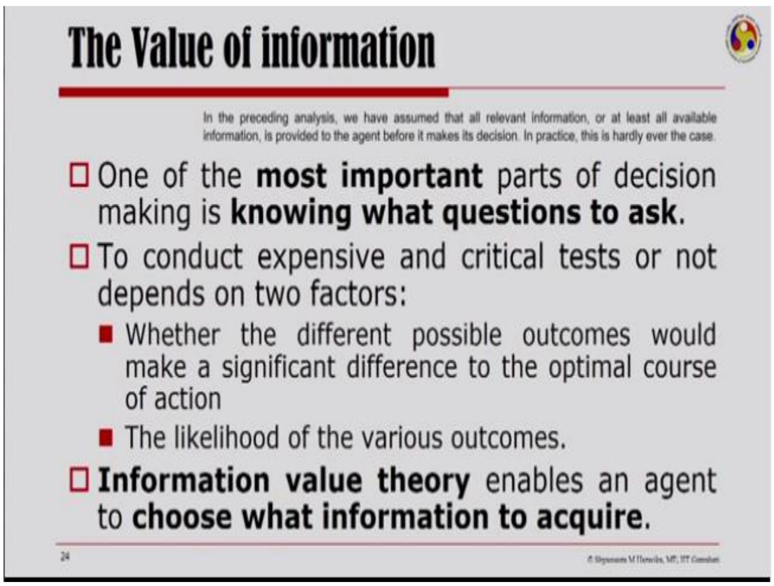

Now in the preceding analysis we have assume that all relevant information or at least all available information is provided to the agent before it makes its decision. In practice this is hardly ever the case, any decision that you make under uncertainty is based on number of questions that you keep on asking. So one of the most important parts of decision making is knowing what question to ask.

To conduct expensive and critical tests or not depend on 2 factors whether the different possible outcomes would make a significant difference to the optimal course of action. And what would be the likelihood of the various outcomes, information value theory enables an agent to choose what information to acquire.

(Refer Slide Time: 31:18)

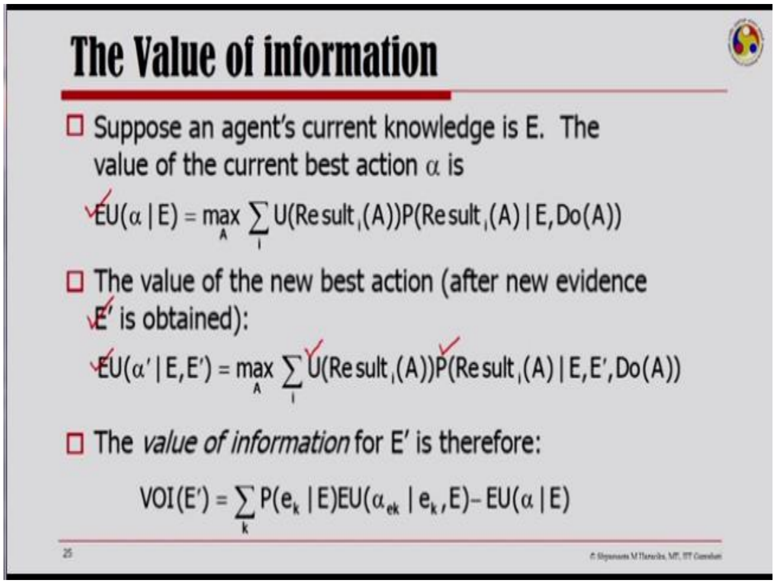

So suppose an agent's current knowledge is E, the value of the current best action alpha is given as expected utility alpha given E which is the maximum of the product of the utility and the probability of the result of A given the evidence E and doing the action A, this we have seen up till now. Now let us say we have new evidence that is obtain now let us call it E prime, so now I am interested in computing the expected utility which is no longer alpha, so which alpha prime give E and E prime.

Now the utility of action A remains the same but the probability changes, so the probability of result of A is now not dependant on E alone. But probability of A given E and E prime and of course doing the action A the difference that I obtain, here is called the value of information for E prime.

