Lecture Note 16: Uncertainty, Risk Preference, and Expected Utility Theory

David Autor, MIT and NBER

14.03/14.003, Microeconomic Theory and Public Policy, Fall 2016

1 Risk Aversion and Insurance: Introduction

- A significant hole in our theory of consumer choice developed in 14.03/14.003 to date is that we have only modeled choices that are devoid of *uncertainty*: everything is known in advance. That's convenient, but not particularly plausible.
  - Prices change
  - Income fluctuates
  - Bad stuff happens
- Most decisions are *forward-looking*: these decisions depend on our beliefs about what is the optimal plan for present and future. Inevitably, such choices are made in a context of uncertainty. There is a risk (in fact, a likelihood) that not all scenarios we hoped for will be borne out. In making plans, we should take these contingencies and probabilities into account—and there is no doubt that people *do* take these things into account. If we want a realistic model of choice, we need to model how uncertainty affects choice and well-being.
- This model should help to explain:
  - How do people choose among "bundles" that have uncertain payoffs, e.g., whether to fly on an airplane, whom to marry?
  - Insurance: Why do people want to buy it?
  - How (and why) the *market* for risk operates? (Markets for risk include life insurance, auto insurance, gambling, futures markets, warranties, bonds, etc.)

1.1 A few motivating examples

1. People don't seem to want to play actuarially fair games. Such a game is one in which the cost of entry is equal to the expected payoff:

$$E(X) = P_{win} \cdot [\text{Payoff}|\text{Win}] + P_{lose} \cdot [\text{Payoff}|\text{Lose}] .$$

   - Most people would not enter into a $\$1,000$ dollar heads/tails fair coin flip.

2. People won't necessarily play actuarially *favorable* games:

   - You are offered a gamble. We'll flip a coin. If it's heads, I'll give you \$10 million dollars. If it's tails, you owe me \$9.8 million.

Its expected monetary value is :

$$\frac{1}{2}\cdot 10-\frac{1}{2}\cdot 9.8=\$0.1\text{ million}(\$100,000)$$

Want to play?

3. People won't pay large amounts of money to play games with huge upside potential. Example "St. Petersburg Paradox."

   - Flip a coin. I'll pay you in dollars 2^n, where n is the number of tosses until you get a head:

   $$X_1=\$2,X_2=\$4,X_3=\$8,...X_n=2^n.$$

   - What is the expected value of this game?

   $$E(X)=\frac{1}{2}2+\frac{1}{4}4+\frac{1}{8}8+...\frac{1}{2^n}2^n=\infty.$$

   - How much *would* you be willing to pay to play this game? [People generally do not appear willing to pay more than a few dollars to play this game.]
   - What is the variance of this gamble? $V(X)=\infty$.

- The fact that a gamble with infinite expected monetary value has (apparently) limited *utility value* suggests something pervasive and important about human behavior: *As a general rule, uncertain prospects are worth less in utility terms than certain ones, even when expected tangible payoffs are the same.*

- To have a coherent model of choice under uncertainty, we need to be able to say how people make choices when:

  – Consumers value outcomes (as we have modeled all along) *and*
  – Consumers have feelings/preferences about the riskiness of those outcomes

We'll introduce the notion of risk aversion, insurance, and insurance *markets* in several steps. First, I'll review some basic probability theory, which will likely already be familiar. Next, I'll provide an *optional* formal development of so-called Von Neumann-Morgenstern expected utility theory (AKA, Expected Utility Theory). You don't have to spend time with the formal development, but you are welcome to do so. Third, I'll provide an *informal* discussion of the Expected Utility property. Following that, I'll show how preferences that satisfy the Expected Utility property can be used to formalize notions of risk preference—specifically, risk averse, risk neutral, and risk seeking preferences. From there, we will use these tools to understand how and why insurance markets work, and why risk is a good (or bad) that consumers will want to trade.

Here's a noteworthy feature of markets for risk that you should contemplate as you study this material: in settings with risk, there may be gains from trade—that is, potential Pareto improvements—*even when* all consumers have identical preferences and endowments.

It is also important to note before we launch in that the models we are going to develop in this lecture note incorporate uncertainty in a very particular way. While we will will explore uncertainty in the probability of certain outcomes occurring, we are going to maintain the assumption that consumers know two important things:

1. They know the set of possible outcomes

2. They know the probabilities of each of the outcomes

These are strong assumptions, but probably not unreasonable for certain important settings.

