WOODROW WILSON SCHOOL OF PUBLIC AND INTERNATIONAL AFFAIRS
MPP CAMP 2002: MICROECONOMICS

Lecture 3 Choice Under Uncertainty
Expected Utility

Reading KR chapter 6.4 (p. 185 – 192)

1. Examples

The Cuban Missile Crisis
The Cuban missile crisis began when Khrushchev's decision to send missiles to Cuba. Ideally (from Khrushchev's perspective), the missiles would be ready before the US found out about them. That would have changed the balance of power in the cold war and would have represented an important victory for the USSR. However, he may have recognized a risk that the US would find out about the missiles before they were ready to deploy and destroy or force their removal. We could model Khrushchev's decision as follows:

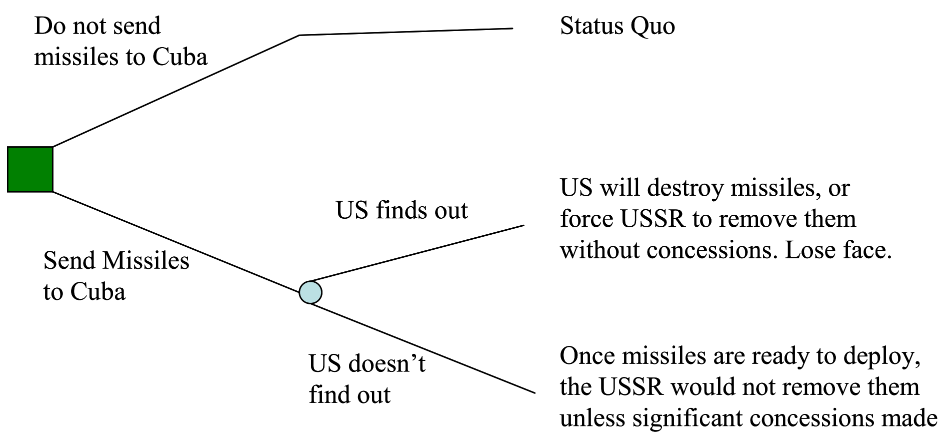

The diagram above is known as a *decision tree*. Decision trees have the following characteristics:

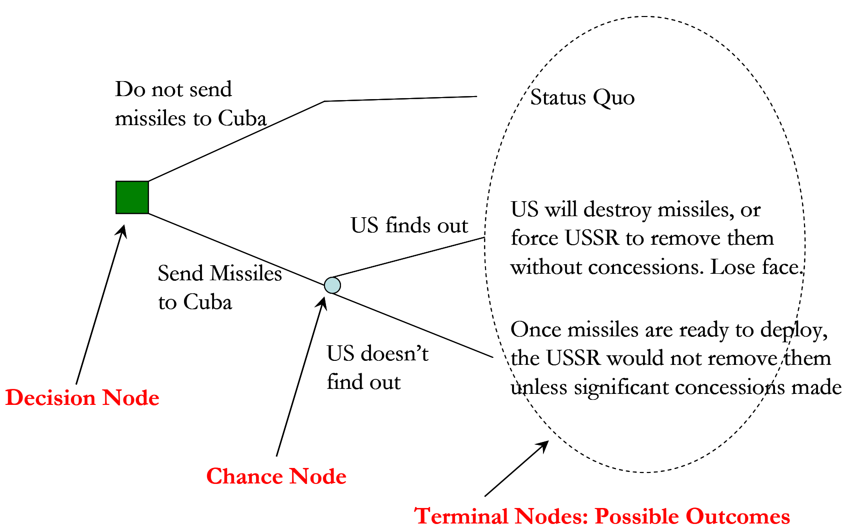

We have modeled the set **A** of choice elements as {Do not send missiles to Cuba, Send missiles to Cuba}. We can reasonably guess that Khrushchev preferences over the outcomes as follows:

(Missiles ready to deploy before US finds out) R (Status Quo) R (Missiles removed/destroyed);

or in utility terms:

u(Missiles ready to deploy before US finds out) > *u*(Status Quo) > *u*(Missiles removed/destroyed).

