Lecture 13: Choice Under Uncertainty

1 Where Are We/Logistics

- Last time, we wrapped up consumer theory
- Today, we move on to choice under uncertainty

- today and next Tuesday are my final two lectures,
  then Lones will be taking over
- The last homework for my half, on choice under uncertainty,
  which actually be due TUESDAY November 2,
  since many of you have a macro exam Monday;
  and your exam for this class is Thursday evening, Nov 4, two weeks from tonight,
  in Ingraham 19

- so far, we've been thinking about our choice set X as being different consumption bundles –
  even if it's multi-dimensional, we're still choosing a single combination of goods to consume
- now we're going to shift gears and start thinking about preferences over *uncertain* outcomes

- Before we start... any questions?

2 Choice Under Uncertainty

- Everything we've done so far in consumer theory has been done under the assumption that you know exactly what is going to happen
- You don't have control over prices, but you know them;

  you know the exact utility you'll get from each combination of goods;

  and you deterministically select the bundle that you like the most

- But obviously, lots of important decisions in life are made with less certainty than that
- The mundane: you decide whether to carry an umbrella, not knowing for sure if it will rain
- The life-altering: you choose a grad school, or a job, or a partner,

  trying to make the best decision you can with the information you have,

  but knowing it's not certain which one will turn out to be the best possible match for you

- Over our final two lectures, we'll build up some theory about decision-making

  when there is uncertainty about the outcome you'll get

- With consumer theory, you may recall our strategy:

  we started with preferences as a primitive,

  described the fairly reasonable assumptions under which they would lead to someone behaving as if maximizing a numerical utility function,

  and then explored the consequences of utility maximization

- With choice under uncertainty, we'll follow a similar strategy:

  we'll start off with preferences over portfolios of possible outcomes,

  without assuming those preferences have a particular structure;

  add some reasonable-looking assumptions,

  and show that implies they can be represented by a particular type of utility function;

  and then explore the consequences of maximizing that type of utility function

- A couple things I'll mention before we get started

2.1 objective probability

- We'll be focusing on uncertainty with **objective** probabilities –

  how you value the possibility of winning \$100 with probability 40%,

  when you *know* that the probability of winning truly is exactly 40%

- Not all probabilities are like that

- We could think about how you would value winning \$100 if it rains on Sunday,

  or if Trump gets re-elected in 2024,

  or if a vaccine-resistant COVID strain emerges before Christmas

- These are different, because these are uncertain events where there isn't a clear, known probability of each outcome

- There's a similar, complementary theory for that type of uncertainty,[1]

  which we call "subjective" uncertainty, and the results are fairly similar,

  but we'll be sticking to objective probabilities

2.2 behavioral critiques

- Probably more than any other area of mainstream economics,

  expected utility maximization has come under critique from behavioral economics

- There's lots of experimental and field evidence

  that seems to be in conflict with the results we'll be talking about

- I'll talk about this very briefly next week,

  and give a rather standard defense of what we're doing

  but it's not something we'll go deep into

- You should just be aware that,

  while every model is a simplification and can't be expected to perfectly predict behavior in every instance,

  expected utility theory is probably more often visibly violated than the other parts of consumer theory we've done so far

- and with that, let's get started

[1] see Savage (1954), *The Foundations of Statistics* ch. 2-3, and Anscombe and Aumann (1963), "A Definition of Subjective Probability," *Annals of Mathematical Statistics* 34.1. These models are more complicated, but largely boil down to the same idea: if you have preferences over combinations of probabilistic outcomes that satisfy certain axioms, then it's as if you're maximizing the expected value of the utility of the realized outcome, taken with respect to some probabilities that we can interpret as the subjective probabilities you assign to the outcomes.

3 Lotteries and Von Neumann-Morgenstern Utility

- To begin, we'll suppose there are just a **finite** number of different possible outcomes, or *prizes*, $\{x_1, x_2, \ldots, x_N\}$
- (For now, these could be different amounts of money or different actual goods – x_1 is \$50, x_2 is a sports car, x_3 is a pony – although it's typical to think of them as just being different amounts of money)
- we'll define a *lottery* as a probability distribution over these outcomes – a lottery L is a probability p_1 of ending up with x_1, p_2 of ending up with x_2, and so on, with $p_i \geq 0$ for each i and $\sum_i p_i = 1$
- What we're driving towards is what's called an *expected utility representation*: a function u defined on the prizes, such that

$$U(L) \quad = \quad \sum_i p_i u(x_i)$$

  that is, so your preferences over lotteries are represented by a utility function which is the expected value of the utility you put on the prize you win
- But we're not going to *assume* that preferences take that form; we're going to derive it as a result
- We're going to assume you have complete and transitive preferences over lotteries, make a couple more assumptions about your preferences over lotteries, and get the *result* that your preferences match a utility function of this form