2 Five Simple Statistical Notions

Definition 1. Probability distribution

Define states of the world $1, 2...n$ with probability of occurrence $\pi_1, \pi_2...\pi_n$.
A valid probability distribution satisfies:

$$\sum_{i=1}^{n} \pi_i = 1, \text{ or } \int_{-\infty}^{\infty} f(x)\partial x = 1 \text{ and } f(x) \geq 0 \; \forall \; x.$$

In this equation $f(x)$ is the 'probability density function' (PDF) of the continuous random variable x, meaning that $f(x)$ is essentially the probability of randomly drawing the given value x (so, $f(x)$ is just like the π_i in the discreet case). [Note that the probability of drawing any specific value from a continuous distribution is zero since there are an infinite number of possibilities. Depending on the distribution, however, some ranges of values will be much more likely than others.]

-

Definition 2. Expected value or "expectation"
The mean of a random variable (a notion that we've used all semester).
Say each state i has payoff x_i. Then

$$E(x) = \sum_{i=1}^{n} \pi_i x_i \text{ or } E(x) = \int_{-\infty}^{\infty} x f(x)\partial x.$$

Example: Expected value of a fair dice roll is $E(x) = \sum_{i=1}^{6} \pi_i i = \frac{1}{6} \cdot 21 = \frac{7}{2}$. We've been using expectations throughout the semester, so no doubt this is familiar.

-

Definition 3. Variance (dispersion)
Gambles with the same expected value may have different dispersion.
We'll measure dispersion with variance.

$$V(x) = \sum_{i=1}^{n} \pi_i \left(x_i - E(x)\right)^2 \text{ or } V(x) = \int_{-\infty}^{\infty} \left(x - E(x)\right)^2 f(x) \partial x.$$

In dice example, $V(x) = \sum_{i=1}^{n} \pi_i \left(i - \frac{7}{2}\right)^2 = 2.92$.

-

Definition 4. Independence.
A case in which the probabilities of two (or multiple) outcomes do not depend upon one another. If events A and B are independent, then $\Pr(A \text{ and } B) = \Pr(A) \cdot \Pr(B)$, and similarly, $E[A \cdot B] = E[A] \cdot E[B]$.
Example: The probability of flipping two sequential heads with a fair coin is $\Pr(H \text{ and } H) = \Pr(H) \cdot \Pr(H) = 0.25$. These probabilities are *independent*, meaning that if you flipped heads the first time, you are no more or less likely to flip heads the second time.

Example: The probabilities of seeing lightning and hearing thunder in an afternoon are *not independent* of one another. If you see lightning, you're reasonably likely to hear thunder, and vice versa.

-

Definition 5. Law of Large Numbers
In repeated, independent trials with the same probability p of success in each trial, the chance that the percentage of successes differs from the probability p by more than a fixed positive amount $e > 0$ converges to zero as number of trials n goes to infinity for every positive e.

Example: If you flip a fair coin 100 times, the probability of getting heads more than $\geq 51\%$ of the time (that is, 51 or more times) is reasonably high. If you flip a fair coin 100,000 times, the probability of getting heads more than $\geq 51\%$ of the time (that is, 510,00 or more times) is vanishingly small.

•

Dispersion and risk are closely related notions. Holding constant the expectation of X, more dispersion means that the outcome is "riskier"—it has both more upside and more downside potential.

Consider four gambles:

1. \$0.50 for sure. $V(L_1) = 0$.

2. Heads you receive \$1.00, tails you receive 0.
$V(L_2) = 1 \times [0.5 \times (1 - .5)^2 + 0.50 \times (0 - .5)^2] = 0.25$

3. 4 independent flips of a coin, you receive \$0.25 on each head.
$V(L_3) = 4 \times \left[\frac{1}{2}(0.25 - 0.125)^2 + \frac{1}{2} \times (0 - 0.125)^2\right]) = 0.0625$

4. 100 independent flips of a coin, you receive \$0.01 on each head.
$V(L_4) = 100 \times \left[\frac{1}{2}(0.01 - 0.005)^2 + \frac{1}{2}(0 - 0.005)^2\right] = 0.0025$

All four of these "lotteries" have same expected value (50 cents), but they have different levels of risk.

A key statistical result, which I will not prove here, is that the variance of n identical *independent* gambles is $\frac{1}{n}$ times the variance of one of the gambles. What this means in practice is that pooling a large number of independent, identical gambles reduces the aggregate riskiness of those gambles. This is closely related to the previous example of flipping a coin 100 versus 100,000 times. The more independent gambles in the pool—the more flips of a fair coin—the greater the certainty with which you can forecast the mean outcome.

2.1 Lottery Details

- The basic building block of our theory is the concept of a *lottery.*

  Definition 6. A simple lottery L is a list $L = (p_1, ...p_N)$ with $p_n \geq 0$ for all n and $\Sigma_n p_n = 1$, where p_n is interpreted as the probability of outcome n occurring.