Investment Decision

Once upon a time, there was an MPP student who used to keep all her savings stashed inside her mattress. Upon learning about discounting and present value, and visiting the Monex web site, she decided that it might be a good idea to invest her money if only to protect it from inflation. She then asked her trustworthy microeconomics instructor for advice on where to invest her hard-earned cash for 6 months. The instructor first recommended an investment in bonds from an Argentine province where even the governor's salary hasn't been paid in the past 6 months. "Sure there is a high risk of default and you may lose all the money that you've been saving for 10 years," he said, "but if you invest your $1,000 savings today and the company doesn't default, you'll get 10,500 in six months!" The upside was appealing, but concerned about the risk of this investment, the student asks for a safer option. The instructor then promptly suggests frozen pork belly future contracts (available at the Chicago Mercantile Exchange). "In the worst case scenario, you can always take delivery and eat pork bellies for a whole year." The mystified student then goes home and models her problem as a decision tree:

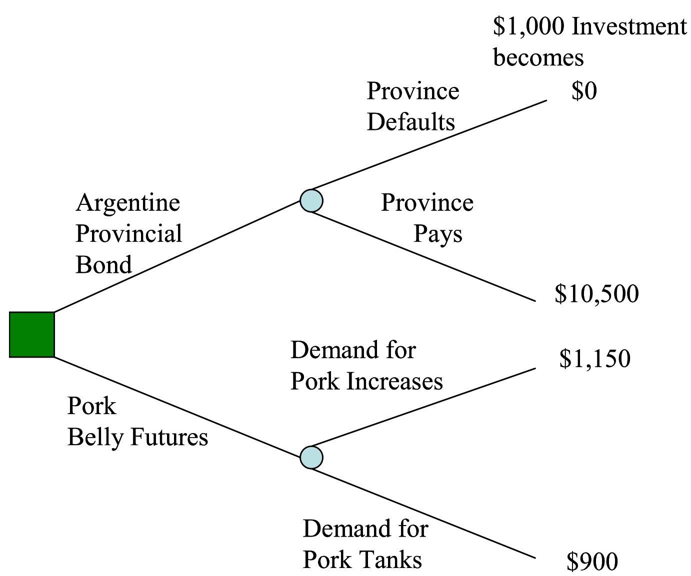

Notice that our previous model breaks down because we no longer have a one-to-one relationship between consequences (outcomes) and choices. Rather, one choice can have more than one outcome depending on a chance (or random) event, and it is no longer clear which choice maximizes utility. Therefore, we need to enrich our model if we wish to understand how choice is made when the action-consequence relationship is intermediated by chance.

2. Probability and Expected Value

Probability is an intuitive concept about the likelihood of the occurrence of specific outcomes in uncertain *events, experiments*, or *lotteries*; mathematics has simply modeled this intuitive concept so that we can apply it for things such as decision theory and statistics. We commonly associate precise probability values only to experiments such as tossing a coin or picking balls out of an urn, where the value of the probability can be objectively derived (if there are 10 balls in an urn and one is blue, the probability of picking the blue ball is 1/10). However, people generally have a notion of probability in mind for *any* uncertain event. For example, if I ask "Who will win the Princeton-Yale football game and you answer "Princeton", you believe that it is more likely that Princeton will win. In fact although most people wouldn't say "I'm 87.4% sure Princeton will win," common statements such as "it's a lot more likely that Princeton will win" also quantify probabilities. Our model of probability applies to both types of uncertainty: the objective kind (balls in an urn) or the subjective type (individual beliefs about the likelihood of events).

The basic mathematical model of probability goes as follows: first, identify the set of all possible outcomes of the experiment or event. In the case of the Princeton-Yale football game, the possible outcomes are Princeton wins, Yale wins, or the game is canceled[1]. Group these elements in a set, called **Ω** (omega) [2]. This set is called the *sample space.* If the event we're interested in is a coin toss, **Ω** = {Heads, Tails}, drawing balls from an urn, **Ω** = {Blue Ball, Red Ball}, etc. In economics, the elements of **Ω** are commonly referred to as "states of the world."

A *probability measure* P is a function (map) from all relevant subsets of **Ω** (meaning those subsets that we're interested in making probability statements about) to the interval [0,1] (or [0%,100%]) satisfying:

A) P(**Ω**) = 1 (100%)

B) P($\varnothing$) = 0

C) Probabilities of *disjoint* subsets (without common elements) of **Ω** can be added meaningfully. As special cases, the probabilities associated with each element of **Ω** can be added, and by A) they add up to 1. Similarly, P(**B**) + P($\mathbf{B}^c$) = 1.

