6.825 Techniques in Artificial Intelligence

Decision Making under Uncertainty

- How to make one decision in the face of uncertainty

In the next two lectures, we'll look at the question of how to make decisions, to choose actions, when there's uncertainty about what their outcomes will be.

6.825 Techniques in Artificial Intelligence

Decision Making under Uncertainty

- How to make one decision in the face of uncertainty
- In a deterministic problem, making one decision is easy

When we were looking at deterministic, logical representations of world dynamics, it was easy to figure out how to make a single decision: you would just look at the outcome of each action and see which is best.

6.825 Techniques in Artificial Intelligence

Decision Making under Uncertainty

- How to make one decision in the face of uncertainty
- In a deterministic problem, making one decision is easy
- Planning is hard because we considered long sequences of actions

What made planning hard, was that we had to consider long sequences of actions, and we tried to be clever in order to avoid considering all of the exponentially many possible sequences of actions.

Decision Making under Uncertainty

- How to make one decision in the face of uncertainty
- In a deterministic problem, making one decision is easy
- Planning is hard because we considered long sequences of actions
- Given uncertainty, even making one decision is difficult

When there's substantial uncertainty in the world, we are not even sure how to make one decision. How do you weigh two possible actions when you're not sure what their results will be?

Decision Making under Uncertainty

- How to make one decision in the face of uncertainty
- In a deterministic problem, making one decision is easy
- Planning is hard because we considered long sequences of actions
- Given uncertainty, even making one decision is difficult

In this lecture, we'll look at the foundational assumptions of decision theory, and then see how to apply it to making single (or a very small number of) decisions. We'll see that this theory doesn't really describe human decision making, but it might still be a good basis for building intelligent agents.

A short survey

1. Which alternative would you prefer:
   A. A sure gain of $240
   B. A 25% chance of winning $1000 and a 75% chance of winning nothing

- 2. Which alternative would you prefer:
  - C. A sure loss of $750
  - D. A 75% chance of losing $1000 and a 25% chance of losing nothing

- 3. How much would you pay to play the following game: We flip a coin. If it comes up heads, I'll pay you $2. If it comes up tails, we'll flip again, and if it comes up heads, I'll pay you $4. If it comes up tails, we'll flip again, and if it comes up heads, I'll pay you $8. And so on, out to infinity.

Please stop and answer these questions. Don't try to think about the "right" answer. Just say what you would really prefer.

Decision Theory

- A calculus for decision-making under uncertainty

Decision theory is a calculus for decision-making under uncertainty. It's a little bit like the view we took of probability: it doesn't tell you what your basic preferences ought to be, but it does tell you what decisions to make in complex situations, based on your primitive preferences.

Decision Theory

- A calculus for decision-making under uncertainty
- Set of primitive outcomes

So, it starts by assuming that there is some set of primitive outcomes in the domain of interest. They could be winning or losing amounts of money, or having some disease, or passing a test, or having a car accident, or anything else that is appropriately viewed as a possible outcome in a domain.

Decision Theory

- A calculus for decision-making under uncertainty
- Set of primitive outcomes
- Preferences on primitive outcomes: $A \succ B$

Then, we assume that you, the user, have a set of preferences on primitive outcomes. We use this funny curvy greater-than sign to mean that you would prefer to have outcome A happen than outcome B.

Decision Theory

- A calculus for decision-making under uncertainty
- Set of primitive outcomes
- Preferences on primitive outcomes: $A \succ B$
- Subjective degrees of belief (probabilities)

We also assume that you have a set of subjective degrees of belief (which we'll call probabilities) about the likelihood of different outcomes actually happening in various situations.

Decision Theory

- A calculus for decision-making under uncertainty
- Set of primitive outcomes
- Preferences on primitive outcomes: A f B
- Subjective degrees of belief (probabilities)
- Lotteries: uncertain outcomes

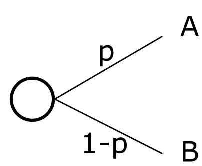

With probability p, outcome A occurs; with probability 1 - p; outcome B occurs.

Then, we'll model uncertain outcomes as lotteries. We'll draw a lottery like this, with a circle indicating a "chance" node, in which, with probability p, outcome A will happen, and with probability 1-p, outcome B will happen.

Axioms of Decision Theory

If you accept these conditions on your preferences, then decision theory should apply to you!

Decision theory is characterized by a set of six axioms. If your preferences (or your robot's preferences!) meet the requirements in the axioms, then decision theory will tell you how to make your decisions. If you disagree with the axioms, then you have to find another way of choosing actions.

Axioms of Decision Theory

If you accept these conditions on your preferences, then decision theory should apply to you!

- Orderability: $A \succ B$ or $B \succ A$ or $A \approx B$

The first axiom is orderability. It says that for every pair of primitive outcomes, either you prefer A to B, you prefer B to A, or you think A and B are equally preferable. Basically, this means that if I ask you about two different outcomes, you don't just say you have no idea which one you prefer.

Axioms of Decision Theory

If you accept these conditions on your preferences, then decision theory should apply to you!

- Orderability: $A \succ B$ or $B \succ A$ or $A \approx B$
- Transitivity: If $A \succ B$ and $B \succ C$ then $A \succ C$

Transitivity says that if you like A better than B, and you like B better than C, then you like A better than C.

Axioms of Decision Theory

If you accept these conditions on your preferences, then decision theory should apply to you!

- Orderability: $A \succ B$ or $B \succ A$ or $A \approx B$
- Transitivity: If $A \succ B$ and $B \succ C$ then $A \succ C$
- Continuity: If $A \succ B \succ C$ then there exists p such that $L_1 \approx L_2$

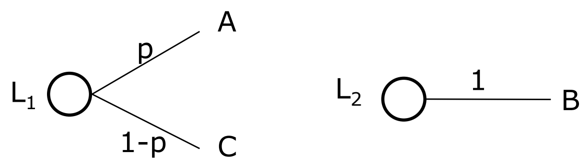

Continuity says that if you prefer A to B and you prefer B to C, then there's some probability that makes the following lotteries equivalently preferable for you. In the first lottery, you get your best outcome, A, with probability p, and your worst outcome, C, with probability 1-p. In the second lottery, you get your medium-good outcome, B, with certainty. So, the idea is that, by adjusting p, you can make these lotteries equivalently attractive, for any combination of A, B, and C.

More Axioms of Decision Theory

- Substitutability: If $A \succ B$, then $L_1 \succ L_2$

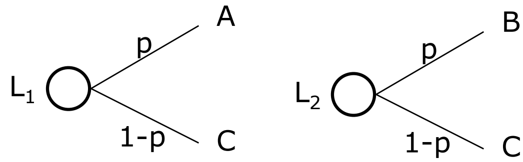

Substitutability says that if you prefer A to B, then given two lotteries that are exactly the same, except that one has A in a particular position and the other has B, you should prefer the lottery that contains A.

More Axioms of Decision Theory

- Substitutability: If A f B , then L_1 f L_2

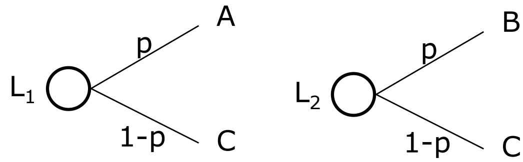

- Monotonicity: If A f B and p > q, then L_1 f L_2

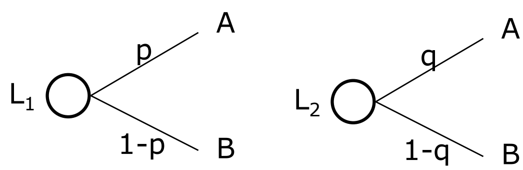

Monotonicity says that if you prefer A to B, and if p is greater than q, then you should prefer a lottery that gives A over B with higher probability.

Last Axiom of Decision Theory

- Decomposibility: $L_1 \approx L_2$

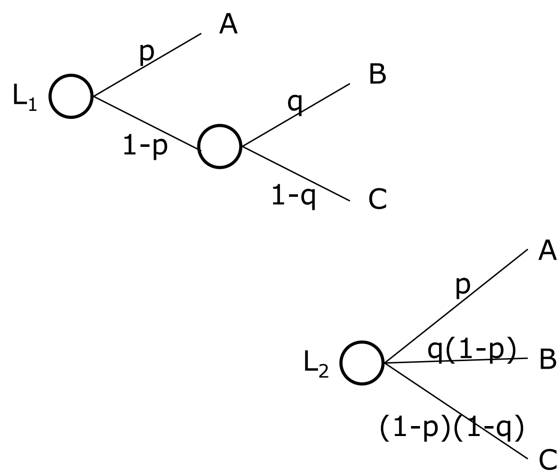

The last axiom of decision theory is decomposability. It, in some sense, defines compound lotteries. It says that a two-stage lottery, where in the first stage you get A with probability p, and in the second stage you get B withprobability q and C with probability 1-q is equivalent to a single-stage lottery with three possible outcomes: A with probability p, B with probability (1-p)*q, and C with probability (1-p)*(1-q).

Main Theorem

If preferences satisfy these six assumptions, then there exists U (a real valued function) such that:

- If A f B , then U(A) > U(B)
- If $A \approx B$, then U(A) = U(B)

So, if your preferences satisfy these six axioms, then there exists a real-valued function U such that if you prefer A to B, then U(A) is greater than U(B), and if you prefer A and B equally, then U(A) is equal to U(B). Basically, this says that all possible outcomes can be mapped onto a single utility scale, and we can work directly with utilities rather than collections of preferences.

Main Theorem

If preferences satisfy these six assumptions, then there exists U (a real valued function) such that:

- If $A \text{ f } B$, then U(A) > U(B)
- If $A \approx B$, then U(A) = U(B)

Utility of a lottery = expected utility of the outcomes

$$U(L) = p \cdot U(A) + (1 - p) \cdot U(B)$$

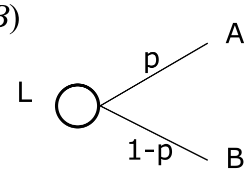

One direct consequence of this is that the utility of a lottery is the expected utility of the outcomes. So, the utility of our standard simple lottery L is p times the utility of A plus (1-p) times the utility of B. Once we know how to compute the utility of a simple lottery like this, we can also compute the utility of very complex lotteries.

Main Theorem

If preferences satisfy these six assumptions, then there exists U (a real valued function) such that:

- If $A \succ B$, then U(A) > U(B)
- If $A \approx B$, then U(A) = U(B)

Utility of a lottery = expected utility of the outcomes

$$U(L) = p \cdot U(A) + (1 - p) \cdot U(B)$$

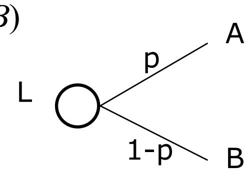

Interestingly enough, if we make no assumptions on the behavior of probabilities other than that they have to satisfy the properties described in these axioms, then we can show that probabilities actually have to satisfy the standard axioms of probability.