- In a simple lottery, the outcomes that may result are certain.

- A more general variant of a lottery, known as a *compound lottery*, allows the outcomes of a lottery to themselves be simple lotteries.

  Definition 7. Given K simple lotteries $L_k = \left(p_1^k, ...p_N^k\right)$, $k = 1..., K$, and probabilities $\alpha_k \geq 0$ with $\Sigma_k \alpha_k = 1$, the compound lottery $(L_1, ...L_K; \alpha_1, ...\alpha_K)$, is the risky alternative that yields the simple lottery L_k with probability α_k for $k = 1, ..., K$.

- For any compound lottery $(L_1, ...L_K; \alpha_1, ...\alpha_K)$, we can calculate a corresponding *reduced lottery* as the simple lottery $L = (p_1, ...p_N)$ that generates the same ultimate distribution over outcomes. So, the probability of outcome n in the reduced lottery is:

$$p_n = \alpha_1 p_n^1 + \alpha_2 p_n^2 + ... + \alpha_k p_n^k.$$

  That is, we simply add up the probabilities, p_n^k, of each outcome n in all lotteries k, multiplying each p_n^k by the probability α_k of facing each lottery k.

- We now study the decision maker's *preferences over lotteries.*

- The basic premise of the model that follows is what philosophers would call a 'consequentialist' viewpoint: for any risky alternative, the decision maker cares only about the outcomes and their associated probabilities, or in technical terms, the *reduced lottery* over final outcomes. By assumption, the decision maker is indifferent to the (possibly many) compound lotteries underlying these reduced lotteries.

- This compound lottery assumption states that the 'frame' or order of lotteries is unimportant. So consider the following two stage lottery:

  - Stage 1: You flip a coin: heads or tails.
  - Stage 2:
    If the Stage 1 flip drew heads, you flip the coin again. Heads yields \$1.00, tails yields \$0.75.
    If the Stage 1 flip drew tails, you roll a dice with payoffs $\$0.10, \$0.20, ...\$0.60$ corresponding to outcomes $1 - 6$.

- Now consider a single state lottery, where:

  - We spin a pointer on a wheel with 8 areas, 2 areas of 90^0 representing \$1.00, and \$0.75, and 6 areas of 30^0 each, representing $\$0.10, \$0.20, ...\$0.60$ each.
  - This single stage lottery has the same payouts at the same odds as the $2-$stage lottery.
  - The 'compound lottery' axiom says the consumer is indifferent between these two.
  - Counterexamples? [This is not an innocuous set of assumptions.]
  - [Is this realistic? Hard to develop intuition on this point, but research shows that this assumption is often violated.]

- Implicitly, we are assuming that what enters into the decision maker's utility function is the final outcomes of these lotteries—the actual bundles consumed—not the probabilities along the way. If that assumption is correct from a utility perspective, then compound lotteries can be collapsed into simple lotteries so long as both the compound and simple lottery give rise to the identical set of consumption bundles with identical probabilities of consumption. (Another way to say this: the consumer does not consume the probabilities, only the realized outcomes.)

3 [Optional] Risk preference and expected utility theory[1]

[This section derives the Expected Utility Theorem. I will not cover this material in class and I will not hold you responsible for the technical details.]

3.1 Preferences over lotteries

- Consider the set of alternatives the decision maker faces, denoted by $\mathcal{L}$ to be the set of all simple lotteries over possible outcomes N.

- We assume the consumer has a rational preference relation $\succsim$ on $\mathcal{L}$, a *complete* and *transitive* relation allowing comparison among any pair of simple lotteries (I highlight the terms *complete* and *transitive* to remind you that they have specific meaning from axiomatic utility theory, given at the beginning of the semester). [This could also be called an axiom—or even two axioms!]

- **Axiom 1. Continuity**. *Small changes in probabilities do not change the nature of the ordering of two lotteries. This can be made concrete here (I won't use formal notation b/c it's a mess). If a "bowl of miso soup" is preferable to a "cup of Kenyan coffee," then a mixture of the outcome "bowl of miso soup" and a sufficiently small but positive probability of "death by sushi knife" is still preferred to "cup of Kenyan coffee."*

- Continuity rules out "lexicographic" preferences for alternatives, such as "safety first." Safety first is a lexicographic preference rule because it does not *trade-off* between safety and competing alternatives (fun) but rather simply requires safety to be held at a fixed value for any positive utility to be attained.

[1]This section draws on Mas-Colell, Andreu, Michael D. Winston and Jerry R. Green, *Microeconomic Theory*, New York: Oxford University Press, 1995, chapter 6. For those of you considering Ph.D. study in economics, MWG is the only text that covers almost the entire corpus of modern microeconomic theory. It is the Oxford English Dictionary of modern economic theory. Most economists keep it on hand for reference; few read it for pleasure.

