14.123 Microeconomic Theory III
Spring 2009

MIT 14.123 (2009) by Peter Eso
Lecture 6: Beyond EU

1. State-Dependent EU
2. Subjective EU and Ellsberg's Paradox
3. Rabin's Puzzle & Measuring riskiness

Read: Finish MWG Chapter 6, assigned readings.

Expanding Expected Utility

- In many applications the "choice over objective-probs lotteries" framework is not appropriate.
  - It may matter what causes the payoff, not just its level and probability. State-of-Nature Representation of Uncertainty.
  - Choices may be given without explicit probabilities of the outcomes. Subjective Probabilities.
- Examples:
  - Insurance: Suppose the probability of an accident is 1%. "$100 with 1% chance" ?⇔? "$100 if accident happens".
  - Betting on the Winner of the 2009 Champions' League (or next year's Super Bowl, or Best Actress at the Oscars…).

State Dependent EU

- Suppose there is a (finite) set of states, S.
  Each state $s \in S$ has probability π_s which is known (for now).
- In each state, the outcome belongs to a (finite) set X.
- The set of alternatives is a vector of objective lotteries over X in each state, $\Delta^{|S|}$, where Δ is the $(|X|-1)$ dimensional simplex.
- Compound lotteries are reduced to simple ones in each state.
- However, lotteries are not reduced further by aggregating the probability of a given outcome across states in which it occurs.

  E.g., $X = \{a,b\}$, $S = \{s_1,s_2,s_3\}$ each state has 1/3 chance.
  Lottery (a,a,b) is not reduced to $2/3 \cdot a + 1/3 \cdot b$, because payoff a in state s_1 is not the same as payoff a in state s_2.

State Dependent EU

- THM: If a preference relation $\succsim$ over $\Delta^{|S|}$ is continuous, complete, transitive, and satisfies the Independence Axiom, then it can be represented by a state-dependent expected utility function.
  A degenerate lottery $(x_1,\ldots,x_S)$ is evaluated by $\sum_s \pi_s\, u_s(x_s)$.
- This result follows directly from the EU Theorem. Instead of a single utility index u on X, here we determine a vector of utilities, $(u_1,\ldots,u_{|S|})$. Each payoff $x_i \in X$ may have a different utility index in each state (they are treated as different outcomes).
- This theorem is not particularly deep.

Application: Insurance

- Two states, S = {"no accident", "accident"}, probs π_1, π_2, resp. W/o insurance, the outcome is $(w,w\text{-}D)$, where w is the initial wealth, D the damage. Let $u_1(0) = u_2(0) = 0$, $u_1'(x) > u_2'(x)\ \forall x>0$; and both u_i' are decreasing. Fair insurance is available.