Survey Question 1

We're going to use the survey questions to motivate a discussion of utility functions and, in particular, the utility of money.

Survey Question 1

Which alternative would you prefer:

A. A sure gain of $240

B. A 25% chance of winning $1000 and a 75% chance of winning nothing

So, the first survey question was whether you'd prefer option A, a sure gain of $240, or option B, a 25% chance of winning $1000 and a 75% chance of winning nothing.

Survey Question 1

Which alternative would you prefer:

A. A sure gain of $240

B. A 25% chance of winning $1000 and a 75% chance of winning nothing

85% prefer option A to option B

When I polled one class of students, 85% preferred option A. This is consistent with results obtained in experiments published in the psychology literature.

Survey Question 1

Which alternative would you prefer:

A. A sure gain of $240

B. A 25% chance of winning $1000 and a 75% chance of winning nothing

85% prefer option A to option B

- U(B) = .25 U($1000) + .75 U($0)
- U(A) = U($240)
- U(A) > U(B)

So, what do we know about the utility function of a person who prefers A to B? We know that the utility of B is .25 times the utility of $1000 plus .75 times the utility of $0. (We're not going to make any particular assumptions about the utility scale; so we don't know, for example, that the utility of $0 is 0). We know that the utility of A is the utility of $240. And we know that the utility of A is greater than the utility of B.

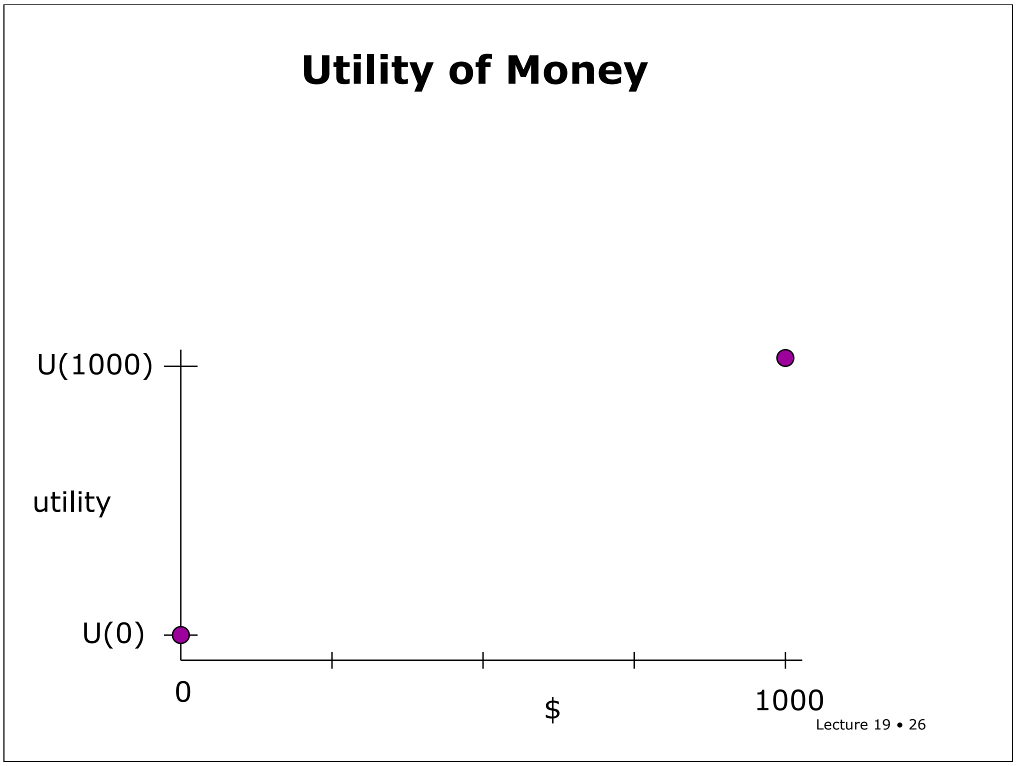

Let’s look in detail at how utility varies as a function of money. So, here’s a graph with money on the x axis and utility on the y axis. We’ve put in two points representing $0 and $1000, and their respective utilities.

Utility of Money

- U(B) = .25 U($1000) + .75 U($0)

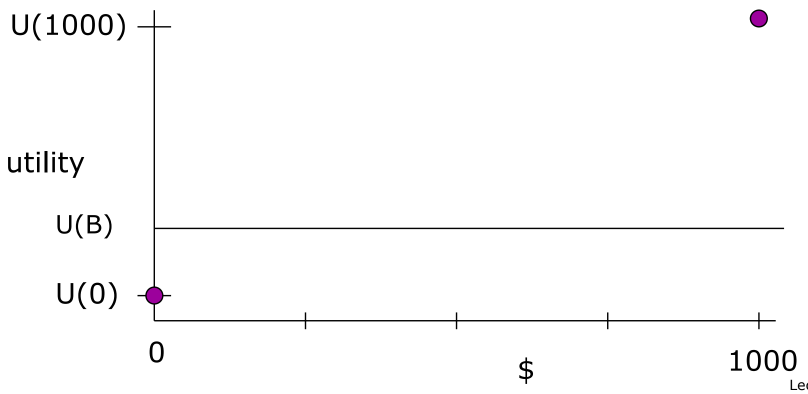

Now, let’s think about the utility of option B. It’s an amount that is 1/4 of the way up the y axis between the utility of $0 and the utility of $1000.

Utility of Money

- $U(B) = .25\ U(\$1000) + .75\ U(\$0)$
- $U(A) = U(\$240)$

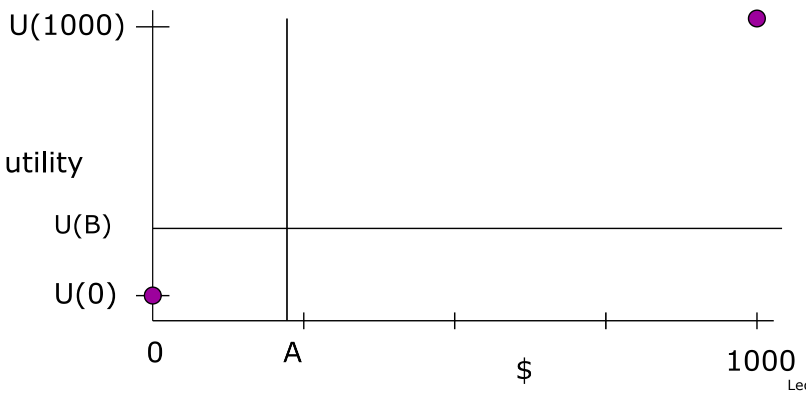

The utility of A is the utility of $240; we don't know exactly what its value is, but it will be plotted somewhere on this vertical line (with x coordinate 240).

Utility of Money

- U(B) = .25 U($1000) + .75 U($0)
- U(A) = U($240)
- U(A) > U(B)

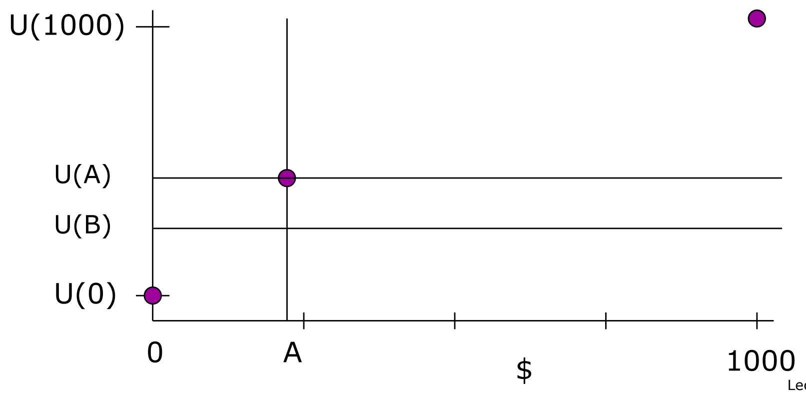

Now, since we know that the utility of A is greater than the utility of B, we just pick a y value for the utility of A that's above the utility of B, and that lets us add a point to the graph at coordinates A, U(A).

Utility of Money

- U(B) = .25 U($1000) + .75 U($0)
- U(A) = U($240)
- U(A) > U(B)

concave utility function
risk averse

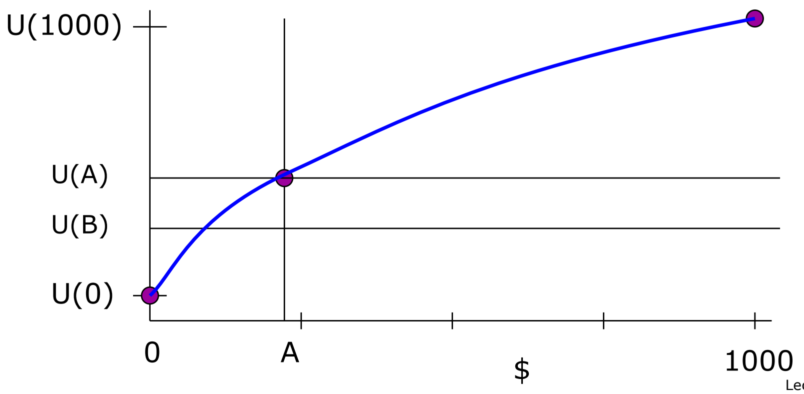

If we plot a curve through these three points, we see that the utility function is concave. This kind of a utility curve is frequently referred to as “risk averse”. This describes a person who would in general prefers a smaller amount of money for sure, rather than lotteries with a larger expected amount of money (but not expected utility). Most people are risk averse in this way.

Utility of Money

- U(B) = .25 U($1000) + .75 U($0)
- U(A) = U($240)
- U(A) > U(B)

concave utility function
risk averse

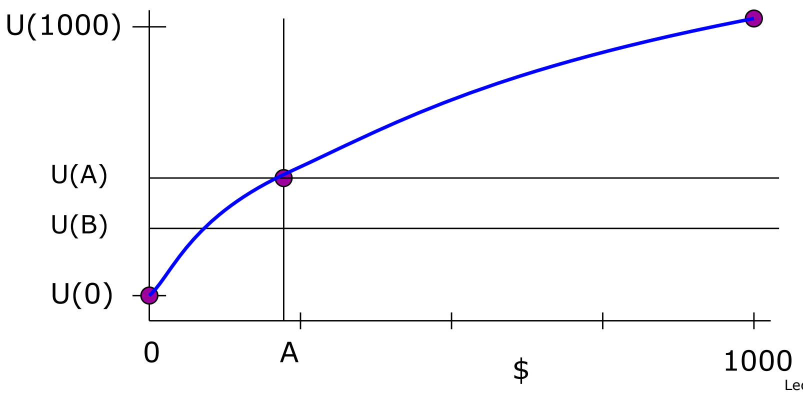

It’s important to remember that decision theory applies in any case. It’s not necessary to have a utility curve of any particular shape (in fact, you could even prefer less money to more!) in order for decision theory to apply to you. You just have to agree to the 6 axioms.