- The second key building block of our theory about preferences over lotteries is the so-called *Independence Axiom.*

- **Axiom 2. Independence.** *The preference relation* $\succsim$ *on the space of simple lotteries* $\pounds$ *satisfies the independence axiom if for all* $L, L', L'' \in \pounds$ *and* $\alpha \in (0,1)$*, we have*

$$L \succsim L' \text{ if and only if } \alpha L + (1-\alpha) L'' \succsim \alpha L' + (1-\alpha) L''.$$

- In words, when we mix each of two lotteries with a third one, then the preference ordering of the two resulting mixtures does not depend on (is *independent of*) the particular third lottery used.

- Example: If a bowl of miso soup is preferred to cup of Peets coffee, then the lottery (bowl of miso soup with 50% probability, steak dinner with 50% probability) is preferred to the lottery (cup of Peets coffee with 50% probability, steak dinner with 50% probability).

3.2 Expected utility theory

- We now want to define a class of utility functions over risky choices that have the "expected utility form." We will then prove that if a utility function satisfies the definitions above for *continuity* and *independence* in preferences over lotteries, then the utility function has the expected utility form.

- It's important to clarify now that "expected utility theory" does *not* replace consumer theory, which we've been developing all semester. Expected utility theory extends the model of consumer theory to choices over risky outcomes. Standard consumer theory continues to describe the utility of consumption of specific *bundles.* Expected utility theory describes how a consumer might select among risky bundles. [This paragraph will be repeated below in the non-optional section of the lecture note.]

  Definition 8. *The utility function* $U : \pounds \to \mathbb{R}$ *has an expected utility form if there is an assignment of numbers* $(u_1, ...u_N)$ *to the* N *outcomes such that for every simple lottery* $L = (p_1, ..., p_N) \in \pounds$ *we have that*

$$U(L) = u_1 p_1 + ... + u_N p_N.$$

- A utility function with the expected utility form is called a Von Neumann-Morgenstern (VNM) expected utility function.

- The term *expected utility* is appropriate because with the VNM form, the utility of a lottery can be thought of as the expected value of the utilities u_n of the N outcomes.

- In other words, a utility function has the expected utility form if and only if:

$$U\left(\sum_{k=1}^{K}\alpha_k L_k\right) = \sum_{k=1}^{K}\alpha_k U\left(L_k\right)$$

for any K lotteries $L_k \in \pounds$, $k = 1, ..., K$, and probabilities $(\alpha_1, ..., \alpha_K) \geq 0$, $\Sigma_k \alpha_k = 1$.

- Intuitively, a utility function has the expected utility property if the utility of a lottery is simply the (probability) weighted average of the utility of each of the outcomes.

- A person with a utility function with the expected utility property flips a coin to gain or lose one dollar. The utility of that lottery is

$$U\left(L\right) = 0.5U\left(w+1\right) + 0.5U\left(w-1\right),$$

where w is initial wealth.

- Q: Does that mean that

$$U\left(L\right) = 0.5\left(w+1\right) + 0.5\left(w-1\right) = w?$$

No. We haven't actually defined the utility of an *outcome*, and we certainly don't want to assume that $U\left(w\right) = w$.

3.3 Proof of expected utility property

Proposition. *(Expected utility theory) Suppose that the rational preference relation $\succsim$ on the space of lotteries $\pounds$ satisfies the continuity and independence axioms. Then $\succsim$ admits a utility representation of the expected utility form. That is, we can assign a number u_n to each outcome $n = 1, ..., N$ in such a manner that for any two lotteries $L = (p_1, ...P_N)$ and $L' = (p'_1, ...P'_N)$, we have $L \succsim L'$ if and only if*

$$\sum_{n=1}^{N} up_n \geq \sum_{n=1}^{N} u_n p'_n$$

Proof. Expected Utility Property (in five steps) □

Assume that there are best and worst lotteries in $\pounds$, $\bar{L}$ and $\underline{L}$.

1. If $L \succ L'$ and $\alpha \in (0,1)$, then $L \succ \alpha L + (1-\alpha) L' \succ L'$. This follows immediately from the independence axiom.

2. Let $\alpha, \beta \in [0,1]$. Then $\beta \bar{L} + (1-\beta)\underline{L} \succ \alpha \bar{L} + (1-\alpha)\underline{L}$ if and only if $\beta > \alpha$. This follows from the prior step.

3. For any $L \in \pounds$, there is a unique α_L such that $\left[\alpha_L \bar{L} + (1-\alpha_L)\underline{L}\right] \sim L$. Existence follows from continuity. Uniqueness follows from the prior step.