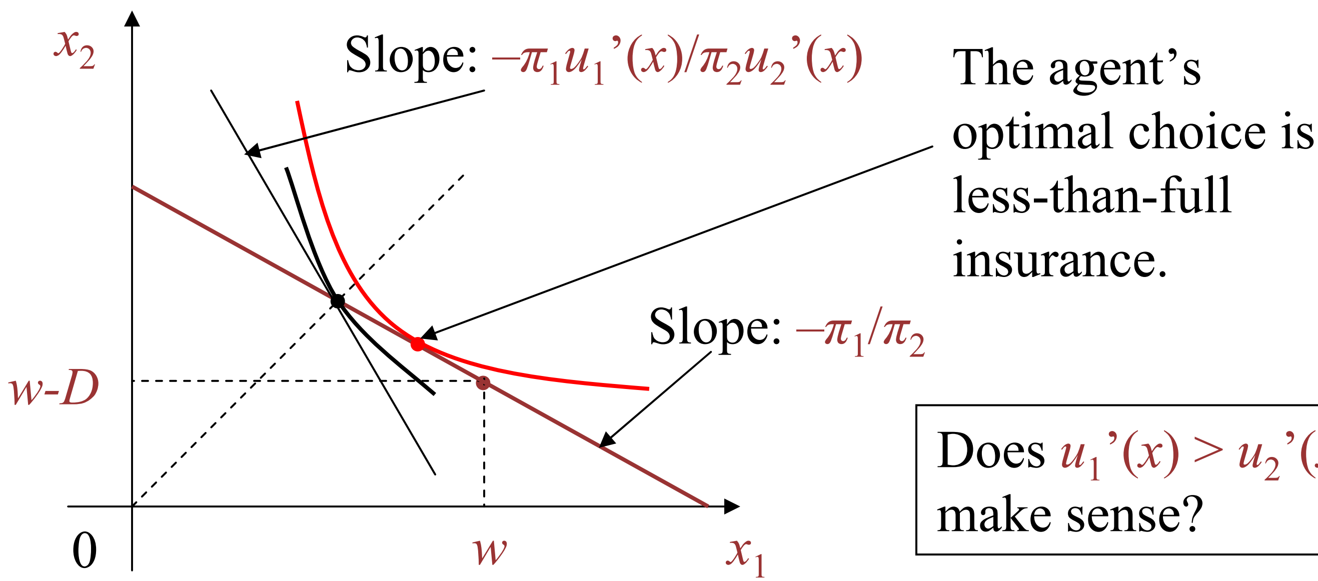

Does $u_1'(x) > u_2'(x)$ make sense?

A Subjective Framework

- In many decision problems with uncertainty, the "lotteries" we choose from do not come with objectively defined probabilities.
- Example: Bet on the winner of Best Actress at the Oscars. Five nominees (ex ante), people may disagree on the odds. They still have preferences over bets on the winner.
- Framework: There are "states of nature" ("The winner is X"). A "gamble" is a set of objective lotteries, each one corresponding to a state. Two sources of uncertainty: (a) risk within a state; (b) uncertainty over which state will occur.
- Set of alternatives: $\Delta^{|S|}$ = set of objective lotteries in each state. Lotteries are reduced in each state, but not across states.

Utility Representations

- Fix a preference relation $\succsim$ on $\Delta^{|S|}$ and assume it is continuous, complete, transitive, and satisfies the Independence Axiom.
- Suppose, for a moment, that the probability of state s is π_s .
  $\rightarrow \succsim$ has a state-dependent expected utility representation:
  $(x_1,\ldots,x_{|S|}) \succsim (y_1,\ldots,y_{|S|})$ iff $\sum_s \pi_s\, u_s(x_s) \geq \sum_s \pi_s\, u_s(y_s)$.
- Now suppose that each state has probability $\pi_s' = 1/|S|$. The same preference $\succsim$ (which is given without reference to the probabilities of the states) is represented by $V(x) = \sum_s v_s(x_s)$; s.t. $v_s(x_s) \equiv \pi_s\, u_s(x_s)$.
- The modeler does not know the decision-maker's assessment of the π_s's. But the decision-maker's behavior may reveal that assessment.

Subjective Expected Utility

- How can the modeler "tease out" the decision maker's (agent's) subjective probability assessment regarding the states?
- If the decision-maker really has state-dependent expected utility, then there is no way.

  Example: If $(\$1,0) \succsim (0,\$1)$ then either the agent thinks state s_1 is more likely than state s_2, or the agent's marginal utility for money is greater in s_1 than s_2, or both.

- Subjective EU Idea: Assume that the decision-maker has state-independent risk-preferences. His vNM utility function can be determined by how he evaluates objective lotteries (useful we kept them in the model). Then, get subjective π_s's from the preferences.