- To do this, we will need to make a couple additional assumptions;

  and to do this, we'll need a bit more language

- First of all, we need to define compound lotteries

- If L is a lottery giving a 50-50 shot at winning \$100 or nothing,

  and L' is a lottery giving a 50-50 shot at winning \$200 or nothing,

  then we can consider the *compound lottery* where you have, say, a 50% chance of getting lottery L and a 50% chance of getting lottery L',

  and this is just

$$L'' \quad = \quad \frac{1}{2}\left(\frac{1}{2}(0) \oplus \frac{1}{2}(\$100)\right) \oplus \frac{1}{2}\left(\frac{1}{2}(0) \oplus \frac{1}{2}(\$200)\right) \quad = \quad \frac{1}{2}(0) \oplus \frac{1}{4}(\$100) \oplus \frac{1}{4}(\$200)$$

- That is, you only care about the probability of each final outcome,

  so we can think about lotteries whose prizes are other lotteries, called compound lotteries,

  just the same as *simple lotteries* (lotteries whose prizes are certain outcomes)

- The fundamental property of preferences that have an expected utility representation is that

$$U(pL \oplus (1-p)L') \quad = \quad pU(L) + (1-p)U(L')$$

  whether L and L' are certain outcomes or lotteries themselves,

  but again, we're not assuming that,

  we want to get that as a result

- We assume your preferences over the set of all *lotteries* is complete and transitive,

  and we also make two additional assumptions

3.1 Continuity

- **continuity**: this can be defined a few different ways,
  which wind up being equivalent
- I think the most useful way is to say that for any lotteries L, L', and L'',
  the two subsets of $[0, 1]$ defined by
  $$\left\{p \ : \ pL \oplus (1-p)L' \ \succsim \ L''\right\}$$
  and
  $$\left\{p \ : \ pL \oplus (1-p)L' \ \precsim \ L''\right\}$$
  are both closed[2]
- That is, if we have a sequence of probabilities $p^1, p^2, \cdots, \to \bar{p}$,
  and if $p^j L \oplus (1-p^j)L' \succsim L''$ for each j,
  then $\bar{p}L \oplus (1-\bar{p})L' \succsim L''$
- This looks a lot like continuous preferences the way we've defined them before,
  where preferences were continuous in the bundle you consumed,
  rather than the probabilities of the different outcomes
- which may make it feel innocuous,
  but it can actually feel a bit strong when applied to lotteries
- here's why:

- Continuity implies that for any three lotteries $L \succ L' \succ L''$,
  there's some probability $p \in (0, 1)$ such that $pL \oplus (1-p)L'' \sim L'$
- Suppose there are three prizes – winning \$10, nothing, and death
- This says that if you prefer \$1 to nothing,
  then you prefer a high probability of \$1, plus a small risk of death, to nothing
- Maybe that really reflects your preferences, maybe not;
  but it's worth emphasizing that continuity here really does have content –
  it means you really do have to perceive very small probabilities as very small in significance

[2] These notes originally had only the first requirement – I apologize for the confusion

3.2 Independence

- The other condition we impose is called **independence**:

  this says that if $L \succsim L'$, then for any probability p and any other lottery L'',

  $pL + (1-p)L'' \succsim pL' + (1-p)L''$

- If I prefer sushi to steak,

  then I prefer any probability of sushi to the same probability of steak,

  regardless of what I get the rest of the time

- (There's no analog to this in regular consumer theory,

  and it wouldn't really make sense,

  since different goods could be complements or substitutes;

  here, since we're talking about probabilities of different outcomes,

  you never get to consume more than one of the prizes,

  so there's no notion of complementarities to worry about)

- one of the things this implies is that you don't get disutility just from uncertainty itself

- for example, if $L \succsim L'$, this means that any mix $pL \oplus (1-p)L' \succsim L'$,

  so you're always happy to take on any probability of getting a better outcome –

  the chance of the better outcome is all you care about,

  not the fact that you're now unsure of what you'll get

3.3 expected utility theorem

- These properties are all we need for the big result

- **Prop. (von Neumann and Morgenstern, 1947)** Suppose a consumer has preferences over lotteries which are complete and transitive and satisfy continuity and independence. Then they are represented by a utility function of the form

$$U(L) \quad = \quad \sum_i p_i u(x_i)$$

where u assigns a constant utility value to each prize.