Despite the fancy symbols, the underlying concept is very simple: we want to express numerically the idea of likelihood, such that a 60% probability associated with an outcome means that such outcome is more likely to happen than one to which we associate only a 10% probability. Part A) of the definition is simply restating the fact that **Ω** contains all possible outcomes of the experiment and defining 100% as certainty: if we define **Ω** correctly and include all possible outcomes of the experiment, with certainty one of its elements will occur. Therefore, 100% is the upper limit of probability. Although we may hear things such as "I'm 110% sure," in probability 100% is as certain as things get. B) is the other side of the same argument, and defines that an impossible event as having zero probability. Again, if we define **Ω** correctly, we know that one of its elements must occur with certainty, even if we are not sure which, so a set with no elements must have probability zero. C) is tricky because it is too intuitive. It doesn't make any sense that the probability of either Princeton winning or the game being canceled could different from 1 minus the probability of Yale winning, but that is exactly why we formally make this part of the model, so that when things get more complicated we don't inadvertently end up with such an inconsistency.

A *random variable* is a function from **Ω** to $\mathbb{R}$. For example, suppose I bet $1 that a coin comes up heads. The event we're interested in is the coin toss and **Ω** = {Heads, Tails}. If the outcome is heads, I get $1; if the outcome is tails, I lose $1. The amount of money I have after the experiment is a random variable: it could be $1 or -$1. So if we call this random variable 'X', X(Heads) = 1; X(Tails) = -1. The *expected value* or *expectation* of a random variable is the value that occurs on average. Therefore, the expected value of my bet is zero (½ x 1 + ½ x –1), since, on average, ½ the time I win a dollar and ½ the time I lose a dollar. Suppose the probability of Princeton winning is 90% and the probability of a tie or cancellation is 1%. I make a $1 bet with a friend saying that Yale will win; she insists Princeton will prevail. We agree that nobody wins or loses in case the game is canceled. Again, my cash position after the game is a random variable, and its expected value is 0.90 x –1 + 0.01 x 0 + 0.09 x 1 = –0.81. As you can see, expected values are weighted averages: it's the average of the payoffs weighted by the probabilities of getting each payoff.

---

[1] I was informed that college football games can no longer end in a draw.

[2] *Sets* are collections of items. Those items that belong to a set are referred to as its *elements*. We represent elements of a set between { }. The *empty set* contains no elements and is represented by the symbol $\varnothing$. We say **B** is a *subset* of **A** if all elements in **B** are also in **A**. The *complement* of a subset **B** of **A**, denoted $\mathbf{B}^c$, consists in all elements of **A** not in **B**. Therefore, if we add the elements of **B** to those of $\mathbf{B}^c$, we get back the set **A**. Note that **B** and $\mathbf{B}^c$ are *disjoint* – meaning they have no elements in common.

3. Expected Utility
Returning to our examples, how do Khrushchev and the MPP student decide between choices with uncertain outcomes? Our model starts by reframing the choice problem as a choice between lotteries. "Lotteries" is just a name for choices whose outcomes may be uncertain. Recall that some choices – such as Khrushchev not sending missiles to Cuba – do not involve uncertainty. Such choices are called "degenerate lotteries," as the probability of their unique outcome is 100%. Therefore, the set **A** of feasible actions becomes the set of lotteries, but it is still the same idea as under choice without uncertainty: it represents the available choices to the decision maker, and we continue to assume that the decision maker is fully aware of all the elements of **A**. *The uncertainty in this model is in the choice-consequence relationship, not on what choices are available.*

Next we model the decision maker's preferences between different lotteries. We retain the *completeness* and *transitivity* assumptions from the previous model, but must add two others. Unfortunately, showing how they relate to the final model of choice is somewhat complicated, but I believe it's nonetheless important to see all the underlying assumptions of economic models and where they might fail.

3. The *Substitution Axiom* is a type of "independence of irrelevant alternatives" property and is easiest understood with an example. Recall the investment decision example. Suppose that the MPP student decided to buy the Argentine provincial bond (therefore chose the bond over the pork bellies), and suppose I offer her the following proposition: (A) give me the $1,000 and I'll take it to Atlantic City. With 90% probability, I'll lose all the money and you get zero. With 10% probability, I make a killing, keep my profits, and buy $1,000 of Argentine bonds for you. OR (B) Give me the $1,000 and I'll take it to Atlantic City. With 90% probability, I'll lose all the money and you get zero. With 10% probability, I make a killing, keep my profits, and buy $1,000 of pork belly futures for you. If you are forced to choose between only (A) and (B), this property says that if you chose Argentine bonds over pork bellies, you must choose (A) over (B), implying that the fact that now you have a 90% probability of losing your money either way doesn't change your preference between gambling with Argentine bonds or pork bellies.