Risk neutrality

- U(B) = .25 U($1000) + .75 U($0) = U($250)
- U(A) = U($240)

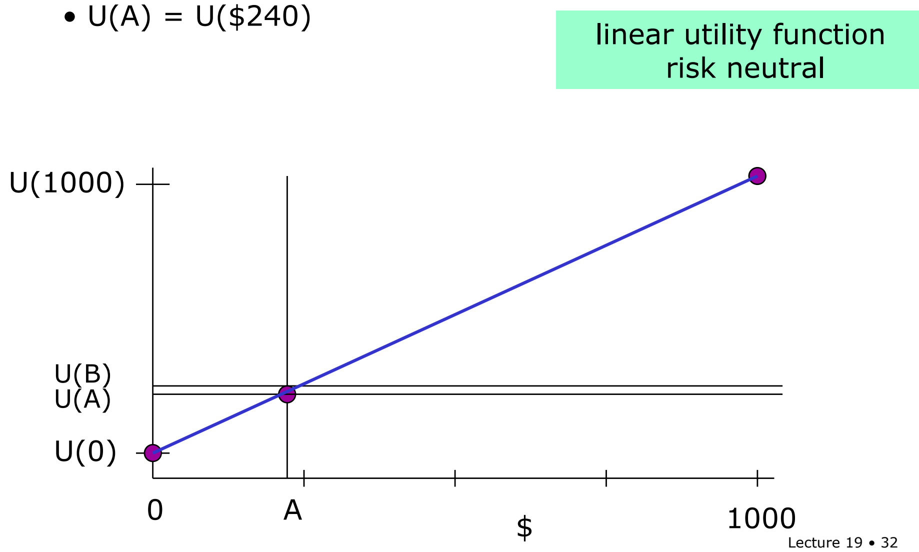

In contrast to the risk averse attitude described in the previous graph, another attitude is risk neutrality. If you are "risk neutral", then your utility function over money is linear. In that case, your expected utility for a lottery is exactly proportional to the expected amount of money you'll make. So, in our example, the utility of B would be exactly equal to the utility of $250. And so it would be greater than the utility of A, but not a lot.

Risk neutrality

- U(B) = .25 U($1000) + .75 U($0) = U($250)
- U(A) = U($240)

linear utility function
risk neutral

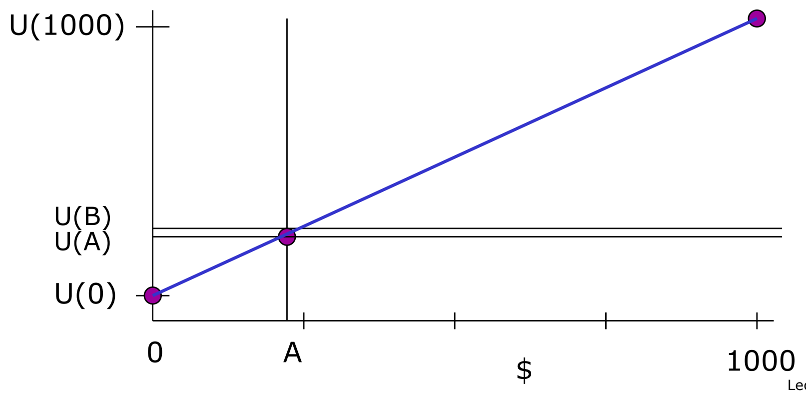

People who are risk neutral are often described as “expected value” decision makers.

Why Play the Lottery?

Do you think that decision theory would ever recommend to someone that they should play the lottery?

Why Play the Lottery?

Consider a lottery ticket:

- Expected payoff always less than price
- Is it ever consistent with utility theory to buy one?

In any real lottery, the expected amount of money you'll win is always less than the price. I once calculated the expected value of a $1 Massachusetts lottery ticket, and it was about 67 cents.

Why Play the Lottery?

Consider a lottery ticket:

- Expected payoff always less than price
- Is it ever consistent with utility theory to buy one?

It's kind of like preferring lottery B to A, below:

A. A sure gain of $260

B. A 25% chance of winning $1000 and a 75% chance of winning nothing

To stay consistent with our previous example, imagine that I offer you either a sure gain of $260, or a 25% chance of winning $1000 and a 75% chance of winning nothing.

Why Play the Lottery?

Consider a lottery ticket:

- Expected payoff always less than price
- Is it ever consistent with utility theory to buy one?

It's kind of like preferring lottery B to A, below:

A. A sure gain of $260
B. A 25% chance of winning $1000 and a 75% chance of winning nothing

Wanting to buy a lottery ticket is like preferring B to A. B is a smaller expected value ($250) than A, but it has the prospect of a high payoff with low probability (though in this example the payoff is a lot lower and the probability a lot higher than it is in a typical real lottery).

Why Play the Lottery?

Consider a lottery ticket:

- Expected payoff always less than price
- Is it ever consistent with utility theory to buy one?

It’s kind of like preferring lottery B to A, below:

A. A sure gain of $260
B. A 25% chance of winning $1000 and a 75% chance of winning nothing

In utility terms:

- U(B) = .25 U($1000) + .75 U($0)
- U(A) = U($260)
- U(A) < U(B)

We can describe the situation in terms of your utility function. The utility of B is .25 times the utility of $1000 plus .75 times the utility of $0. The utility of A is the utility of $260. And the utility of A is less than the utility of B.

Risk seeking

- U(B) = .25 U($1000) + .75 U($0)
- U(A) = U($260)
- U(A) < U(B)

convex utility function
risk seeking

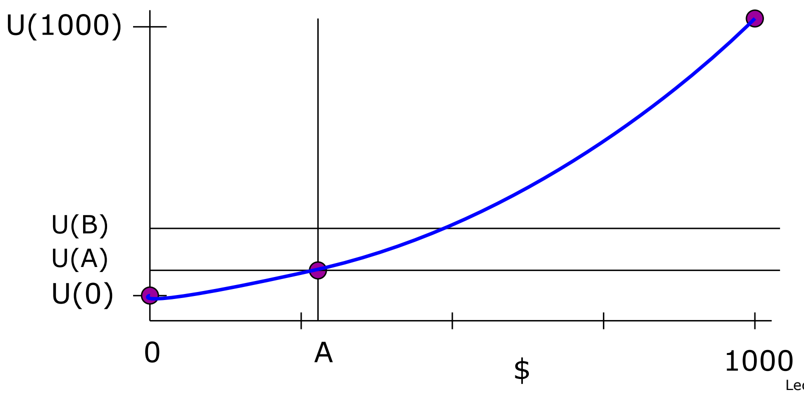

We can show the situation on a graph as before. The utility of B remains one quarter of the way between the utility of $0 and the utility of $1000. But now, A is slightly more than $250, and its utility is lower than that of B. We can see that this forces our utility function to be convex. In general, we'll prefer a somewhat riskier situation to a sure one. Such a preference curve is called “risk seeking”.

Risk seeking

- U(B) = .25 U($1000) + .75 U($0)
- U(A) = U($260)
- U(A) < U(B)

convex utility function
risk seeking

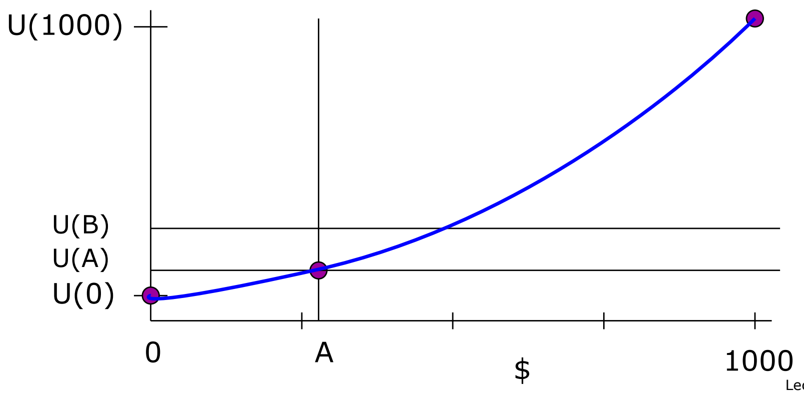

It's possible to argue that people play the lottery for a variety of reasons, including the excitement of the game, etc. But here we've argued that there are utility functions under which it's completely rational to play the lottery, for monetary concerns alone. You could think of this utility function applying, in particular, to people who are currently in very bad circumstances. For such a person, the prospect of winning $10,000 or more might be so dramatically better than their current circumstances, that even though they're almost certain to lose, it's worth $1 to them to have a chance at a great outcome.

Survey Question 2

Which alternative would you prefer:

- C. A sure loss of $750
- D. A 75% chance of losing $1000 and a 25% chance of losing nothing

Now, let's look at survey question 2. It asks whether you'd prefer a sure loss of $750 or a 75% chance of losing $1000 and a 25% chance of losing nothing.

Survey Question 2

Which alternative would you prefer:

- C. A sure loss of $750
- D. A 75% chance of losing $1000 and a 25% chance of losing nothing

91% prefer option D to option C

The vast majority of students in a class I polled preferred option D to option C. This is also consistent with results published in the psychology literature. People generally hate the idea of accepting a sure loss, and would rather take a risk in order to have a chance of losing nothing.

Survey Question 2

Which alternative would you prefer:

- C. A sure loss of $750
- D. A 75% chance of losing $1000 and a 25% chance of losing nothing

91% prefer option D to option C

- U(D) = .75 U(-$1000) + .25 U($0)
- U(C) = U(-$750)
- U(D) > U(C)

If you prefer option D to option C, we can characterize the utility function as follows. The utility of D is .75 times the utility of -$1000 plus .25 times the utility of $0. The utility of C is the utility of -$750. And the utility of D is greater than the utility of C.

Risk seeking in losses

- U(D) = .75 U(-$1000) + .25 U($0)
- U(C) = U(-$750)
- U(D) > U(C)

convex utility function
risk seeking

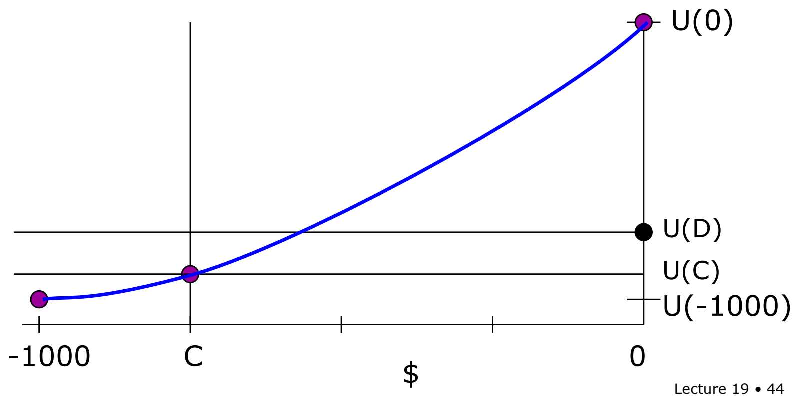

By the same arguments as last time, we can see that these preferences imply a convex utility function, which induces risk seeking behavior. It is generally found that people are risk seeking in the domain of losses. Or, at least, in the domain of small losses.