(Refer Slide Time: 32:40)

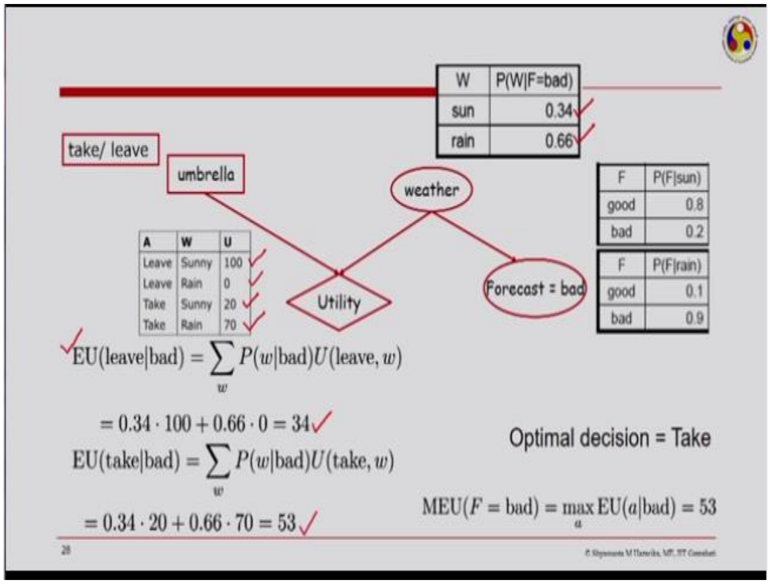

So let us look at our decision network the one that we have been discussing about whether you want to take the umbrella out while you are going out, given the utility as shown and the probability that it may rain being 0.3. Now let us say that I have additional information which is the forecast of the weather. So here the forecast clearly tells me that it would be the case that it maybe sunny or it would be the case that it maybe rainy.

But then I have also added to it the very fact that the forecast itself could be good or bad. So under this new scenario if I want to figure out how much of new information do I have then first I would love to figure out the probability of the weather given the forecast. So let us say I am interested in the probability of the weather given forecasting, so I select for evidences which is whether it is sunny or raining, so I try to figure out the probabilities of the weather.

And then I thereafter try to get to the new information given to me, now one is to realize that this was the original probabilities of the weather which was that probability that it would rain was 0.3, probability that it would be sunny was there for 1 – 0.3 which is 0.7. But with this new information that forecasting has come and under the scenario that forecast is bad I now know that the probability that it would be sunny is 0.2.

And the probability that it would be raining is 0.9, so that is the new evidence that I have. We join the probability of the weather and the probability of forecast being bad given the weather and then we normalize. So we have these probabilities of weather and forecast being bad and we then generate the normalize values which is that the probability of weather forecast being bad that is sunny is 0.34 and probability that it would rain forecast being bad is 0.66.

Given these forecasts I now have to replace the probability of the weather which is just 0.3 to this new probability forecastable which is of 0.34 and 0.66.

(Refer Slide Time: 35:54)