Subjective Expected Utility

- DEF: Given $\succsim$ with utility-representation $(u_1,\ldots,u_{|S|})$, the induced state-contingent preference relation is $\succsim_s$ such that for $p,q \in \Delta$, $p \succsim_s q$ iff $\sum_{x\in X} p_s(x)u_s(x) \geq \sum_{x\in X} q_s(x)u_s(x)$.
- DEF: $\succsim$ is state-uniform if $\succsim_s = \succsim_{s'}$ for all $s,s' \in S$.
- THM: Suppose that the set of alternatives is $\Delta^{|S|}$, where S is a finite set of states. If $\succsim$ is continuous, complete, transitive, state-uniform and satisfies the Independence Axiom, then there exist probability weights $(\pi_1,\ldots,\pi_{|S|})$ and utility function u such that

$$(x_1,\ldots,x_{|S|}) \succsim (y_1,\ldots,y_{|S|}) \text{ iff } \sum_s \pi_s\, u(x_s) \geq \sum_s \pi_s\, u(y_s).$$

- Richer model, more assumptions; probabilities from preferences.

Ellsberg's Paradox

- Two bags, each has two balls. The "$20 gamble" pays $20 if you guess the color of a randomly-selected ball correctly.
- Bag *A* contains 1 red and 1 green ball. How much would you pay for the "$20 gamble"? Average = $9.74, Fox and Tversky (1995).
- Bag *B* also contains two balls, but the color mix is not specified. How much would you pay for the "$20 gamble"? Avg = $8.53.
- Subjective EU $\Rightarrow$ The agent has a probability assessment π_R = Pr(Red) for each bag. No matter what π_R is for Bag *B*, one cannot do worse guessing the color of the ball in Bag *B* than guessing it in Bag *A*. Behavior is inconsistent with Subjective EU.

Ellsberg's Paradox

- Proposed Solution: Ambiguity Aversion.
- The bet on Bag B is "ambiguous": the probabilities are not given. The decision-maker cannot rule out any distribution (prior). Perhaps s/he evaluates a gamble as its minimum expected utility over all possible distributions.
- Idea dates back to Wald (1950), *Statistical decision functions*.
- Gilboa and Schmeidler (JME 1989), using subjective framework, give a set of axioms for an agent's preferences to be represented by a vNM utility function u and a set of priors P such that $V(x) = \min_{\pi \in P} \sum_s \pi_s u(x_s)$.
- The set of priors captures the extent of ambiguity aversion.

Ellsberg's Paradox

- More recently, Klibanoff et al. (2005, ECMA) axiomatize:

  EU preferences over lotteries, subjective EU over second-order lotteries. Representation: $V(x) = E_{\mu}[\phi(\sum_s \pi_s u(x_s))]$, where μ is a (subjective) distribution of prob. weights π that the agent considers possible, and ϕ is increasing, concave.

- Bag A: Two states (ball is R or G), μ puts prob. 1 on $\pi_R = 1/2$.
  Bag B: μ might put positive prob on a range of π_R's.
  Prefer the "\$20 gamble" with Bag A if ϕ is concave.

- But: In Fox-Tversky experiment, ambiguity aversion is found only if the subject is asked both questions (in comparison), not when s/he only faces the "\$20 gamble" with either Bag A or Bag B.
  The agent is suspicious that the experimenter is an adversary?

Final Comments on EU Models

- The Expected Utility model of preferences over risky lotteries builds on compelling axioms and admits a simple representation. Framework can have objective or subjective probabilities.
- Models with risk and risk aversion explain a variety of observed phenomena. Many strange effects can be accommodated by taking into account background risk or state-dependent utility.
- There are few paradoxes that raise fundamental questions about framing and how people understand choice problems (ambiguity).
- Apparent violations of rational choice models inspired work in Behavioral Economics (from Rabin to Gül & Pesendorfer) and Decision Theory (Machina, Epstein, recently Klibanoff et al, …).

Rabin's Puzzle

- Suppose we flip a fair coin; you win \$100 if H, lose \$90 if T.
- Do you take this bet given your current wealth? Higher wealth?
- Repeating the experiment with different wealth levels and stakes should calibrate your risk aversion – if your behavior is consistent with some vNM utility function.
- THM (Rabin, 2000 ECMA): An agent that has concave vNM utility and turns down a bet to win \$100 or lose \$90 with equal probabilities at all wealth levels should also turn down a bet to win $\$x$ or lose \$800 with equal probabilities for any $x \in \mathbb{R}$.