4. The function $U : \pounds \to \mathbb{R}$ that assigns $U(L) = \alpha_L$ for all $L \in \pounds$ represents the preference relation $\succsim$.
Observe by Step 3 that, for any two lotteries $L, L' \in \pounds$, we have

$$L \succsim L' \text{ if and only if } \left[\alpha_L \bar{L} + (1-\alpha_L)\underline{L}\right] \succsim \left[\alpha_{L'} \bar{L} + (1-\alpha_{L'})\underline{L}\right].$$

Thus $L \succsim L'$ if and only if $\alpha_L \geq \alpha_{L'}$.

5. The utility function $U(\cdot)$ that assigns $U(L) = \alpha_l$ for all $L \in \pounds$ is linear and therefore has the expected utility form.
We want to show that for any $L, L' \in \pounds$, **and** $\beta[0,1]$, **we have** $U(\beta L + (1-\beta)L') = \beta U(L) + (1-\beta)U'(L)$.
By step (3) above, we have

$$\begin{aligned} L &\sim U(L)\bar{L} + (1-U(L))\underline{L} = \alpha_L \bar{L} + (1-\alpha_L)\underline{L} \\ L' &\sim U(L')\bar{L} + (1-U(L'))\underline{L} = \alpha'_L \bar{L} + (1-\alpha'_L)\underline{L}. \end{aligned}$$

By the Independence Axiom,

$$\beta L + (1-\beta)L' \sim \beta\left[U(L)\bar{L} + (1-U(L))\underline{L}\right] + (1-\beta)\left[U(L')\bar{L} + (1-U(L'))\underline{L}\right].$$

Rearranging terms, we have

$$\begin{aligned} \beta L + (1-\beta)L' &\sim [\beta U(L) + (1-\beta)U(L')]\bar{L} + \left[\beta(1-U(L)) + (1-\beta)\left(1-U(L)'\right)\right]\underline{L} \\ &= [\beta U(L) + (1-\beta)U(L')]\bar{L} + [1-\beta U(L) + (\beta-1)U(L')]\underline{L}. \end{aligned}$$

By step (4), this expression can be written as

$$\begin{aligned} & [\beta\alpha_L + (1-\beta)\alpha_{L'}]\bar{L} + [1-\beta\alpha_L + (\beta-1)\alpha'_L]\underline{L} \\ =\ & \beta\left(\alpha_L\bar{L} + (1-\alpha_L)\underline{L}\right) + (1-\beta)\left(\alpha_{L'}\bar{L} + (1-\alpha'_L)\underline{L}\right) \\ =\ & \beta U(L) + (1-\beta)U(L'). \end{aligned}$$

This establishes that a utility function that satisfies continuity and the Independence Axiom, has the expected utility property: $U(\beta L + (1-\beta) L') = \beta U(L) + (1-\beta) U(L')$

[End of optional self-study section.]

4 The Expected Utility property

- The key to our model of risk preference is the assumption that preferences over risk can be described by the so-called Von Neumann-Morgenstern (VNM) Expected Utility Property (derived formally in the optional material above).

- Preferences that satisfy VNM Expected Utility theory have the property:

$$U(\beta L + (1-\beta) L') = \beta U(L) + (1-\beta) U(L'),$$

  where L and L' are bundles with $L \neq L'$ and $\beta \in (0,1)$.

- This equation says that for a person with VNM preferences, the utility of consuming two bundles L and L' with probabilities β and $(1-\beta)$, respectively, is equal to β times the utility of consuming bundle L plus $(1-\beta)$ times the utility of consuming bundle L'. Thus, the utility function is *linear* in probabilities though *not necessarily linear* in preferences over the bundles. [Note: VNM does *not* imply that $U(2L) = 2 \times U(L)$. As we'll see below, that equation would *only* hold for risk neutral preferences.]

- A person who has VNM EU preferences over lotteries will act as if she is maximizing *expected utility*—a weighted average of utilities of each state, where weights equal probabilities.

- If this model is correct, then we don't need to know exactly how people feel about risk *per se* to make strong predictions about how they will optimize over risky choices.

- [If the model is not entirely correct—which it surely is not—it may still provide a useful description of the world and/or a normative guide to how one should analytically structure choices over risky alternatives.]

- To use this model, two ingredients are needed:

  1. First, a utility function that assigns bundles an ordinal utility ranking. Note that such functions are defined up to an affine (i.e., positive linear) transformation. This means they are required to have more structure (i.e., are more restrictive) than standard consumer utility functions, which are only defined up to a monotone transformation.

- These utility functions will capture consumer preferences for risk as well as consumer preferences for bundles.