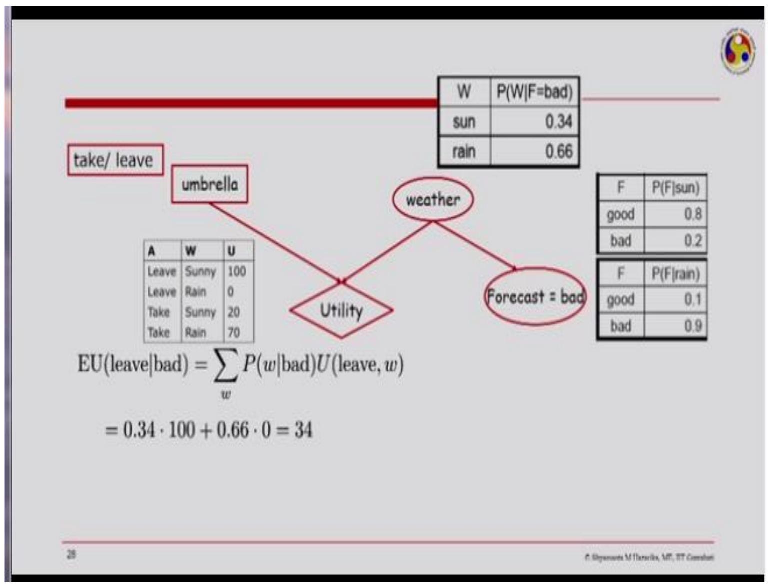

Given that table to me I would now able to compute the expected utility of leaving the umbrella when I have a bad forecast. And it is about leaving it is about 2 things one utility of 100 and the other utility of 0, so I multiply that with sunny 0.34 into 100 and raining 0.66 into o that is 34. That is the utility of leaving the umbrella when going out when the forecast is bad and the next one is about taking it even if I know the forecast is bad.

And then the maximum utility that I get is 0.34 into 20 the utility when I take it out when it is sunny and 0.66 into 70 which is when I take it out when it is rainy and I have the utility as 53. So under these scenarios the decision would be then to take the umbrella with you when you go out and the maximum expected utility is 53.

(Refer Slide Time: 37:37)

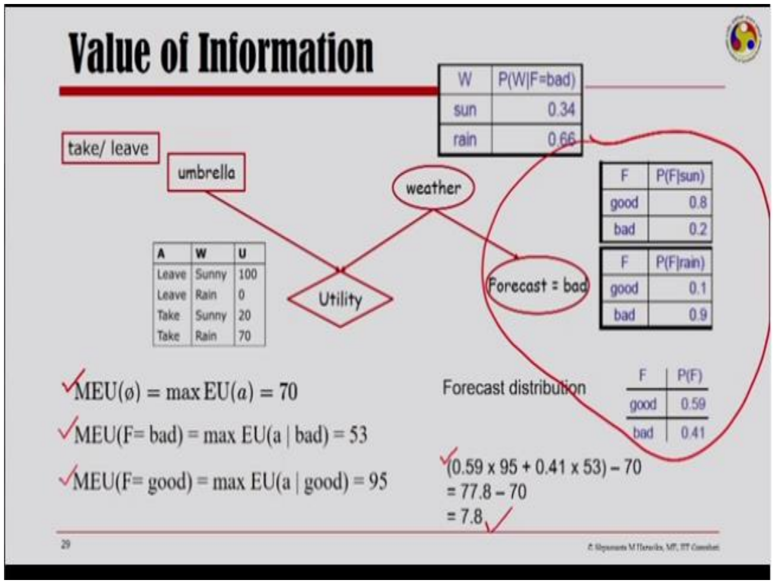

Now if you look for what is the value of information that you have gain in this then here is what you need to look at very carefully. Under no additional information the maximum expected utility that I compute it was 70 under an additional information about the forecast and under the scenario that the forecast is bad. Here it does not mean that the forecast is about a bad weather or a good weather, all it means here is the forecast itself is bad.

The reliability on the forecast is not something that you can be very sure about, so under such a scenario I have a maximum expected utility of 53. One I have not done here which you can try on your own is to look at this whole decision network and figure out what would be the maximum expected utility when the forecast is good and you should crosscheck that you would get a figure of 95.

So the maximum expected utility when the forecast is good is 95 and under such a scenario if you look at a forecast distribution let us say the good and bad is about 0.59 and 0.41. Then we could see that 95 into 0.59 that is what the good portion about 0.41 into 53 is the utility when the forecast is bad. So this is the total utility because of new information which is about the forecast here that has come in, so this portion of information that has come in.

So given that information on to this decision network we have a value of information which is 7.8.