   In symbols: Let x, y and z be lotteries and p be a probability.
   If xRy, then $[px + (1-p)z]R[py + (1-p)z]$.
   In terms of our example, x=buy Argentine bond (a lottery, since its payoffs are uncertain), y=buy pork belly futures (another lottery), z=all money lost (a degenerate lottery), p = 10%.
   So if [Argentine Bonds]R[pork bellies], then [10%(Argentine Bonds) + 90%(all money lost)] R [10%(pork bellies) + 90%(all money lost)].

4. The *Continuity Axiom* states that if you have three alternatives (lotteries), you can always manipulate probabilities in such a way that you can make a lottery between the most and least preferred alternatives preferred to the lottery where you get the intermediate alternative for sure, or vice versa. So for example, if [Argentine Bonds] R [pork belly] R [gambling], then there is some (large) probability p such that [p(Argentine Bonds) + $(1-p)$(gambling)] R [pork belly], and there is some (small) probability q such that [pork belly] R [q(Argentine Bonds) + $(1-q)$(gambling)].

If preferences satisfy assumptions 1 – 4, then xRy if and only if $E(u(x)) \geq E(u(y))$. This is called a *von Neuman and Morgenstern* utility function. In words, lottery x is preferred to lottery y is the expected value of the utility I get out of lottery x is at least as great as the expected value of the utility I get out of lottery y. As in our original choice model, the decision maker will pick the choice which gives him or her the highest expected utility (which we just said is the equivalent of his or her most preferred alternative).

Two observations are in order before we revisit our earlier examples. First, we are treating subjective probability as if it were objective. In reality, when people make probability assessments, there is uncertainty about the probability itself. I think it's 90% likely that Princeton will win, but I'm only 50% confident on the 90%, etc. Second, because probabilities are cardinal numbers, von Neuman-Morgenstern utility is no longer ordinal. This is in a sense stated in the continuity axiom: If you truly dislike gambling, p will have to be extremely small for you to be willing to make a choice where you may end up gambling.

Consider again Khrushchev's choice and suppose that he assessed the probability of the US finding out at 70%. In addition, we need to assign utility numbers to the outcomes. Recall that a random variable is a function that ultimately gives you some number, which in this case is a utility number. We assume that Khrushchev would be much happier with the shift in balance of power than with the status quo, but that he was almost indifferent between the status quo and the Americans finding out. We then include this information in the decision tree below:

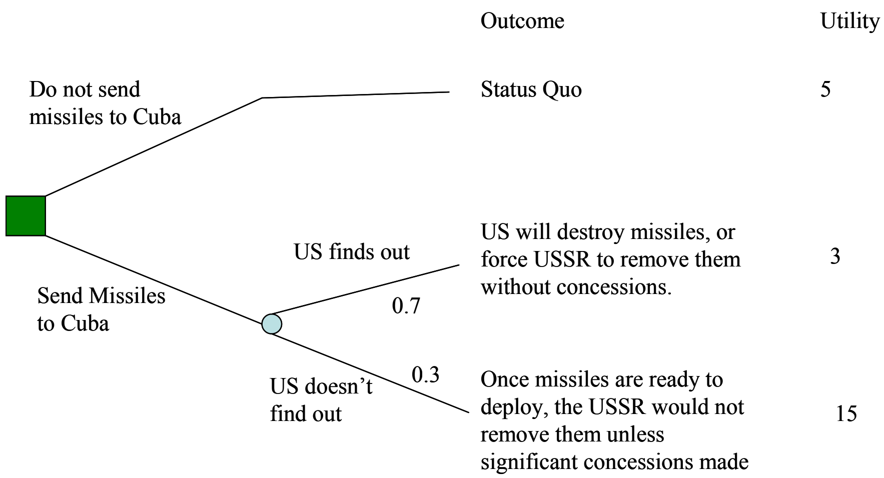

Which lottery does Khrushchev choose? The expected utility of the "do not send missiles" (degenerate) lottery is 1 x 5 =5. The expected utility of the "send missiles to Cuba" lottery is 0.7 x 3 (2.1) + 0.3 * 15 (4.5) = 6.6. Since 6.6 > 5, this model would explain the observation that, indeed, Khrushchev would choose to send the missiles to Cuba.