- Actually, this is often stated as if-and-only-if –

  if you have a utility function of this form,

  then your preferences satisfy continuity and independence;

  and if they satisfy continuity and independence, they have a utility function of this form

- We'll skip the proof – it's in MWG (pp 176-178) if you're interested,

  but the brief outline:

  - It's simplest if we assume that there's a best lottery $\overline{L}$ –

    a lottery you weakly prefer to all others –

    and a worst lottery $\underline{L}$ that you weakly prefer all other lotteries to
  - (This isn't completely inoccuous – there's a finite number of *prizes*,

    but an infinite number of potential lotteries,

    and we haven't imposed a lot of structure on your preferences yet –

    but it keeps the proof simple)
  - We can then assign utility of 0 to the worst lottery $\underline{L}$, 1 to the best lottery $\overline{L}$,

    and for any other lottery L, find the number p such that $L \sim p\overline{L} \oplus (1-p)\underline{L}$,

    and assign the lottery utility of p
  - Such a p exists by continuity; we'd then show it's unique,

    that this utility function actually represents preferences,

    and that it takes the expected utility form

3.4 u is unique up to affine transformations

- When we proved a utility representation for general preferences,

  we noted that any monotone transformation of that utility function was another valid one

- Here, that's not the case at all:

- **Prop.** Suppose U is an expected utility representation of preferences over lotteries. Then V is also an expected utility representation of the same preferences if and only if there are constants $a \in \mathbb{R}$ and $b \in \mathbb{R}_{++}$ such that $V(\cdot) = a + bU(\cdot)$.

- That is, expected utility representations are uniquely determined up to positive linear transformations

4 Switch over to $X \subseteq \mathbb{R}$

- So far, we've been dealing with lotteries over a finite "prize space" –

  say, lotteries where you could win nothing, or \$10, or \$50, but not every amount in between

- Next, we're going to switch over to thinking of the prize space being $\mathbb{R}$

  so that lotteries are now probability distributions over $\mathbb{R}$

- A lottery could still be a distribution over a finite number of different prizes,

  each with positive probability,

  or be a continuous distribution over outcomes

- If a lottery is represented by a CDF F,

  an expected utility representation would give

$$U(F) = \int u(x)dF(x) = \int u(x)f(x)dx$$

- The function $u(\cdot)$ is referred to as the Bernoulli utility function

- We'll always assume u is strictly increasing – more money is strictly better

5 Risk Aversion

- Let δ_x be degenerate lottery putting all weight on a point x.[3]
- A decision-maker is **risk averse** if for any lottery F with expected value E_F,

$$\delta_{E_F} \succsim F$$

  and **strictly risk averse** if this preference is strict whenever F is non-degenerate

- This means a decisionmaker is risk-averse if

$$u\left(\int x dF(x)\right) \quad \geq \quad \int u(x) dF(x)$$

  which is exactly Jensen's inequality, the fundamental property of concave functions;

  so an expected-utility-maximizer is risk-averse if and only if u is concave

- similarly, an expected-utility maximizer is **risk-neutral** if u is linear,

  in which case they just maximize the expected value of the lottery;

  and an expected-utility maximizer is **risk-loving** if u is convex,

  in which case they prefer an uncertain outcome to a certain outcome with the same EV

- economists assume u is strictly increasing,

  and focus almost exclusively on the cases of risk-neutral or risk-averse decisionmakers,

  meaning u is weakly concave

- this automatically implies u is continuous –

  all increasing, concave functions are continuous –

  and for convenience, we'll typically assume it's differentiable as well

[3] δ_x can be thought of as the Dirac delta function – the limit of a Normal density with mean x as the variance goes to zero, which has the properties that $\int \delta_x(t)dt = 1$, $\delta x(t) = 0$ for $t \neq x$, and $\int g(t)\delta_x(t)dt = g(x)$ for any function g.