(Refer Slide Time: 40:06)

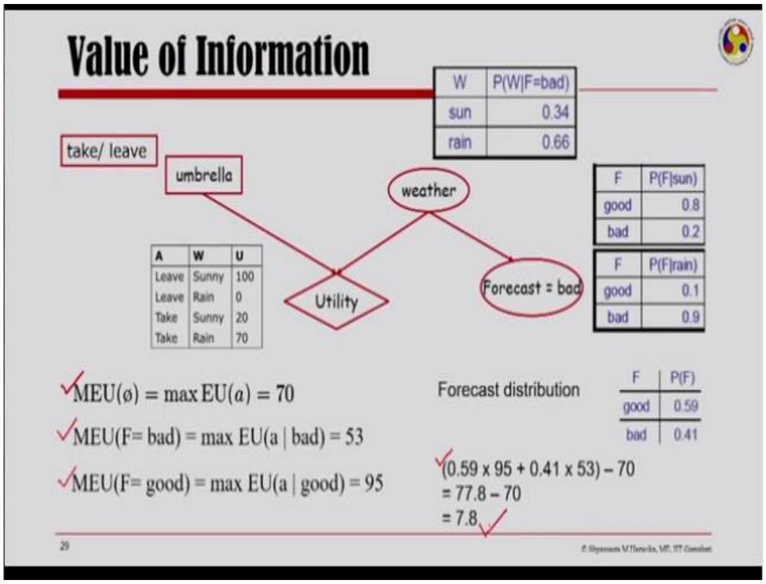

So what we can see from here is the following without any information we have a maximum expected utility of 70 with information which was bad information. Our maximum expected utility went down to 53 with additional information which was good information and our maximum expected utility went to 95 and the value of information is 7.8.

(Refer Slide Time: 40:43)

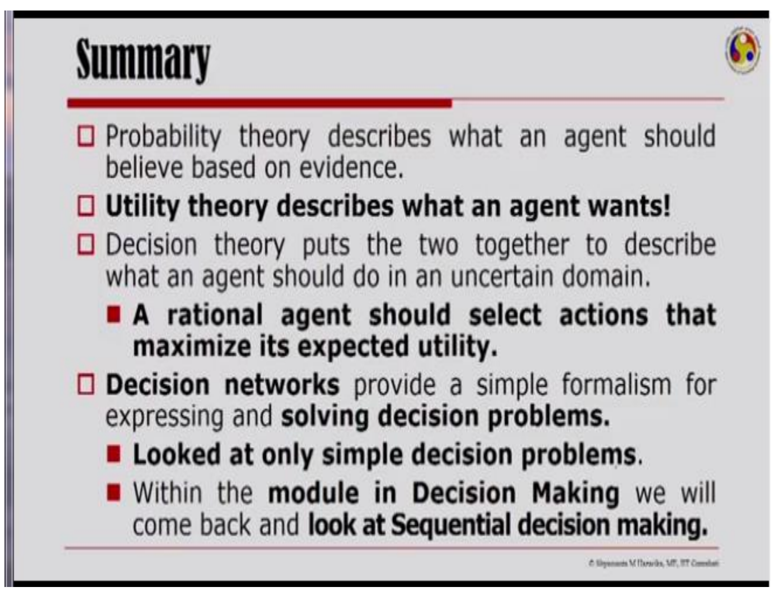

To summarize we have looked at probability theory, that describes what an agent should believe based on evidence. And today we have looked at utility theory, now utility theory actually is a description of what an agent wants. In terms of giving some amount of measure to the actions that it takes. Decision theory that we have been talking of puts the 2 together to describe what an agent should do in an uncertain domain.

And we have seen that a rational agent selects actions that maximizes it expected utility. So bringing together utility and probability the decision networks somehow ensure that the rational agent select actions that maximizes it expected utility. And we have looked at very simple example here of whether someone would like to take an umbrella while going out given the probabilities of rainy or sunny day and then what would be the utility of taking the umbrella with him or her.

Decision networks provides a formalism for expressing and solving such decision problems. As I was mentioning we had looked at only simple decision problems or what are called one shot decisions. While we will be doing decision making in one of these modules we will come back and look at a very interesting problem in decision making which is called sequential decision making, thank you very much.