For our theory of choice, it is expected utility, not expected monetary payoff that counts. Therefore, in our investment decision example we must transform monetary amounts into utility by choosing a utility function for money. One commonly used function is square root (as is the case in the textbook), so u(payoff) = $\sqrt{payoff}$. We must also assign probabilities to the outcomes of each lottery.

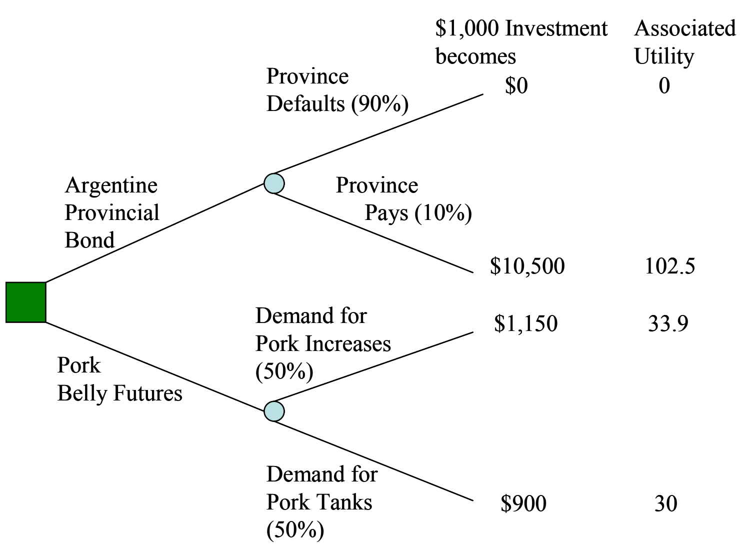

Which investment does the MPP choose? The expected utility of investing in the Argentine provincial bond is 0.9 x 0 + 0.1 x 102.5 = 10.3, while investing in pork belly futures gives an expected utility of 0.5 x 33.9 + 0.5 x 30.0 = 32.0. Since 32.0 > 10.3, clearly, the MPP buys pork belly futures.

4. Attitudes towards Risk

One key fact to note about the example above is that decisions are made on the expected value of *utility*, not monetary payoffs. In fact, if you calculate, the expected dollar payoff from the Argentine provincial bond is actually higher than the pork belly future ($1,050 vs. $1,025). The MPP's preferences are said to display *risk aversion*, since she showed a preference for the investment with less uncertainty. "Less uncertainty" here roughly means less variation in her monetary position between states of the world. The difference in her monetary position if she bought Argentine provincial bonds was $10,500, while the if she bought pork belly futures the difference was $250. Therefore, 'uncertainty' is a statement about the difference in payoffs in each state of the world rather than a difference in how likely each state is. The 'sure thing' not risky at all because the difference in your monetary position in any state of the world is zero (if you don't gamble, regardless of how the roulette turns out, you still have your dollar).

Note that we chose a utility function, square root, which modeled risk aversion. In contrast, suppose the utility function was u(payoff) = payoff. Then expected utility is equal to expected dollar payoff and the student would have bought the Argentine provincial bonds. In that case, the individual would be *risk neutral*. A *risk loving* individual is one who would prefer a riskier lottery even if the expected monetary payoff of that lottery were lower.

The mistake made in class was that individuals who are risk averse prefer lotteries with lower risk, so if the *monetary payoff* of two lotteries is the same, they go for the less risky one. But that means that their *expected utility is higher with the less risky lottery*. If the expected utility is the same, the individual is indifferent between the lotteries.

5. Criticisms of Expected Utility Theory.

Allais Paradox (From Kreps 1988)

The Allais paradox is one of the first, and certainly the best known criticism of standard models of choice under uncertainty. His results were replicated many times, and the example in Kreps (1988) is from Kahneman and Tversky (1979).

Individuals were asked to choose between two lotteries:

Lottery A			
Probability (%)	33	66	1
Prize ($)	2,500	2,400	0

Lottery B	
Probability (%)	100
Prize ($)	2,400

Then they were told to choose between two other lotteries:

Lottery C		
Probability (%)	33	67
Prize ($)	2,500	0

Lottery D		
Probability (%)	34	66
Prize ($)	2,400	0

Kahneman and Tversky observe that 82% of subjects picked lottery B over lottery A, while 83% chose lottery C over lottery D. Therefore, at least 65% of their subjects choose B in the first case and C in the second. But this is inconsistent with our expected utility theory:

If A R B, it must be the case that $u(2,400) \geq 0.33 \bullet u(2,500) + 0.66 \bullet u(2,400) + 0.01 \bullet u(0)$. Regardless of what function we choose, u(2,400) has to be the same on either side of the equation, so we can move one item from the right hand side to the left hand side and get $0.34 \bullet u(2,400) \geq 0.33 \bullet u(2,500) + 0.01 \bullet u(0)$. Now if we add to each side 0.66 times u(0), which we are allowed by the substitution axiom, we get $0.66 \bullet u(0) + 0.34 \bullet u(2,400) \geq 0.33 \bullet u(2,500) + 0.67 \bullet u(0)$. But this means those who chose lottery A in the first case would have to choose lottery D in the second, a contradiction to actual observations.