2. Second, the VNM assumptions. These make strong predictions about the maximizing choices consumers will take when facing risky choices (i.e., probabilistic outcomes) over bundles, which are of course ranked by this utility function.

   - These assumptions are discussed at length above, but these assumptions guarantee that consumers have well behaved preferences over lotteries, similarly to how the axioms of consumer choice guaranteed that consumers have well-defined preferences over bundles of good. Specifically, we assume that consumer preferences over lotteries are continuous (i.e. small changes in probabilities do not change the ordering of preferences) and independent (i.e. if we mix each of two lotteries with a third one, then the preference ordering of the two resulting mixed lotteries does not depend on (is *independent of*) the particular third lottery used)

- It's important to clarify now that expected utility theory does *not* replace consumer theory, which we've been developing all semester. Expected utility theory extends the model of consumer theory to choices over risky outcomes. Standard consumer theory continues to describe the utility of consumption of specific *bundles*. Expected utility theory describes how a consumer might select among risky bundles.

5 Expected Utility Theory and Risk Aversion

- We started off to explain risk aversion. What we have done to far is lay out expected utility theory, which is a set of (relatively restrictive) axioms about how consumers make choices among risky bundles.

- Where does risk aversion come in? It is going to come in with the specification of the utility function

- Consider the following three utility functions characterizing three different expected utility maximizers:

-
$$\mathbf{u}_1(\mathbf{w}) = \mathbf{w}$$

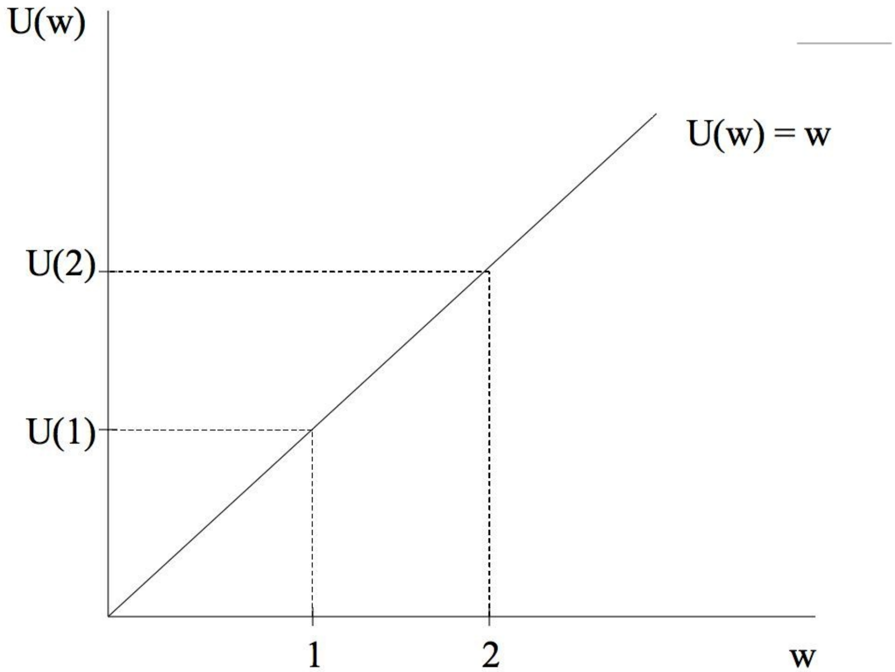

- $$\mathbf{u}_2(\mathbf{w}) = \mathbf{w}^2$$

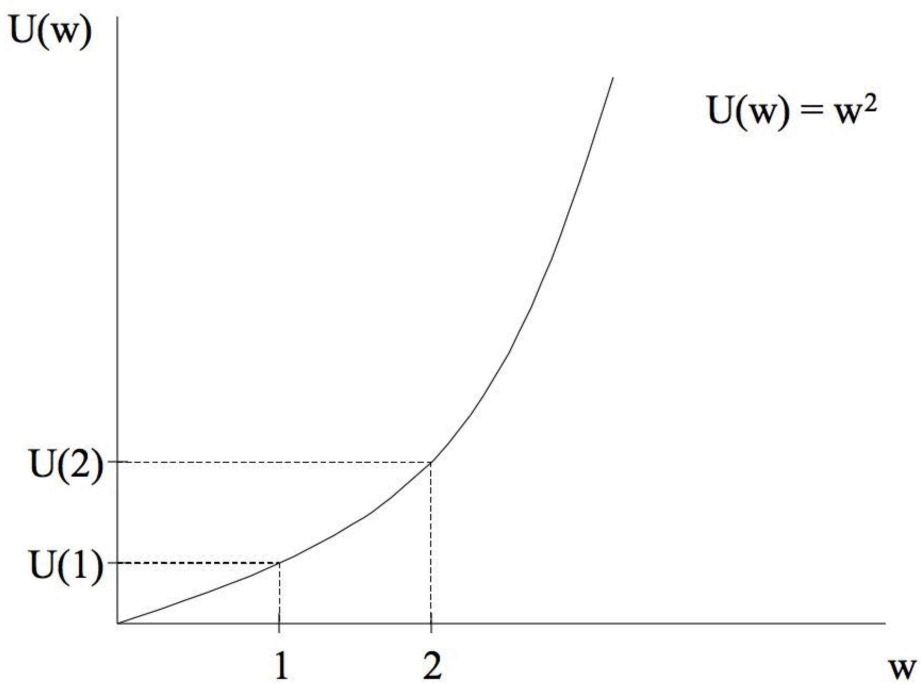

- $$\mathbf{u}_3(\mathbf{w}) = \sqrt{w}$$

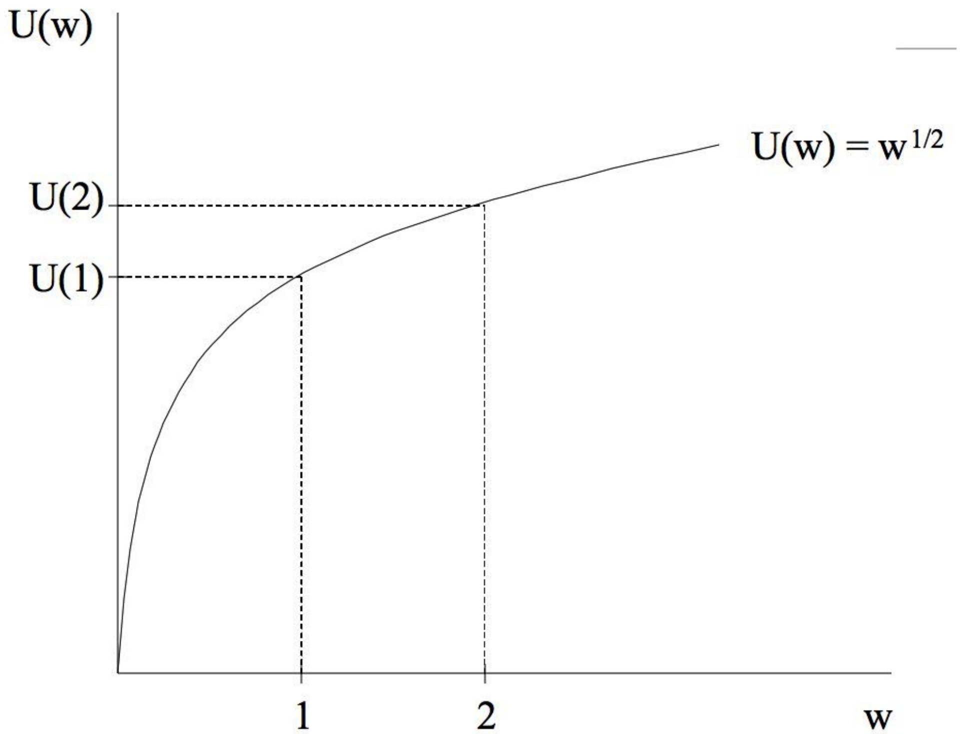

- Consider a lottery where the consumer faces 50/50 odds of either receiving two dollars or zero dollars. The expected monetary value of this lottery is \$1.
- How do these three consumers differ in their risk preference?
- First notice that $u_1(1) = u_2(1) = u_3(1) = 1$. That is, they all value *one dollar with certainty* equally.
- Now consider the *Certainty Equivalent* for a lottery L that is a 50/50 gamble over \$2 versus \$0. The certainty equivalent is the amount of cash that the consumer would be willing to accept with certainty in lieu of facing lottery L.
  - Step 1: What is the expected utility value?
    1. $u_1(L) = .5 \cdot u_1(0) + .5 \cdot u_1(2) = 0 + .5 \cdot 2 = 1$
    2. $u_2(L) = .5 \cdot u_2(0) + .5 \cdot u_2(2) = 0 + .5 \cdot 2^2 = 2$
    3. $u_3(L) = .5 \cdot u_3(0) + .5 \cdot u_3(2) = 0 + .5 \cdot 2^{.5} = .71$
  - Step 2: What is the "Certainty Equivalent" of lottery L for these three utility functions—that is, the cash value that the consumer would take in lieu of facing these lotteries? To find this, we calculate the dollar value that gives the consumer the same utility as the lottery.