Risk seeking in losses

- U(D) = .75 U(-$1000) + .25 U($0)
- U(C) = U(-$750)
- U(D) > U(C)

convex utility function
risk seeking

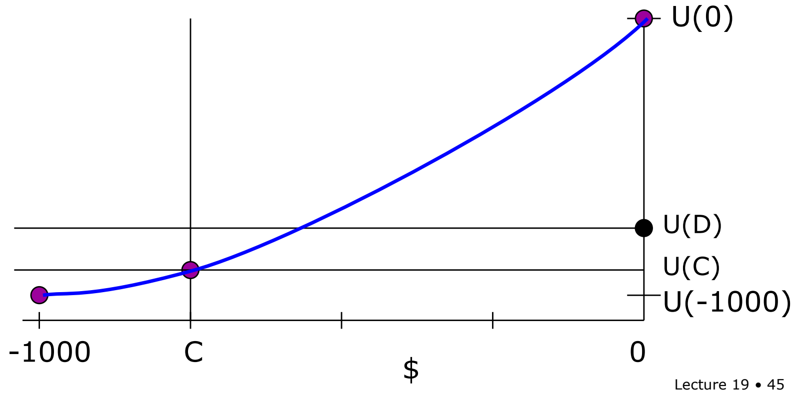

An interesting case to consider here is that of insurance. You can think of insurance as accepting a small guaranteed loss (the insurance premium) rather than accepting a lottery in which, with a very small chance, a terrible thing happens to you. So, perhaps, in the domain of large losses, people's utility functions tend to change curvature and become concave again.

Risk seeking in losses

- U(D) = .75 U(-$1000) + .25 U($0)
- U(C) = U(-$750)
- U(D) > U(C)

convex utility function
risk seeking

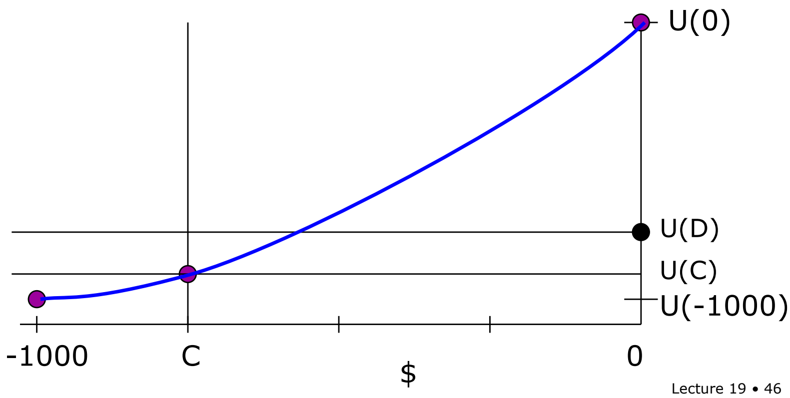

It's often possible to cause people to reverse their preferences just by changing the wording in a question. If I were selling insurance, I would have asked whether you'd rather accept the possibility of losing $1000, or pay an insurance premium of $750 that guarantees you'll never have such a bad loss. That change in wording might make option C more attractive than D.

Human irrationality

Most people prefer A in question 1 and D in question 2.

We've seen that decision theory can accommodate people who prefer option A to B, and people who prefer option D to C. In fact, most people prefer A and D. But we can show that it's not so good to both prefer A to B and D to C.

Human irrationality

Most people prefer A in question 1 and D in question 2.

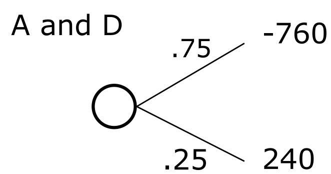

The easiest way to see it is to examine the total outcomes and probabilities of option A and D versus B and C.

Human irrationality

Most people prefer A in question 1 and D in question 2.

A and D

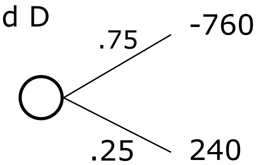

If you pick A and D, then with probability .75 you have a net loss of $760 and with probability .25 you have a net gain of $240.

Human irrationality

Most people prefer A in question 1 and D in question 2.

A and D

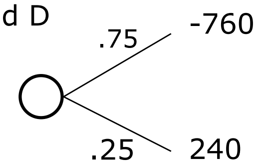

B and C

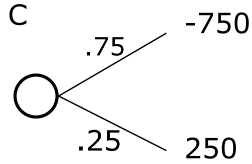

On the other hand, if you pick B and C, then with probability .75 you have a net loss of $750 and with probability .25, you have a net gain of $250.

Human irrationality

Most people prefer A in question 1 and D in question 2.

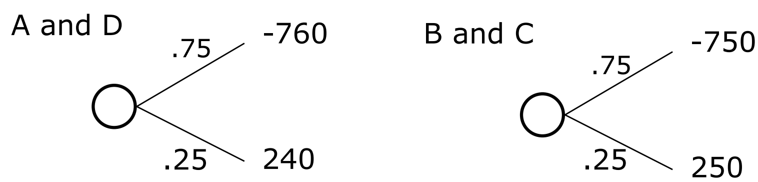

So, with B and C, it’s like getting an extra $10, no matter what happens. So it seems like it’s just irrational to prefer A and D.

Human irrationality

Most people prefer A in question 1 and D in question 2.

A and D

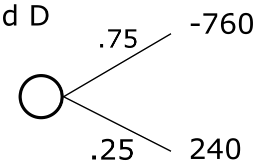

B and C

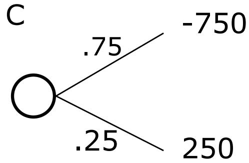

One student in my class made a convincing argument that it isn't irrational at all. That being given each single choice, it's okay to pick A in one and D in the other. But if you know you're going to be given **both** choices, then it would be unreasonable to pick both A and D.

St. Petersburg Paradox

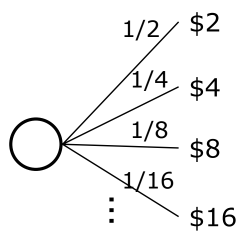

Just for fun, I asked how much you would pay to play this game, in which you get $2 with probability 1/2, $4 with probability 1/4, etc. This is called the St. Petersburg Paradox.

St. Petersburg Paradox

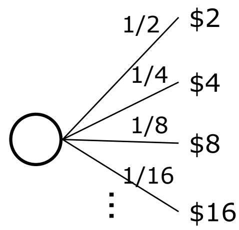

Expected value = 1 + 1 + 1 + … = ∞

It’s a paradox because the game has an expected dollar amount of infinity (1/2 * 2 is 1; 1/4 * 4 is 1; etc). However, most people don’t want to pay more than about $4 to play it. That’s a pretty big discrepancy.

St. Petersburg Paradox

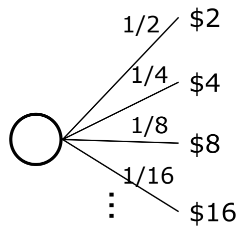

Expected value = 1 + 1 + 1 + … = ∞

It was this paradox that drove Bernoulli to think about concave, risk averse, utility curves, which you can use to show that although the game has an infinite expected dollar value, it will only have a finite expected utility for a risk averse person.

Buying a Used Car

Okay. Now we're going to switch gears a little bit, and see how utility theory might be used in a (somewhat) practical example.

Buying a Used Car

- Costs $1000
- Can sell it for $1100, $100 profit

Imagine that you have the opportunity to buy a used car for $1000. You think you can repair it and sell it for $1100, making a $100 profit.

Buying a Used Car

- Costs $1000
- Can sell it for $1100, $100 profit
- Every car is a lemon or a peach
- 20% are lemons
- Costs $40 to repair a peach, $200 to repair a lemon

You have some uncertainty, though, about the state of the car. It might be a lemon (a fundamentally bad car) or a peach (a good one). You think that 20% of cars are lemons. It costs $40 to repair a peach, and $200 to repair a lemon.

Buying a Used Car

- Costs $1000
- Can sell it for $1100, $100 profit
- Every car is a lemon or a peach
- 20% are lemons
- Costs $40 to repair a peach, $200 to repair a lemon
- Risk neutral

Let's further assume that, for this whole example, that you're risk neutral, so that the utility of $100 is 100. This isn't at all necessary to the example, but it will simplify our discussion and notation.

Buying a Used Car

- Costs $1000
- Can sell it for $1100, $100 profit
- Every car is a lemon or a peach
- 20% are lemons
- Costs $40 to repair a peach, $200 to repair a lemon
- Risk neutral

decision
tree

We can describe this decision problem using a decision tree. Note that we also talk about decision trees in supervised learning; these decision trees are almost completely different from those other ones; don't confuse them, despite the same name!

Buying a Used Car

- Costs $1000
- Can sell it for $1100, $100 profit
- Every car is a lemon or a peach
- 20% are lemons
- Costs $40 to repair a peach, $200 to repair a lemon
- Risk neutral

decision tree

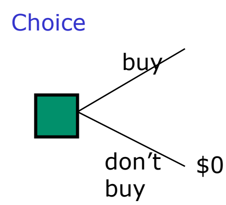

We start the decision tree with a “choice” node, shown as a green square: we have the choice of either buying the car or not buying it. If we don’t buy it, then the outcome is $0.

Buying a Used Car

- Costs $1000
- Can sell it for $1100, $100 profit
- Every car is a lemon or a peach
- 20% are lemons
- Costs $40 to repair a peach, $200 to repair a lemon
- Risk neutral

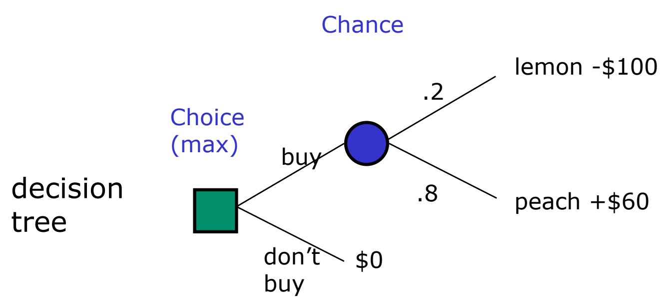

If we do buy the car, then the outcome is a lottery. We'll represent the lottery as a "chance" node, shown as a blue circle. With probability 0.2, the car will be a lemon and we'll have the outcome of losing $100 ($100 in profit minus $200 for repairs). With probability 0.8, the car will be a peach, and we'll have the outcome of gaining $60 ($100 in profit minus $40 for repairs).