- We can define the **certain equivalent** $c(F,u)$ of a lottery F as the amount of riskless money that would be equally good, i.e., by

$$\delta_{c(F,u)} \quad \sim \quad F$$

or

$$U(\delta_{c(F,u)}) \quad = \quad U(F)$$

or

$$u(c(F,u)) \quad = \quad \int u(x)dF(x)$$

- For a risk averse decisionmaker, the certain equivalent is always weakly less than the expected value – strictly less if F is nondegenerate and risk aversion is strict

- And we can define the **risk premium** as the difference –

  that is, how much expected value someone would sacrifice to avoid the risk,

  or $E_F - c(F,u)$

- As an example, let's calculate the risk premium for a small coin-flip gamble –

  a lottery where relative to your starting wealth w, you can either win or lose some small amount y with equal probabilities

- The risk premium δ will be defined by

$$u(w-\delta) \quad = \quad \frac{1}{2}u(w+y)+\frac{1}{2}u(w-y)$$

  (since the expected value is w and the certain equivalent is $w-\delta$)

- If y is small, δ will be small as well, and we can Taylor-expand both sides, giving

$$\begin{aligned} u(w-\delta) &= \tfrac{1}{2}u(w+y)+\tfrac{1}{2}u(w-y) \\ u(w)-\delta u'(w)+\tfrac{1}{2}\delta^2u''(w) &= \tfrac{1}{2}\left(u(w)+yu'(w)+\tfrac{1}{2}y^2u''(w)\right)+\tfrac{1}{2}\left(u(w)-yu'(w)+\tfrac{1}{2}y^2u''(w)\right) \\ -\delta u'(w)+\tfrac{1}{2}\delta^2u''(w) &= \tfrac{1}{2}y^2u''(w) \\ \delta-\tfrac{1}{2}\delta^2\tfrac{u''(w)}{u'(w)} &= \tfrac{1}{2}y^2\left(-\tfrac{u''(w)}{u'(w)}\right) \end{aligned}$$

- For δ small, δ^2 is even smaller, so for y small, the second term vanishes, giving

$$\delta \quad \approx \quad \frac{1}{2}y^2\left(-\frac{u''(w)}{u'(w)}\right)$$

- What do we see?
- first, the risk premium is proportional to the *square* of the size of the gamble – as gambles get small, our aversion to them gets not just small, but very small
- (Another way to think about this: as we "zoom in" on a concave function, it looks linear as we get closer)

- and second, the risk premium is proportional to the ratio $\frac{-u''}{u'}$
- This ratio is called the **Arrow-Pratt Coefficient of Absolute Risk Aversion**
- It's an excellent local measure of "how risk-averse" someone is, and is often written as $A(u, w)$ or $A(w)$.

- we could analogously define the **probability premium** as the probability advantage you would require to accept an even-money bet,

$$u(w) \quad = \quad \left(\frac{1}{2} + \epsilon\right) u(w+y) + \left(\frac{1}{2} - \epsilon\right) u(w-y)$$

and we would find that for small y,

$$\epsilon \quad \approx \quad \frac{1}{4} y \left(-\frac{u''(w)}{u'(w)}\right)$$

- so the advantage we would need to accept a small bet is also proportional to this Coefficient of Absolute Risk Aversion

5.1 Who's more risk averse?

- the Arrow-Pratt coefficient of absolute risk aversion, $A = -\frac{u''}{u'}$,
  is a nice local measure of "how concave" your utility function is,
  and therefore "how risk-averse" you are

- (why not just u''? we said expected utility was defined up to an affine transformation –
  but if we multiply u by a constant, u'' scales up, but u''/u' does not,
  so at least we know it's uniquely defined given vN-M preferences)

- what's nice is that its significance extends beyond small gambles, to large ones;
  its significance is summed up well in Pratt's theorem:

Theorem (Pratt, 1964). *Let u and v be two Bernoulli utility functions. The following definitions of "u is more risk averse than v" are equivalent:*

1. *For any lottery F and certain outcome δ_x, if $F \succsim_u \delta_x$ then $F \succsim_v \delta_x$*
2. *For every F, $c(F,u) \leq c(F,v)$*
3. *$u = g \circ v$ for some increasing, concave function g*
4. *for every x, $A(u,x) \geq A(v,x)$, or $-\frac{u''(x)}{u'(x)} \geq -\frac{v''(x)}{v'(x)}$*

- 1 and 2 are obviously equivalent, since $F \succsim \delta_x$ if and only if $c(F,u) \geq x$,
  so if u gives a lower certain equivalent than v,
  it means u prefers a wider range of certain outcomes to the same lottery F

- to see why 2 and 3 are equivalent, note that since u is increasing,

$$\begin{array}{rcl} c(F,u) & \leq & c(F,v) \\ & \updownarrow & \\ u(c(F,u)) & \leq & u(c(F,v)) \\ & \updownarrow & \\ \int u(x)dF(x) & \leq & u(c(F,v)) \end{array}$$

  (by definition of $c(F,u)$).

