5 Expected Utility Theory

We have been talking about arbitrage models in discrete time. Now we are going to begin talking about utility-based models in discrete time. In this section of the notes, we review some results from the economics of uncertainty. We are going to say that people maximize expected utility subject to budget constraints. This material is covered in several places, including Varian's chapter 11. It should be review for most of you so we will cover it fairly quickly.

5.1 Expected Utility

A consumer's expected utility function is not a primitive in economics. We make assumptions about an agent's preferences in order to derive an expected utility function for him or her. We define expected utility over the space of lotteries, L. Using Varian's notation, $(p \circ x \oplus (1-p) \circ y)$ means receiving x with probability p and y with probability $(1-p)$. The operator $\sim$ implies indifference while $\succeq$ implies weak preference. We assume:

1. Getting a prize with probability = 1 is the same as getting the prize for certain. $(1 \circ x \oplus (1-1) \circ y \sim x)$

2. The consumer doesn't care about the order in which the lottery is described. $(p \circ x \oplus (1-p) \circ y \sim (1-p) \circ y \oplus p \circ x)$

3. A consumer's perception of a lottery depends only on the net probabilities in the lottery, not on how the lottery is packaged. $(q \circ (p \circ x \oplus (1-p) \circ y) \oplus (1-q) \circ y \sim (qp) \circ x \oplus (1-qp) \circ y)$

4. Consumers' preferences over lotteries are:

   - complete, (either $x \succeq y$ or $y \succeq x$ or both $\forall x, y$)

- reflexive, ($x \succeq x \; \forall x$)
- and transitive. (if $x \succeq y$ and $y \succeq z$ then $x \succeq z$)

5. Preferences are continuous. ($\{p \in [0;1] : p\delta_x \oplus (1-p)\delta_y \succeq z\}$ and $\{p \in [0;1] : z \succeq p\delta_x \oplus (1-p)\delta_y\}$ are closed sets for all x; y; and z in L.

6. If people are indifferent about two goods they will be indifferent about lotteries over those goods. ($x \sim y \Rightarrow p\delta_x \oplus (1-p)\delta_z \sim p\delta_y \oplus (1-p)\delta_z$)

Existence of Expected Utility Function. If $(L; \succeq)$ satisfy the above axioms then there is a utility function, u that ranks lotteries according to preferences,

$$\begin{aligned} p\delta_x \oplus (1-p)\delta_y \succ q\delta_w \oplus (1-q)\delta_z \;\Leftrightarrow \\ u(p\delta_x \oplus (1-p)\delta_y) > u(q\delta_w \oplus (1-q)\delta_z); \end{aligned} \tag{68}$$

and satisfies the expected utility property

$$u(p\delta_x \oplus (1-p)\delta_y) = pu(x) + (1-p)u(y): \tag{69}$$

You can read the proof of the theorem in Varian or in other references. It can also be shown that expected utility functions are unique up to an affine (or linear) transformation. These properties make expected utility maximization an extremely useful way to think about people's behavior under uncertainty.

Why do we worry about the existence of an expected utility function? Some of the attractiveness of expected utility maximization is driven by its mathematical tractability. We don't want our models to be determined by tractability alone - we want them to reflect reality as well. Several economists have proposed alternatives to the expected utility paradigm. Mark Machina has made a career out of his alternative to expected utility. He drops some of the assumptions we made above and finds something like local expected utility maximization. Kahnemann and Tversky, two psychologists, have also been trying to replace expected utility with something that is more consistent with behavior. Whenever economic models fail, it is possible that people are simply not maximizing expected utility functions like we want them to. For this reason, many economists are more comfortable assuming that there is no arbitrage in the market than assuming that all agents are maximizing expected utility somehow.

5.2 Risk aversion

What is it about expected utility that makes it so useful for ¯nance? Besides assuming that people are maximizing expected utility functions, we usually assume that their utilities make them risk averse. A risk averse person would rather take a certain amount of money than take a gamble with an expected payo® that is slightly larger than the certain amount. People can also be risk neutral or risk loving, of course.

It turns out that people with concave utility functions are risk averse. This result is expressed with an oft-used inequality

Jensen's Inequality. If $f(x)$ is a strictly concave function (like a risk-averse utility function) then $E[f(x)] < f(E[x])$.

Again, you can see the proof of this result in Varian. The intuition behind this result is what comes out of the diagram that is usually explained in intermediate economics classes.