Buying a Used Car

- Costs $1000
- Can sell it for $1100, $100 profit
- Every car is a lemon or a peach
- 20% are lemons
- Costs $40 to repair a peach, $200 to repair a lemon
- Risk neutral

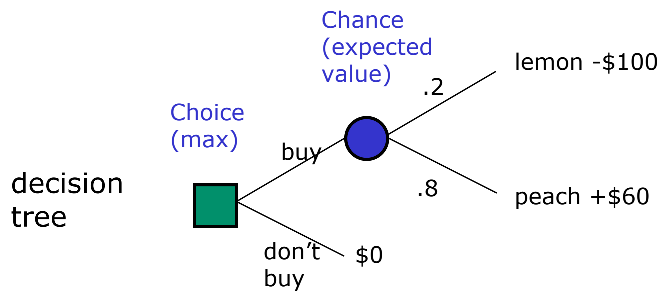

Now, we can use this tree to figure out what to do. We will evaluate it; that is, assign a value to each node. Starting from the leaves and working back toward the root, we compute a value for each node. At chance nodes, we compute the expected value for each leaf (it's nature who is making the choice here, so we just have to take the average outcome). At choice nodes, we have complete control, so the value is the maximum of the values of all the leaves.

Buying a Used Car

- Costs $1000
- Can sell it for $1100, $100 profit
- Every car is a lemon or a peach
- 20% are lemons
- Costs $40 to repair a peach, $200 to repair a lemon
- Risk neutral

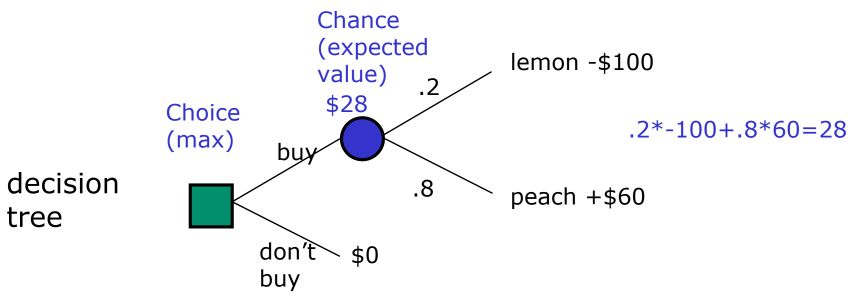

So, we start by evaluating the chance node. The expected value of this lottery is $28, so we assign that value to the chance node.

Buying a Used Car

- Costs $1000
- Can sell it for $1100, $100 profit
- Every car is a lemon or a peach
- 20% are lemons
- Costs $40 to repair a peach, $200 to repair a lemon
- Risk neutral

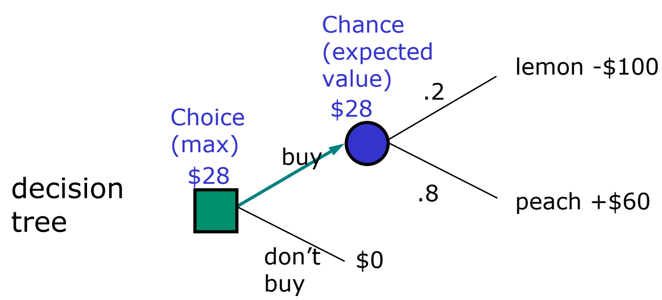

Now, we evaluate the choice node. We'll make $28 if we buy the car and nothing if we don't. So, we choose to buy the car (which we show by putting an arrow down that branch), and we assign value 28 to the choice node.

Expected Value of Perfect Info

How much should you pay for information of what type of car it is (before you buy it)?

Now, while we're sitting at the used car lot, thinking about whether to buy this car, an unhappy employee comes out and offers to sell us perfect information about whether the car is really a lemon or a peach. He's been working on the car and he knows for sure which it is. So, the question is, what is the maximum amount of money that we should be willing to pay him for this information?

Expected Value of Perfect Info

How much should you pay for information of what type of car it is (before you buy it)?

- Reverse the order of the chance and action nodes. The chance node represents uncertainty on what the information will be, once we get it.

We'll draw a decision tree to help us figure this out. We need to draw the nodes in a different order from left to right. In this scenario, the idea is that we first get the information about whether the car is a lemon or a peach, and then we get to make our decision (and the important point is that it can possibly be different in the two different cases).

Expected Value of Perfect Info

How much should you pay for information of what type of car it is (before you buy it)?

- Reverse the order of the chance and action nodes. The chance node represents uncertainty on what the information will be, once we get it.

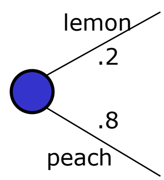

So, we start with a chance node, to describe the chance that the car is a lemon or a peach.

Expected Value of Perfect Info

How much should you pay for information of what type of car it is (before you buy it)?

- Reverse the order of the chance and action nodes. The chance node represents uncertainty on what the information will be, once we get it.

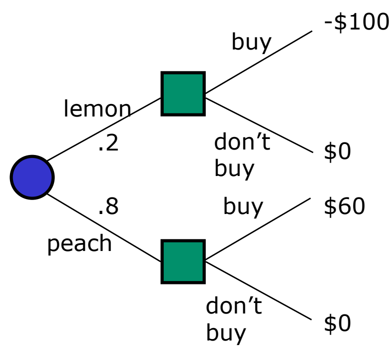

Now, for each of those pieces of information, we include a choice node about whether or not to buy the car.

Expected Value of Perfect Info

How much should you pay for information of what type of car it is (before you buy it)?

- Reverse the order of the chance and action nodes. The chance node represents uncertainty on what the information will be, once we get it.
- C is cost you have to pay for information

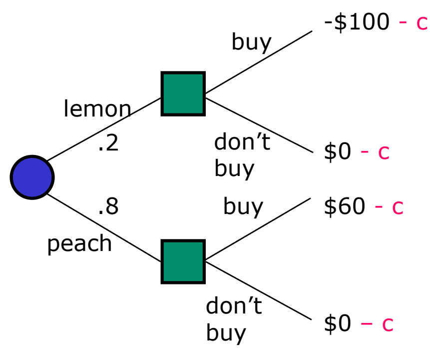

Let's let c be the amount of money we have to pay for this information. When we calculate the outcomes for the leaves, we have to subtract c from the outcomes we calculated before (-100 for buying a lemon, 0 for not buying anything, and +60 for buying a peach).

Expected Value of Perfect Info

How much should you pay for information of what type of car it is (before you take it to be repaired)?

- Reverse the order of the chance and action nodes. The chance node represents uncertainty on what the information will be, once we get it.
- C is cost you have to pay for information

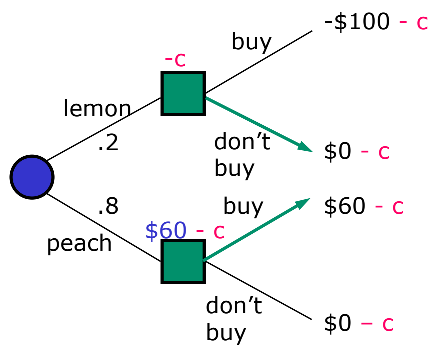

We evaluate the tree as before, working back from the leaves toward the root. We first hit these two choice nodes. In the top choice node, it's clearly better not to buy the car, for a total loss of c dollars. In the bottom choice node, it's better to buy the car, for a net of 60 – c dollars.

Expected Value of Perfect Info

How much should you pay for information of what type of car it is (before you take it to be repaired)?

- Reverse the order of the chance and action nodes. The chance node represents uncertainty on what the information will be, once we get it.
- C is cost you have to pay for information

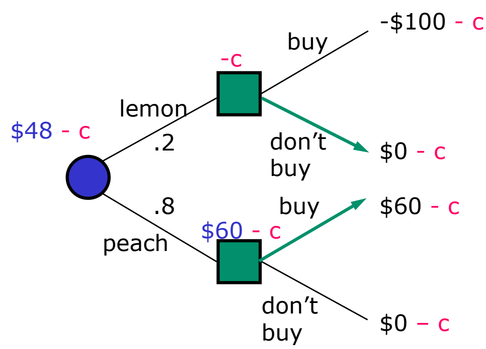

Now, we arrive at the chance node, where we take the expected value of the leaves. In this case, we get 48 – c dollars.

Expected Value of Perfect Info

How much should you pay for information of what type of car it is (before you take it to be repaired)?

- Reverse the order of the chance and action nodes. The chance node represents uncertainty on what the information will be, once we get it.
- C is cost you have to pay for information

c = $20 [EVPI]

ties the expected value with no information

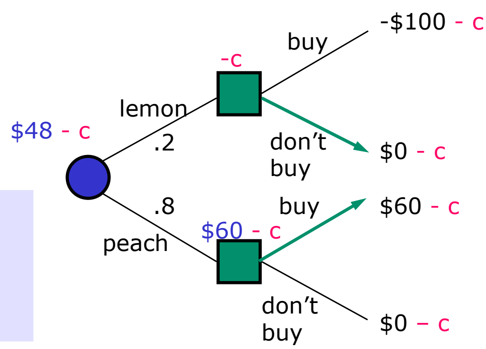

Since the original car-buying deal was worth $28 in the case without any extra information, then we shouldn't pay any more than $20 for perfect information. If we pay that much, then this deal is equivalent to the original one. If we pay less than $20, then this deal is better. We'll say that $20 is the expected value of perfect information.

Guarantee?

- Costs $60
- Covers 50% of repair costs
- If repairs > $100, covers them all

Now let's consider a different scenario. Nobody comes out to offer us information. Instead, a sleazy guy comes out of the office and says he has a special deal, just for us. He'll sell us a guarantee on this car for $60. It will cover %50 of all repair costs, if the repairs cost less than $100. But if the repairs are more than $100, it will cover them all.

Guarantee?

- Costs $60
- Covers 50% of repair costs
- If repairs > $100, covers them all

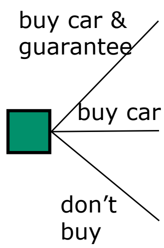

So, should we buy the guarantee? Let's make a decision tree. Now we have three choices. We can buy both the car and the guarantee, we can buy the car only, or we can buy nothing (it's completely unreasonable to buy just the guarantee!).

Guarantee?

- Costs $60
- Covers 50% of repair costs
- If repairs > $100, covers them all

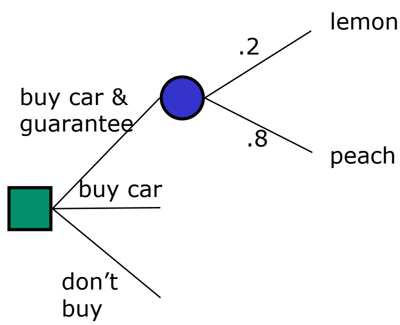

If we buy the car and the guarantee, then we have to consider what the outcomes will be depending on whether or not it's a lemon.