- We can always find a function g such that $u = g \circ v$,

  just by defining $g(x) = u(v^{-1}(x))$ (since v is strictly increasing and therefore invertible);

  plugging in $u = g \circ v$, this is

$$\begin{array}{rcl} \int g(v(x))dF(x) & \leq & g\left(v(c(F,v))\right) \\ & \updownarrow & \\ \int g(v(x))dF(x) & \leq & g\left(\int v(x)dF(x)\right) \\ & \updownarrow & \\ E(g(Y)) & \leq & g\left(E(Y)\right) \end{array}$$

  where $Y = v(X)$; this is Jensen's Inequality, and holds if and only if g is concave.

- and to see why 3 and 4 are equivalent, since u and v are strictly increasing,

  differentiating $u(x) = g(v(x))$ gives $u'(x) = g'(v(x))v'(x)$,

  and again gives $u''(x) = g'(v(x))v''(x) + g''(v(x))(v'(x))^2$, so

$$-\frac{u''}{u'} \quad = \quad -\frac{g'v''}{g'v'} - \frac{g''(v')^2}{g'v'} \quad = \quad -\frac{v''}{v'} - \frac{g''}{g'}v'$$

- Since g and v are increasing, this means $-\frac{u''}{u'} \geq -\frac{v''}{v'}$ if and only if $g'' \leq 0$, or g is concave

- So, having a higher Coefficient of Absolute Risk Aversion means you're more risk-averse than someone else –

  you have a lower certain equivalent for the same risky lottery,

  and you're more prone to prefer a sure prize to a lottery

5.2 Risk preferences and wealth

- We can also use the Pratt result to consider how risk preferences change with wealth level

- We say a Bernoulli utility function has **decreasing absolute risk aversion** if $-\frac{u''(x)}{u'(x)}$ is decreasing in x;

  constant or increasing absolute risk aversion if it's constant or increasing

- Decreasing absolute risk aversion means that as you add wealth –

  or add a constant to all payoffs in a lottery –

  you behave as if you're less risk-averse

- For example, let $L + z$ denote the lottery L, but with a constant z added to every prize

- **Prop.** Suppose $-\frac{u''(x)}{u'(x)}$ is decreasing.

  For any lottery L and constants y and $z > 0$, if $L \succsim \delta_y$, then $L + z \succsim \delta_{y+z}$.

- The proof is that we can evaluate $L + z$ and δ_{y+z} as if we were comparing L to δ_y,

  just using a different Bernoulli utility function, $v(x) = u(x + z)$

- That is, if we let $v(x) = u(x + z)$ and F be the distribution of outcomes under L,

  then $U(L + z) = \int u(x + z)dF(x) = \int v(x)dF(x)$,

  and $U(\delta_{y+z}) = u(y + z) = v(y)$

- So we want to know whether $\int v(x)dF(x) \geq v(y)$,

  and we already know that $\int u(x)dF(x) \geq u(y)$.

- But if A is decreasing in x, then $A(v, x) \leq A(u, x)$ at every x,

  because $-\frac{v''(x)}{v'(x)} = -\frac{u''(x+z)}{u'(x+z)} \leq -\frac{u''(x)}{u'(x)}$,

  so we can think of "richer you" as a different decisionmaker, who's less risk-averse

- so if you have decreasing absolute risk aversion, as you get wealthier,

  your certain equivalents get closer to the expected value of a gamble,

  or you demand a smaller risk premium for a given gamble,

  or a smaller probability premium, and so on

- If you have increasing absolute risk aversion, we'd get the opposite result –

  as you get wealthier, you become *less* tolerant of risk

- (And of course, if your coefficient of absolute risk aversion is *strictly* increasing or decreasing,

  this all holds strictly –

  with strictly increasing absolute risk aversion, if you're indifferent between a lottery L and a sure result δ_x,

  then if I added some money to every outcome, you'd now *strictly* prefer the sure outcome)

- there's also a family of utility functions with *constant* absolute risk aversion

$$u(x) \quad = \quad 1 - e^{-cx}$$

- we can calculate $u'(x) = ce^{-cx}$ and $u''(x) = -c^2 e^{-cx}$, so $-\frac{u''}{u'} = -\frac{-c^2 e^{-cx}}{ce^{-cx}} = c$ is constant

- with CARA utility, your risk preferences remain exactly the same as your wealth changes –

  there are no wealth effects in your preferences over lotteries

- (adding a constant to all outcomes of a lottery just adds the same constant to your certain equivalent, and doesn't change the risk premium you demand;

  or changing your starting point of wealth doesn't change your attitude to risk)

- you'll explore all this on the last homework