To say more about risk aversion, we are going to have to define a risk premium. Suppose we were thinking about a random consumption bundle $\tilde{x} = \bar{x} + \tilde{\epsilon}$, where $\bar{x}$ is a constant and $\tilde{\epsilon}$ is a random variable with an expected value of zero. For now, a risk premium is defined as the value of π that makes true the statement:

$$E[u(\tilde{x})] = u(\bar{x} - \pi). \tag{70}$$

Now for a particular realization of $\tilde{\epsilon}$, we can use a taylor series expansion[4] to argue that

$$u(\bar{x} + \epsilon) \approx u(\bar{x}) + \epsilon u'(\bar{x}) + \frac{\epsilon^2}{2} u''(\bar{x}). \tag{71}$$

Therefore,

$$E[u(\tilde{x})] \approx u(\bar{x}) + \frac{\sigma_\epsilon^2}{2} u''(\bar{x}). \tag{72}$$

Furthermore, if σ_ϵ^2 is "small" then π is also small, so using a taylor expansion again,

$$u(\bar{x} - \pi) \approx u(\bar{x}) - \pi u'(\bar{x}) \tag{73}$$

which means that we can express our risk premium as

$$\pi \approx -\frac{\sigma_\epsilon^2}{2} \frac{u''(\bar{x})}{u'(\bar{x})} = \frac{\sigma_\epsilon^2}{2} R_A; \tag{74}$$

where R_A is known as the absolute risk aversion coefficient. The absolute risk aversion coefficient is a nice way to measure risk aversion. People with higher coefficients are more risk averse than people with lower coefficients.

If you replace the additive error term that we assumed above with a multiplicative

[4] $f(x) = f(a) + f'(a)(x-a) + \frac{f''(a)}{2}(x-a)^2 + \dots + \frac{f^k(a)}{k!}(x-a)^k + \int_a^x \frac{(x-t)^k}{k!} f(t)^{k+1} dt.$

error term, the measure of risk aversion that results is known as relative risk aversion

$$R_R = \bar{x}R_A = -\frac{\bar{x}u''(\bar{x})}{u'(\bar{x})}. \tag{75}$$

You can find a derivation for relative risk aversion in Varian or elsewhere. Next we are going to state (but not prove) an important theorem.

Pratt's theorem. Given 2 utility functions, u^1 and u^2, that are twice differentiable, strictly concave and increasing, the following are equivalent:

1. $R_A^1(\bar{x}) \geq R_A^2(\bar{x})$
2. $\rho^1(\bar{x}; \epsilon) \geq \rho^2(\bar{x}; \epsilon)$
3. u^1 is more concave than u^2.

Proofs of this theorem can be found in lots of places, including Varian. Pratt's theorem tells us three different but equivalent ways to determine if one person is more risk averse than another.

5.3 Utility Functions

There are several utility functions that are used very frequently by economists. We will discuss three of them here. The first type of utility function we will discuss is what is known as the constant relative risk aversion (CRRA) or power utility function. It is parameterized as:

$$u(x) = x^{\alpha}, \quad \alpha \in (0, 1). \tag{76}$$

As the exponent of a particular version of the power utility function goes to zero, it becomes the log utility function,

$$\lim_{\alpha \to 0} \frac{x^{\alpha} - 1}{\alpha} = \log(x), \tag{77}$$

where the logarithm in the function is a natural log. This family of utility functions is called CRRA because its coefficient of absolute risk aversion is

$$R_A = (1 - \alpha)x^{-1}; \qquad (78)$$

giving it a relative risk aversion that is constant.

A second family of utility functions that is commonly used in research is the constant absoulte risk aversion family (CARA). This family is parameterized

$$u(x) = 1 - e^{-\lambda x}; \quad \lambda > 0. \qquad (79)$$

The absolute risk aversion coefficient for this utility function is just λ.

The last type of utility function we will discuss is the quadratic utility function. This function is written

$$u(x) = a + bx - cx^2; \quad b, c > 0. \qquad (80)$$

You can calculate the risk aversion coefficients for this utility function as a homework assignment. A quadratic utility function looks like an inverted parabola. There is always a point at which marginal utility, $u'(x)$, becomes negative. You get to solve for this as a homework problem as well.