Guarantee?

- Costs $60
- Covers 50% of repair costs
- If repairs > $100, covers them all

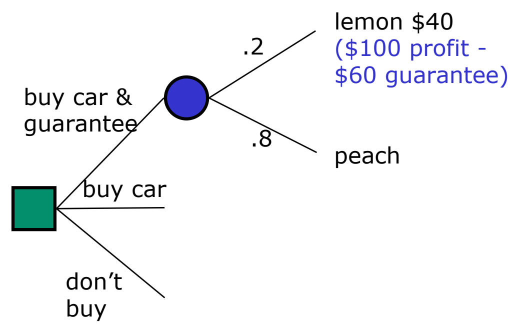

If it's a lemon, we make $40. That's because the guarantee covers all the repairs. So we make $100 profit on the sale of the car, but we have to pay $60 for the guarantee.

Guarantee?

- Costs $60
- Covers 50% of repair costs
- If repairs > $100, covers them all

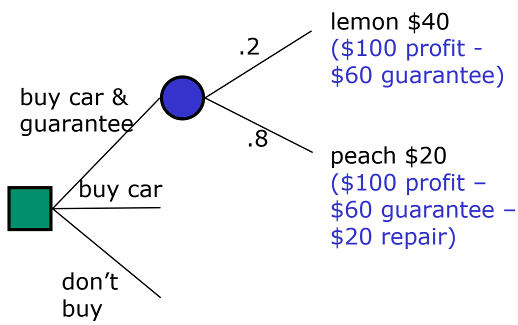

If it's a peach, we make $20. We start with $100 in profit, but we have to pay $60 for the guarantee and $20 for repairs (at least we only have to pay for half of the repairs!).

Guarantee?

- Costs $60
- Covers 50% of repair costs
- If repairs > $100, covers them all

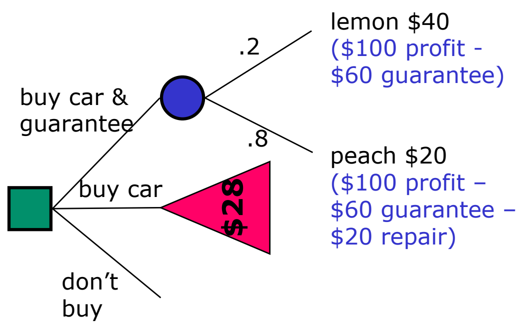

Now, if we decide to buy the car, we can substitute in the chance node and outcomes that we already calculated for buying the car in the first example. I put in a red triangle to stand for the summarization of some other decision tree, with a result of $28.

Guarantee?

- Costs $60
- Covers 50% of repair costs
- If repairs > $100, covers them all

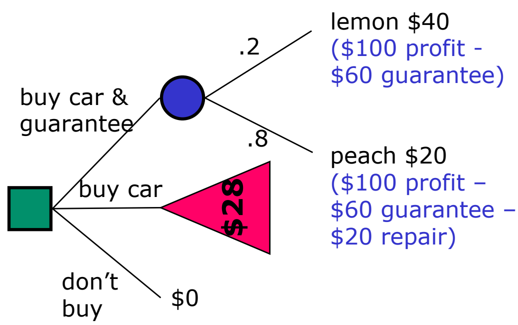

Finally, if we don't buy the car, we have a guaranteed zero.

Guarantee?

- Costs $60
- Covers 50% of repair costs
- If repairs > $100, covers them all

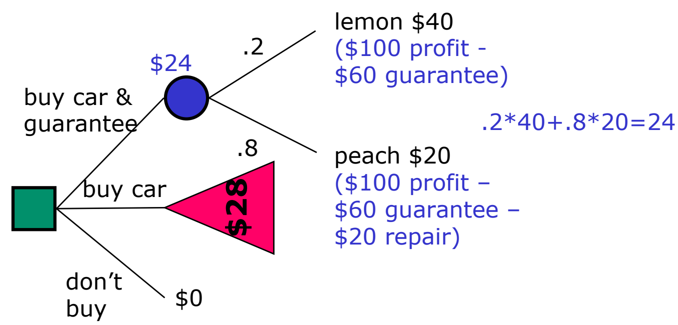

So, we start by figuring out the value of the new chance node. The expected value is $24.

Guarantee?

- Costs $60
- Covers 50% of repair costs
- If repairs > $100, covers them all

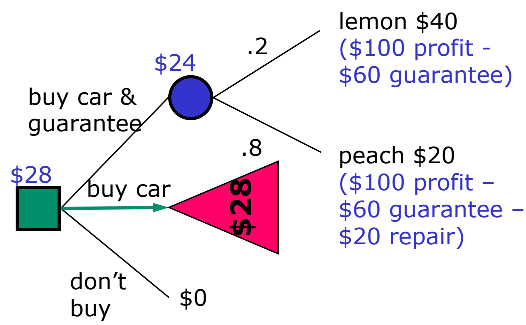

Now, at the choice node, we pick the maximum value course of action, which is to buy the car but not the guarantee. And so, this whole deal still has a value of $28.

Guarantee?

- Costs $60
- Covers 50% of repair costs
- If repairs > $100, covers them all

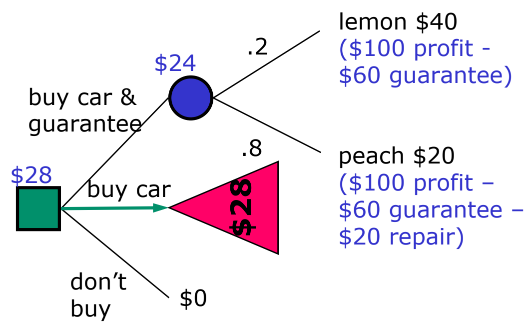

Most people would buy the guarantee in this situation, due to risk aversion.

Interestingly, most people would buy the guarantee in this situation, If you buy the guarantee, you are sure to make at least some profit. If you don't, you could possibly lose $100.

Inspection?

We can have the car inspected for $9

Just for some practice with Bayes' rule, and another example, let's think about one more case. This time there's no offer of perfect information or of a guarantee. But there is a garage across the street that offers to inspect the car for $9. It can't tell for sure whether the car is a lemon or a peach. It will perform a set of tests and tell you whether the car passes or fails.

Inspection?

We can have the car inspected for $9

P("pass" | peach) = 0.9 P("fail" | peach) = 0.1

P("pass" | lemon) = 0.4 P("fail" | lemon) = 0.6

The tests aren't completely reliable. So, the probability that it passes if it's a peach is 0.9. But the probability that it passes if it's a lemon is 0.4.

Inspection?

We can have the car inspected for $9

P("pass" | peach) = 0.9 P("fail" | peach) = 0.1

P("pass" | lemon) = 0.4 P("fail" | lemon) = 0.6

P("pass") =
P("pass" | lemon)P(lemon) + P("pass" | peach)P(peach)

P("pass") = 0.4*0.2 + 0.9*0.8 = 0.8

P("fail") = 0.2

We're going to need to have some related probabilities for use in our decision tree, so let's do it now. We'll need the probability of the car passing the test, which is 0.8. Note that it's only a coincidence that this is also the probability that the car is a peach.

Inspection?

We can have the car inspected for $9

P(“pass” | peach) = 0.9 P(“fail” | peach) = 0.1

P(“pass” | lemon) = 0.4 P(“fail” | lemon) = 0.6

P(“pass”) =
P(“pass” | lemon)P(lemon) + P(“pass” | peach)P(peach)

P(“pass”) = 0.4*0.2 + 0.9*0.8 = 0.8

P(“fail”) = 0.2

P(lemon | “pass”) = P(“pass” | lemon) P(lemon)/P(“pass”)

P(lemon | “pass”) = 0.4 * (0.2 / 0.8) = 0.1

P(lemon | “fail”) = P(“fail” | lemon) P(lemon)/P(“fail”)

P(lemon | “fail”) = 0.6 * (0.2 / 0.2) = 0.6

We’ll also need the probability of lemon given pass, which turns out to be 0.1, and the probability of lemon given fail, which turns out to be 0.6.

Inspection Tree

Let's build the decision tree for the inspection problem, so we can decide whether we should pay for the inspection or not.

Inspection Tree

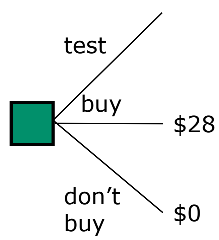

We start with a choice node, with three options: test (which means, we pay for the car to be inspected, and then perhaps buy it or not, depending on the results), buy, which means we just buy the car, which as we've calculated already has a value of $28, and don't buy, which has a value of $0.

Inspection Tree

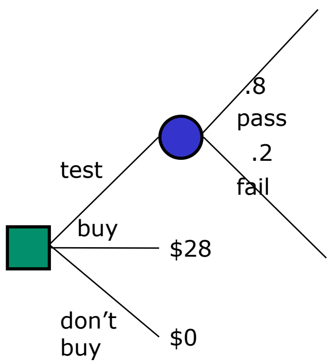

Now, if we decide to test the car, the next thing that happens is that we find out whether it passed or not. We just computed the probability of pass, which is 0.8, so that's the probability we put on this arc (and 0.2 on the fail arc).

Inspection Tree

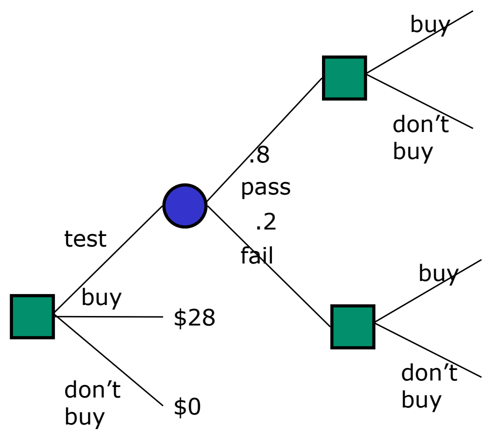

Once we have the information from the garage about whether the car passed or failed, we can decide whether to buy it or not. That results in these two decision nodes.