Framing Effects/Tendency to Simplify the Problem.
(from Rubinstein 1998)
This is another experiment performed by Kahneman and Tversky (1986). Individuals were again asked to pick between two lotteries, both of which involve the spinning of a roulette wheel.

Lottery A

Color	White	Red	Green	Yellow
Probability (%)	90	6	1	3
Prize ($)	0	45	30	-15

Lottery B

Color	White	Red	Green	Yellow
Probability (%)	90	7	1	2
Prize ($)	0	45	-10	-15

58% of Kahneman and Tversky's subjects chose A, while in class 86% of students chose A. Individuals were then asked to consider lotteries C and D below:

Lottery C

Color	White	Red	Green	Blue	Yellow
Probability (%)	90	6	1	1	2
Prize ($)	0	45	30	-15	-15

Lottery D

Color	White	Red	Green	Blue	Yellow
Probability (%)	90	6	1	1	2
Prize ($)	0	45	45	-10	-15

The lottery D dominates C and both in class and in the experiment just about everyone chose D. However, notice that lottery A is essentially the same as lottery C (blue and yellow are combined in A), and lottery B is essentially the same as lottery D (red and green in D are combined in B). This clearly violates expected utility theory since if we tried to predict behavior in the second lottery by individuals choices in the first, we would have failed. Rubinstein argues that decision makers try to simplify problems and similarity relations are used for this purpose. 6 and 7 percent, and likewise 2 and 3 percent are similar, so all the odds appear fairly similar and individuals focus on the payoffs. The only different payoff is green, and a comparison between -10 and 30 lead individuals to choose lottery A.

5. Reading
Please read Thomas Friedman's New York Times Op-ed article on possible states of the world for oil prices if the US invades Iraq.
http://www.nytimes.com/2002/07/31/opinion/31FRIE.html (requires free registration)

For more on choice theory and criticisms, two books are highly recommended:
David Kreps *Notes on the Theory of Choice,* Westview Press, 1988.
Ariel Rubinstein *Modeling Bounded Rationality* MIT Press, 1998
There are links to the original Kahneman and Tversky papers above, and in the Lecture 2 there is a link to a survey paper that also contains criticisms of expected utility.

6. Exercises

1. Read Thomas Friedman's Op-ed article and draw a decision tree for an oil importer in the US trying to decide whether to buy a future contract on oil at $30 (i.e., he commits to buy the oil from the futures seller in six months at $30 per barrel regardless of the going price then). Use the following utility functions:

   $u(\text{price of oil}) = 100 - \sqrt{\text{price of oil}}$ and $u(\text{price of oil}) = -2 \bullet (\text{price of oil})$

2. Finance specialists are frequently recommending that individuals diversify their portfolios. Analyze what would happen to the MPP problem by introducing a third choice: investing $500 on Argentine provincial bonds and $500 on pork belly futures. Re-draw the decision tree and conclude which of the three options the student would choose.

3. Martha Stewart is in trouble because of allegations she had inside information about ImClone stock, which she used to make a substantial amount of money by selling the stock before the market received some bad news (a drug was not going to be approved by the FDA). Write the decision tree of a random market participant for selling ImClone stock, assuming random participants do not sell. Assume the price of ImClone stock would go up $10 if the drug was approved by the FDA, and would go down $20 if it didn't. Write Martha Stewart's decision tree assuming she did have inside information. Maybe Martha Stewart neglected the possibility of being caught by the SEC – or did she? Draw another decision tree that takes into account the probability that she may get caught.

4. (optional) KR problem 6.10, page 194.

5. (optional – adapted from Kreps 1992) Let x be a lottery giving prizes $10 and $20 with probabilities 2/3 and 1/3, respectively, and let y be a lottery giving prizes $5, $15, and $30 with probabilities 1/3, 5/9 and 1/9 respectively. Show that any risk averse expected utility maximizer will (weakly) prefer x to y.