5.4 Stochastic Dominance

We have talked about ways to determine whether one person is more or less risk averse than another person. Now we will shift our emphasis to asking whether a particular lottery is more or less risky than another lottery. Probably the most general way to compare the risk of lotteries is in terms of what is called stochastic dominance.

First Order Stochastic Dominance. The cumulative distribution of payoffs F first order stochastic dominates (FOSD) G iff $G(x) \geq F(x)\ \forall x \in I$, where I is the sample space of x.

First order stochastic dominance is an attractive property because it has been shown that, for all increasing utility functions, $u(x)$,

$$F \text{ FOSD } G \iff E_F[u(x)] \geq E_G[u(x)]; \tag{81}$$

where E_F is the expectation taken under the assumption that F is the distribution of payoffs. To interpret FOSD, remember that $G(x) = \Pr(\tilde{x} \leq x)$ and draw a picture:

A weaker concept than FOSD is second order stochastic dominance (SOSD). To define SOSD, we need to define the function

$$T(x) = \int_I [G(x) - F(x)]dx \tag{82}$$

Second Order Stochastic Dominance. F SOSD G iff $T(x) \geq 0\ \forall x \in I$ and $E_G[x] = E_F[x]$.

Second order stochastic dominance is a weaker concept than FOSD in the sense that FOSD implies SOSD but SOSD does not imply FOSD. For all increasing and strictly concave utility functions,

$$F \text{ SOSD } G \Leftrightarrow E_F[u(x)] \geq E_G[u(x)]. \tag{83}$$

Since SOSD is a weaker concept than FOSD, we need the additional condition that $u(x)$ is concave to get the result that people should prefer payoffs that second order dominate.

Second order stochastic dominance is an attractive property to work with because it corresponds to a frequently used abstraction known as a mean preserving spread. Economists often add a mean preserving spread to their models in order to introduce uncertainty. They are sometimes described as a “sprinkling” of risk. If we define the random variable $\tilde{y}$ as $\tilde{x}$ plus a mean preserving spread then

$$\tilde{y} = \tilde{x} + \tilde{\varepsilon}, \tag{84}$$

where

$$\begin{aligned} E[\tilde{\varepsilon}] &= 0 \\ E[\tilde{\varepsilon} | x] &= 0 \\ Var[\tilde{\varepsilon}] &> 0 \end{aligned} \tag{85}$$

In this case, the distribution of $\tilde{x}$ will second order stochastic dominate the distribution of $\tilde{y}$. A plot of a mean preserving spread can be instructive:

A useful theorem for interpretting SOSD is the Rothschild-Stiglitz theorem.

Rothschild-Stiglitz Theorem. The following conditions are equivalent:

1. F SOSD G.

2. G = F plus noise (mean preserving spread)

3. F and G have the same mean and all risk averters prefer F to G.

5.5 Homework Problems

1. Solve for the absolute risk aversion coefficient and the relative risk aversion coefficient for the quadratic utility function, (80). Solve also for the point at which utility is maximized if there is no constraint on consumption, x. Assuming that there exists some maximum possible value of x, what assumption could you make to rule out satiated consumers.

2. (Number 11.6 from Varian) Esperanza has been an expected utility maximizer ever since she was five years old. As a result of the strict education she received at an obscure British boarding school, her utility function u is strictly increasing, strictly concave, twice differentiable and bounded. Now, at the age of thirty-something, Esperanza is evaluating an asset with stochastic outcome R which is normally distributed with mean μ and variance σ^2. Thus, its density function is given by

$$f(R) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left[-\frac{1}{2} \left(\frac{R-\mu}{\sigma} \right)^2 \right]. \quad (86)$$

   - Show that Esperanza's expected utility from R is a function of μ and σ^2 alone. Thus, show that $E[u(R)] = \phi(\mu; \sigma^2)$.
   - Show that $\phi(\cdot)$ is increasing in μ.
   - Show that $\phi(\cdot)$ is decreasing in σ^2.

   Hint: define a new variable $\varepsilon = (R - \mu)/\sigma$, write down the expression for $E[u(\varepsilon)]$ and sign the derivatives - you may need integration by parts.

3. (Number 11.7 from Varian) Let R_1 and R_2 be the random returns on two assets. Assume that R_1 and R_2 are independently and identically distributed. Show that an expected utility maximizer will divide her wealth between both assets provided she is risk averse; and invest all her wealth in one of the assets if she's risk loving.