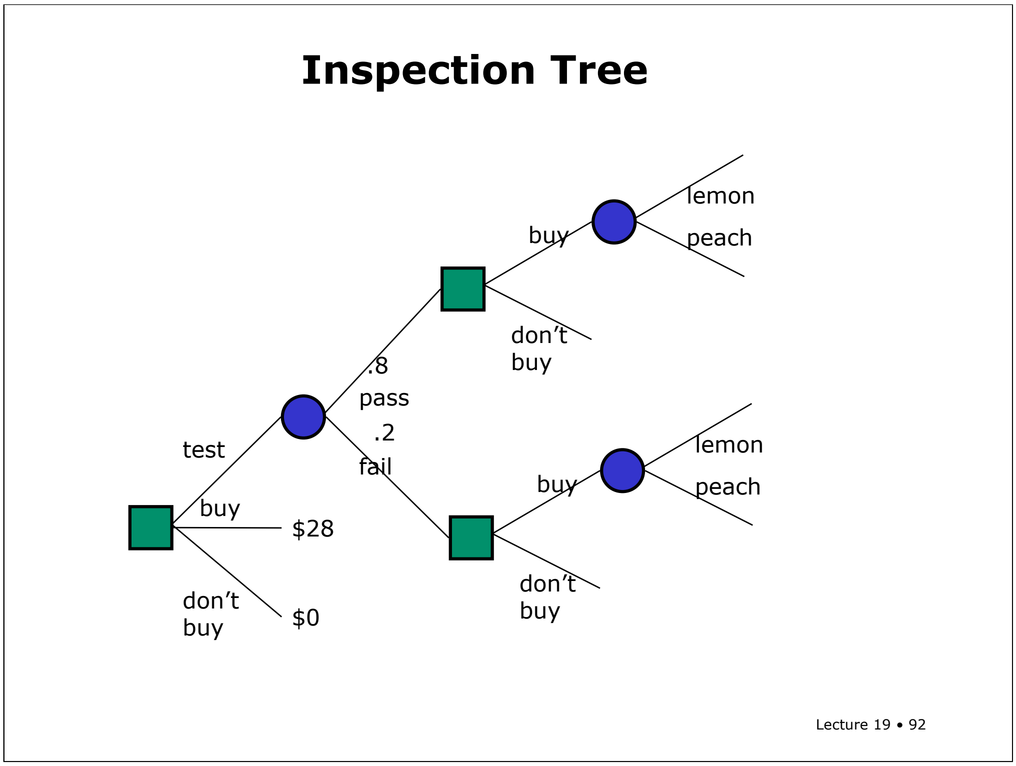

And finally, in each of the cases in which we buy the car, we have to add a chance node to take into account whether the car really is a lemon or not.

Inspection Tree

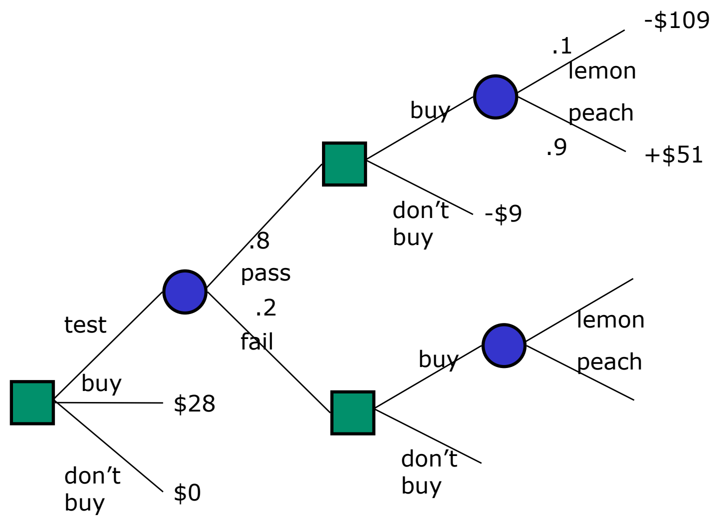

In the top node, we know that the car passed, and so when we put probabilities on the branches, we condition on the information we already know from having come down this path in the tree. So, on the "lemon" branch, we put the probability of lemon given pass, which is 0.1. And, of course we put 0.9 on the other branch. Then we can fill in the outcomes. We have to subtract $9 in every case, because in this whole branch of the tree, we're assuming that we paid for the test.

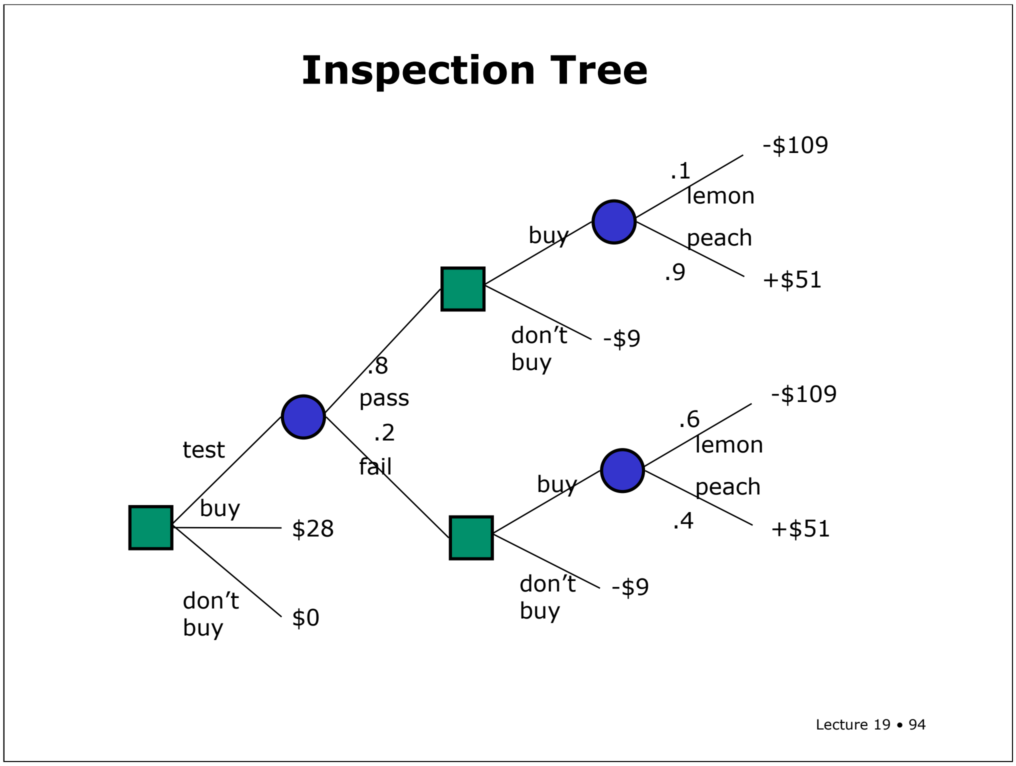

In the bottom choice node, we know that the car failed the test, so we label the lemon branch with the probability of lemon given fail, which is .6, and the other branch with .4. The outcomes are the same.

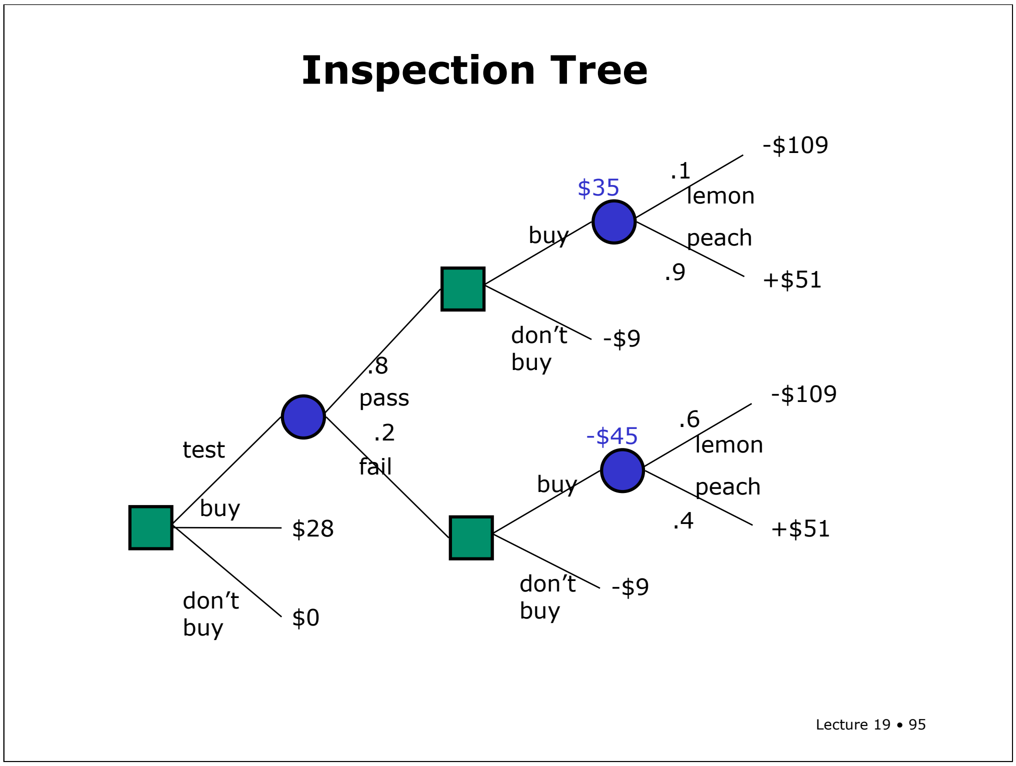

Whew! Now it's time to evaluate this tree. We start by evaluating the chance nodes nearest the leaves. We compute expected values, and get +$35 for the top one and -$45 for the bottom one.

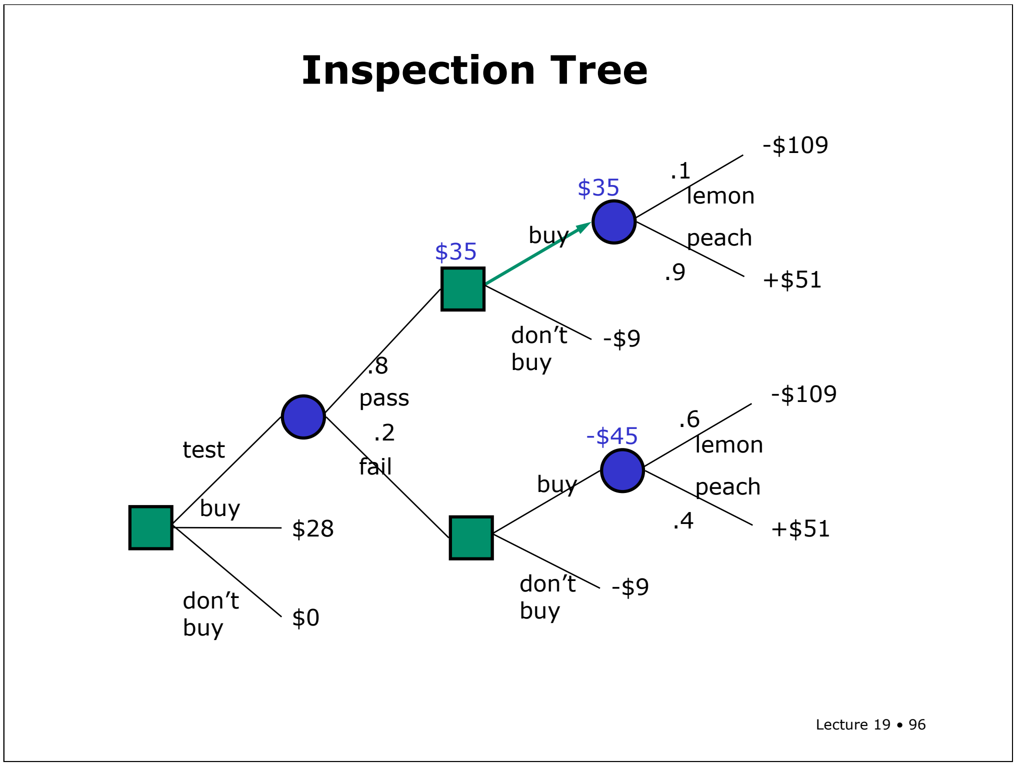

Now, we can do the choice nodes. In the top choice node, it's much better to buy the car, and the node gets value +$35.