*Proof (Recitation)

- $u'' \le 0 \Rightarrow [u(w+100) - u(w+10)]/90 \le [u(w+100) - u(w)]/100$, so

  $$u(w+100) - u(w+10) \le .9\,[u(w+100) - u(w)]. \qquad (*)$$

- Agent turns down the gamble at $w+100$, hence

  $[u(w+200) + u(w+10)]/2 \le u(w+100)$, or

  $u(w+200) - u(w+100) \le u(w+100) - u(w+10)$.

- (*) $\Rightarrow$ $u(w+200) - u(w+10) \le .9\,[u(w+100) - u(w)]$.

  Induction: $u(w+100k) - u(w+100(k-1)) \le .9^{k-1}[u(w+100) - u(w)]$.

- Gain $u(x) - u(w) \le \sum_{k \ge 1} u(w+100k) - u(w+100(k-1))$

  $\le (1 + .9 + .9^2 + \ldots)[u(w+100) - u(w)] \le 10[u(w+100) - u(w)]$.

- Loss: $u(w) - u(w-800) = \sum_{1 \le k \le 8} u(w-100(k-1)) - u(w-100k)$

  $\ge (1 + \ldots + (10/9)^8)\,[u(w+100) - u(w)] > 10[u(w+100) - u(w)]$. ■

Insight from the Puzzle

- As we increase *w* by \$100, if the agent keeps turning down the gamble (+\$100, -\$90 with equal probs) then his utility increment declines by a constant factor (that is, exponentially).
  This is why he does not take the second bet for any large *x*.
- Possible answers:
  (1) Is there anybody who turns down that gamble with <u>any</u> *w*?
  (I became risk neutral to small bets as soon as I got a job.)
  (2) Initial wealth as a "mental state"?
- We pursue answer (1) next.

Measuring Riskiness

Dean Foster and Sergiu Hart, "An Operational Measure of Riskiness", working paper, 2008.

- Measure the riskiness of a gamble without a detailed model of decision making (e.g., no utility fcn, no expected utility).
- Idea: Calculate the "critical level of wealth" at which it is "safe" to accept a gamble (bet).
- Gambles: Positive expected value random variables that have <u>negative realizations</u>.
- When is it "safe" to accept a gamble? The decision maker with a given initial wealth wants to avoid bankruptcy.

Setup

- A gamble g is a random variable with finitely many possible outcomes, $x_1,\ldots,x_m$, each with positive probability p_i , $i=1,\ldots,m$.
- Assume that the expected value of g is positive, $\sum_i p_i x_i > 0$.
- Denote the maximal loss of g by $L(g) = -\min\{x_i\}$, and assume $L(g) > 0$. (There is a negative outcome).
- Denote the set of all such gambles by G.
- Let the decision maker's initial wealth be W_1. In a static setup, the way to avoid bankruptcy is not to take any g with $L(g) > W_1$.
- The paper's approach is more sophisticated: dynamic setup, potentially unknown process $(g_t)_{t>1}$, nature as "adversary".
- Avoid bankruptcy: $\lim_{t\to\infty} W_t = W_\infty > 0$.

An Example

- Suppose g is the 50-50% gamble, win \$100 or lose \$90. (Recall Rabin's puzzle.)
- Suppose that the decision maker faces an iid sequence of such gambles. (Main results will not assume the process is known.)
- Suppose that in period t = 1,2,…, the decision maker can take on any proportion of the gamble, i.e., $\alpha_t g$ in period t, $\alpha_t \in [0,1]$.
- Assume that the agent uses a "simple proportional strategy": Computes a number Q, called "critical wealth" for gamble g, and with wealth W_t in period t, he chooses $\alpha_t = \max\{W_t/Q, 1\}$.

    Interpretation: If W_t is greater than the critical wealth Q, then take the whole gamble, if not, take a portion of it.