1. $CE_1(L) = U_1^{-1}(1) = \$1.00$
2. $CE_2(L) = U_2^{-1}(2) = 2^{.5} = \$1.41$
3. $CE_3(L) = U_3^{-1}(0.71) = 0.71^2 = \$0.51$

- Depending on the utility function, a person would pay \$1, \$1.41, or \$0.51 dollars to participate in this lottery.
- Although the expected monetary value $E(V)$ of the lottery is \$1.00, the three utility functions value it differently:

1. The person with U_1 is *risk neutral*: $CE = \$1.00 = E(Value) \Rightarrow$ Risk neutral
2. The person with U_2 is *risk loving*: $CE = \$1.41 > E(Value) \Rightarrow$ Risk loving
3. The person with U_3 is *risk averse*: $CE = \$0.50 < E(Value) \Rightarrow$ Risk averse

- *What gives rise to these inequalities is the shape of the utility function. Risk preference comes from the concavity/convexity of the utility function*:
- Expected utility of wealth: $E(U(w)) = \sum_{i=1}^{N} p_i U(w_i)$
- Utility of expected wealth: $U(E(w)) = U\left(\sum_{i=1}^{N} p_i w_i\right)$
- Jensen's inequality:
  - $E(U(w)) = U(E(w)) \Rightarrow$ Risk neutral
  - $E(U(w)) > U(E(w)) \Rightarrow$ Risk loving
  - $E(U(w)) < U(E(w)) \Rightarrow$ Risk averse
- So, the core insight of expected utility theory is this: *For a risk averse consumer facing an uncertain set of possible wealth levels,* ***the expected utility of wealth is less than the utility of expected wealth****.*

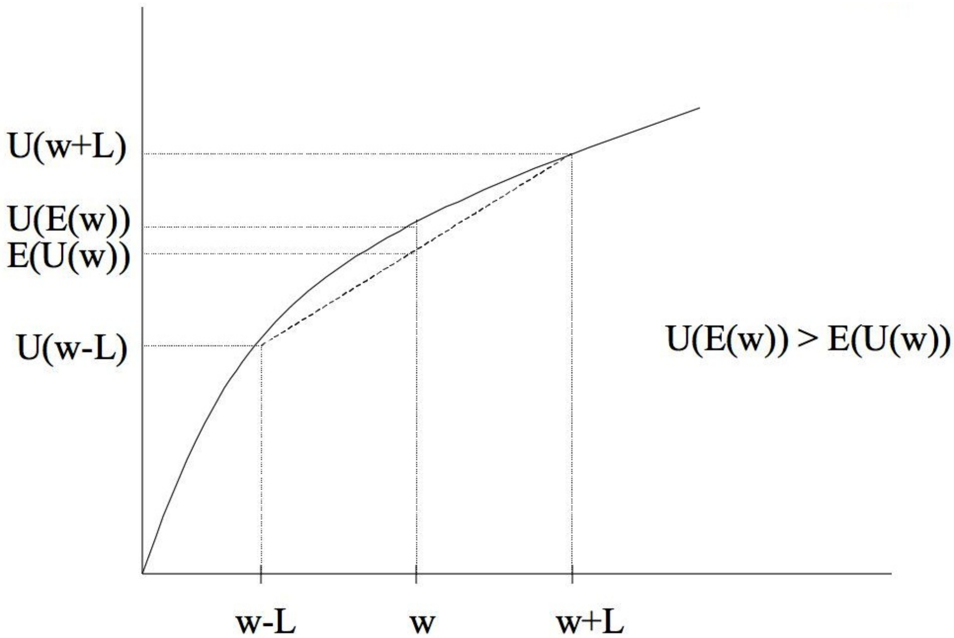

- The reason this is so:
  - If wealth has diminishing marginal utility (as is true if $U(w) = w^{1/2}$), losses cost more utility than equivalent monetary gains provide. This can be seen in the concave shape of the utility function - for lower levels of wealth, the utility function has a steeper slope.
  - Consequently, a risk averse consumer is better off receiving a given amount of wealth *with certainty* than the same amount of wealth *on average* but with variance around this quantity.

6 Conclusions

You may be thinking that we have used a lot of machinery to build a pretty modest conceptual widget. So far, you would be right. But the tool of Expected Utility Theory will prove quite powerful in the lectures to follow. We will first use VNM EU Theory to formally model people's willingness to pay to defray (avoid/reduce) risk. We will then analyze markets for risk, and glimpse the potential for insurance markets to generate Pareto-improving trades among economic agents who each possess identical bundles, preferences, and technologies. We will also analyze implicit markets for risk, as seen in the case of speed limits and the so-called value of a statistical life (VSL). Risk and insurance will then prove foundational in our study of imperfect information in markets, which is the final broad topic of the semester.