Inspection Tree

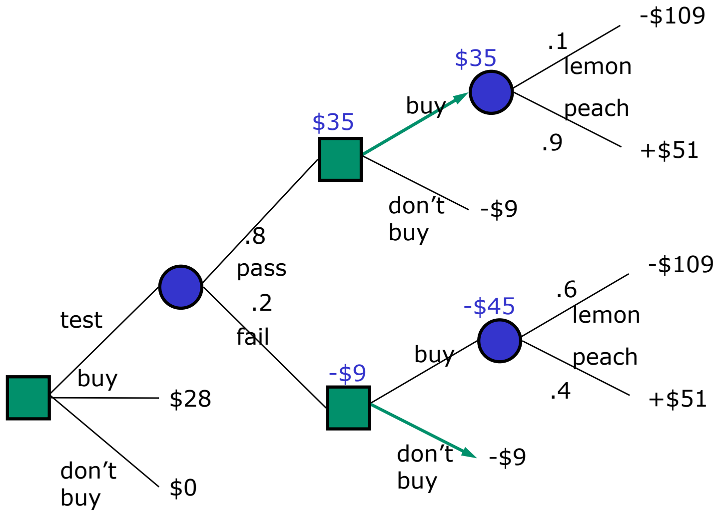

In the bottom choice node, buying the car is a bad idea, so we just end up with -$9, for having paid for the test.

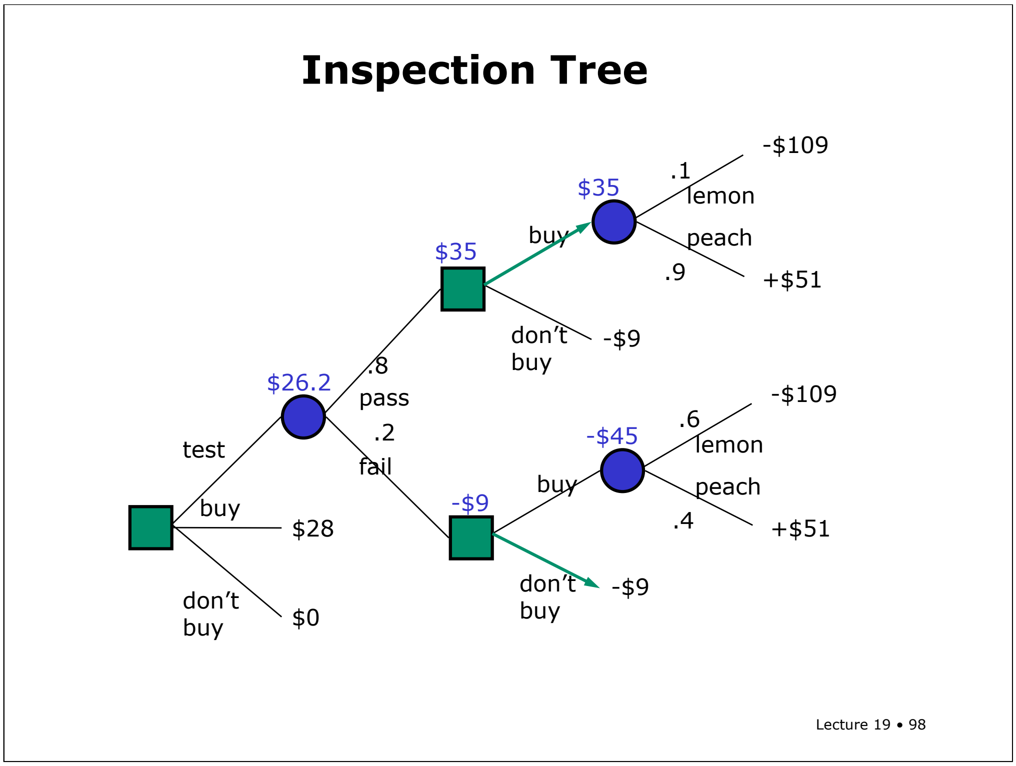

Now we're ready to evaluate the first chance node, by taking the expected value of its children. We get +$26.2.

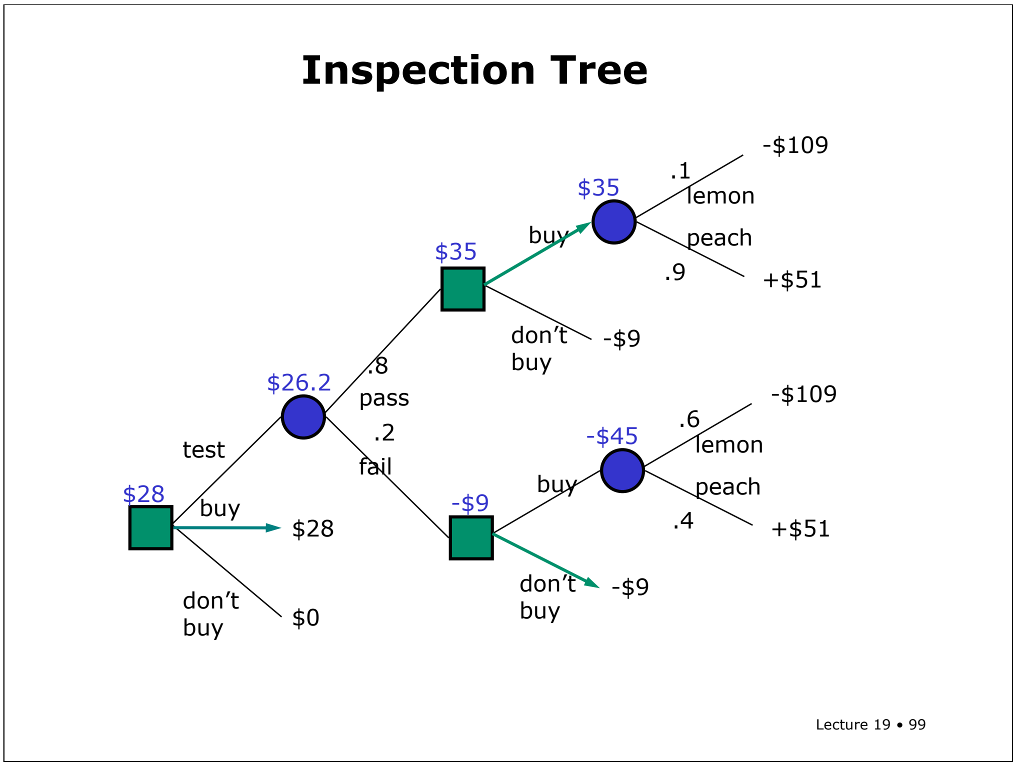

Finally, we can decide what action to take on the first step: we'll just buy the car directly, without having it inspected.

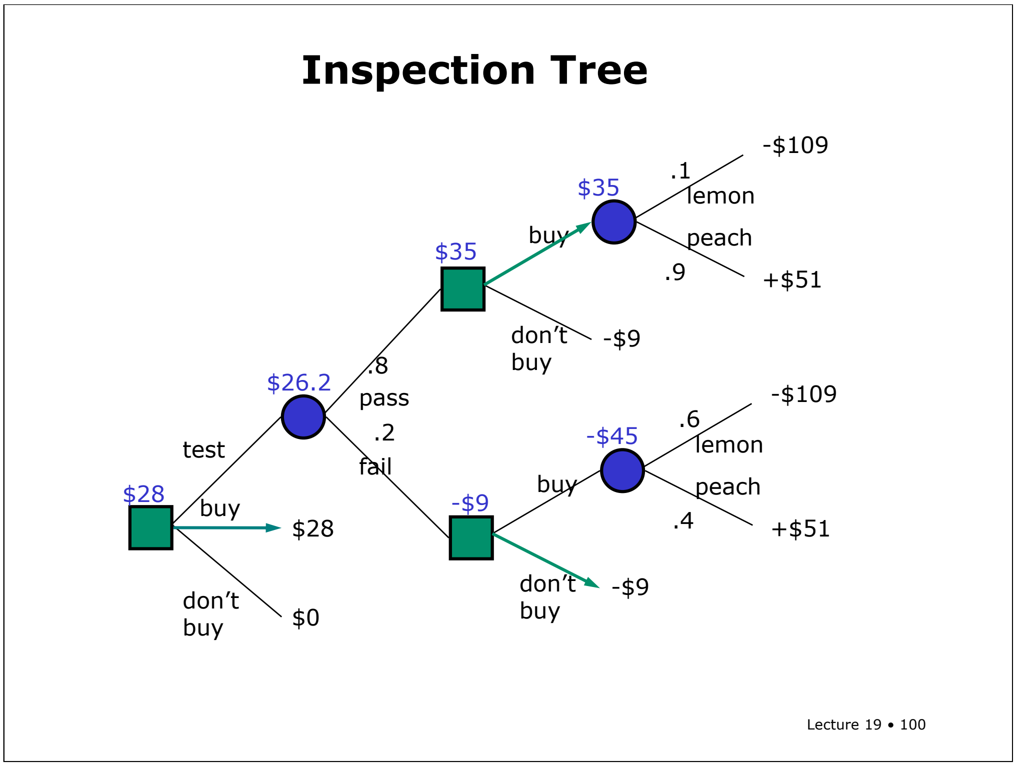

Again, this is probably a case in which a typical, somewhat risk averse person would pay for the test. The worst outcome on the test branch is losing $9, whereas the worst outcome on the buy branch is losing $100.

Recitation Problem

Let’s consider one last scenario in the purchase of used cars. We are going to have the car inspected, and then use the result of the inspection to decide if we will:

- buy the car without a guarantee
- buy the car with a guarantee
- not buy the car

Calculate the decision tree for this scenario. Use all the costs and probabilities from the previous scenarios. What is the expected value? Is it better than just buying the car ($28)?

Here’s one more car buying scenario. You can have the car inspected first, and then decide, based on the result, to buy the car with or without a guarantee, or not at all. What should you do?

Another Recitation Problem

Is it ever useful (in the sense of resulting in higher utility) to pay for information, but take the same action no matter what information you get?