An Example, cont'd

- How will W_t evolve, given g and strategy corresponding to Q ?
- $W_T = W_{T\text{-}1} + \alpha_{T\text{-}1}\, g_{T\text{-}1} = W_{T\text{-}1} + (W_{T\text{-}1}/Q)g_{T\text{-}1} = W_{T\text{-}1}\,(1{+}g_{T\text{-}1}/Q)$
  $= W_1 \prod_{t\leq T} (1{+}g_t/Q)$.
- Suppose, for instance, that $Q = 500$ for the gamble 50-50% chance of +100 and -90. Then, $1{+}g_t/Q = 1.2$ and 0.82 with equal probabilities; i.e., the returns in each period are +20% or -18%. In the long run, by the Law of Large Numbers, the average wealth change per period is $\sqrt{1.2*0.82} = 0.992$. Almost surely, $W_t \to 0$.
- If $Q = 1000$, then $1{+}g_t/Q \in \{1.1, 0.91\}$, the long-run average per-period wealth growth is $\sqrt{1.1*0.91} = 1.005$, so $W_t \to \infty$.
- At $Q = 900$, the long-run average growth is exactly unity.

More Generally

- Suppose that g_t is iid, and the agent uses the "simple proportional strategy" described above, with critical wealth Q.
- THM: Let R solve E[ln(1+g/R)] = 0.
  If $Q > R$ then $W_t \to \infty$; while if $Q < R$ then $W_t \to 0$, almost surely.
  - By compounding returns and using the Law of Large Numbers as we did in the example. ■
- Note: For any g, there is a unique solution R to E[ln(1+g/R)] = 0. This needs a little proof.
- If the gamble is 50-50% chance of gain a and loss b, $0 < b < a$, then $R = ab/(a-b)$.
- We may call R an objective measure of riskiness.

Going Full Tilt

- Relax the iid assumption and assume instead that $(g_t)_{t\geq 1}$ is <u>any process</u> where each g_t belongs to a finite cone G.
- Indeed, assume that the process is generated by an adversary – Nature picks the gambles to maximize the chance of bankruptcy given the agent's strategy.
- Let $\alpha_t \in \{0,1\}$ instead of $[0,1]$. (Nature can scale up/down bets.)
- <u>THM</u>: $\forall g$, let $R(g)$ be the unique solution to $E[\ \ln(1+g/R(g))\ ] = 0$. If a strategy rejects g at $W < R(g)$ then it avoids bankruptcy, i.e., $\lim_{t\to\infty} W_t > 0$ for any sequence $(g_t)_{t\geq 1}$ and any initial wealth W_1.
  - Proof is similar to calculations in the example; uses martingale convergence theorem in the general case.■

Properties of Riskiness R

- <u>THM</u>: For all gambles g, $h \in G$,

(1) If g and h have the same distribution, then $R(g) = R(h)$.

(2) Homogeneity: $R(\lambda g) = \lambda R(g)$.

(3) Riskiness exceeds maximal loss: $R(g) > L(g)$.

(4) Sub-additivity: $R(g+h) \leq R(g) + R(h)$.

(5) Convexity: $\forall \lambda \in (0,1)$, $R(\lambda g+(1-\lambda)h) \leq \lambda R(g) + (1-\lambda)R(h)$.

(6) Independent gambles: For independent random vars g and h,
$\min\{R(g),R(h)\} < R(g+h) < R(g) + R(h)$.

■ All follow from the formula. ■

Connections

- Agent with vNM utility u rejects gamble g at initial wealth W iff $E[u(W+g)] < u(W)$.

  For $u(x) = \ln(x)$, reject g iff $E[\ln(1+g/W)] < 0$, i.e., $W<R(g)$.

  To avoid bankruptcy, reject any gamble that log-utility rejects.
  Or: CRRA(ρ) utility guarantees no-bankruptcy if $\rho \geq 1$.

- Rabin's puzzle: "Reject a gamble at all wealth levels W" ?
  Is W total current wealth (including value of all future earnings, human capital etc.), or gambling wealth (amount ready to lose).

- Compare with riskiness measures used in finance, like VaR (value at risk). Those are even more ad hoc, this is at least motivated with a good story.