Decision Theory

Lotteries and Expected Utility

Luce, D. and H. Raiffa [1957]: *Games and Decisions,* John Wiley chapter 2.5

there are r prizes $1,\dots,r$

a lottery L consists of a finite vector $(p_1,\dots,p_r)$ where p_i is the "probability" of winning prize i

properties of "probabilities" $p_i \geq 0, \sum_{i=1}^{r} p_i = 1$

Definition: the lottery L_i has $p_i = 1$

Preferences $\geq$ are defined over the set of lotteries

order the lotteries so that $L_i \geq L_{i+1}$, that is higher numbered prizes are worse

Usual preference assumptions:

1) transitivity

2) continuity: for each L_i there exists a lottery $\tilde{L}_i$ such that $p_j = 0$ for $j = 2,\ldots,r-1$ and $L_i \sim \tilde{L}_i$

(in words: we can find probabilities of the best and worst prize that are indifferent to any lottery)

Definition: u_i is such that $\tilde{L}_i = (u_i, 0, \ldots, 0, (1 - u_i))$

Assumptions relating to probability:

a compound lottery is a lottery in which the prizes are lotteries

we can write a compound lottery $(q^1, L^1, q^2, L^2, \ldots, q^k, L^k)$ where

q^i is the probability of lottery L^i (not to be confused with L_i)

1) reduction of compound lotteries

preferences are extended from simple lotteries to lotteries over lotteries by the usual laws of probability

example: $L^1 = (p_1^1, p_2^1, \ldots, p_r^1)$, $L^2 = (p_1^2, p_2^2, \ldots, p_r^2)$

$$(q_1, L^1, q_2, L^2) \sim (q^1 p_1^1 + q^2 p_1^2, q^1 p_2^1 + q^2 p_2^2, \ldots, q^1 p_r^1 + q^2 p_r^2)$$

2) substitutability (independence of irrelevant alternatives)

for any lottery L the compound lottery that replaces L_i with $\tilde{L}_i$ is indifferent to L

$$(p_1, p_2, \ldots, p_r) \sim (p_1, L_1, p_2, L_2, \ldots, p_i, \tilde{L}_i, \ldots, p_r, L_r)$$

3) monotonicity

$(p, 0, \ldots, 0, (1-p)) \geq (p', 0, \ldots, 0, (1-p'))$ if and only if $p \geq p'$

Expected utility theory:

Start with a lottery $L = (p_1, \ldots, p_r)$

Using transitivity and continuity L is indifferent to the compound lottery $(p_1\tilde{L}_1, \ldots, p_r\tilde{L}_r)$

Notice that the lotteries $\tilde{L}_i$ involve only the highest and lowest prizes

Now apply reduction of compound lotteries: this is equivalent to the lottery

$L \sim (u, 0, \ldots, 0, (1-u))$ where $u = \sum_{i=1}^{r} p_i u_i$

This says that we may compare lotteries by comparing their “expected utility” and by monotonicity, higher utility is better

Allais Paradox

Take Q = 1 billion dollars US

Decision problem 1:

Q for sure

(or)

.1 x 5Q, .89 x 1Q, .01 x 0Q

Decision problem 2:

.1 x 5Q, .9 x 0Q

(or)

.11 x 1Q, .89 x 0Q

Decision problem 1:

1 x 1Q for sure **[most common choice]**

(or)

.1 x 5Q, .89 x 1Q, .01 x 0Q

Decision problem 2:

.1 x 5Q, .9 x 0Q **[most common choice]**

(or)

.11 x 1Q, .89 x 0Q

So $u(1) > .1u(5) + .89u(1) + .01u(0)$ or $u(5) < 1.1u(1) - .1u(0)$

And $.1u(5) + .9u(0) > .11u(1) + .89u(0)$ or $u(5) > 1.1u(1) - .1u(0)$

Notice that the original problem had Q equal to 1 million US. This doesn't work well anymore because most people make the second choice in the first problem and the first choice in the second problem, which is consistent with expected utility

Two views:

1) this is a big problem [Tversky and Kahneman, 1979]

decent theory due to Machina [1982], Segal [1990]

2) this is a curiousity due to the unusual magnitudes of the payoffs Rubsinsten [1988], Leland [1994]

Subjective Uncertainty

Ellsburg Paradox

Ellsberg [1961]

Two urns: each contains red balls and black balls

Urn 1: 100 balls, how many red or black is unknown

Urn 2: 50 red and 50 black

Choice 1: bet on urn 1 red or urn 2 red

Choice 2: bet on urn 1 black or urn 2 black

Urn 1: 100 balls, how many red or black is unknown

Urn 2: 50 red and 50 black

Choice 1: bet on urn 1 red or urn 2 red **[urn 2]**

Choice 2: bet on urn 1 black or urn 2 black **[urn 2]**

1 says that urn 2 red more likely than urn 1 red

2 says that urn 2 black more likely than urn 2 black

but this is inconsistent with probabilities that add up to 1

Can introduce theory of “ambiguity aversion” as in Schmeidler [1989],

Ghirardato and Marinacci [2000]

Basically probabilities do not add up to one; remaining probability is assigned to “nature” moving after you make a choice and choosing the worst possibility for you. [The stock market always tumbles right after I buy stocks.]

Ellsburg Paradox Paradox

we should be able to break the indifference

Urn 1: 1000 balls, how many red or black is unknown

Urn 2: 501 red and 499 black

Choice 1: bet on urn 1 red or urn 2 black **[urn 2]**

Choice 2: bet on urn 1 black or urn 2 black **[urn 2]**

Combine this into a single choice:

Bet on urn 1 red, urn 1 black or urn 2 black

Ambiguity aversion says go with urn 2 black…

But this is a bad idea: flip a coin to decide between urn 1 red and urn 1 black

Risk Aversion

Jensen's inequality

u is a concave function if and only if $u(Ex) \geq Eu(x)$

that is: you prefer the certainty equivalent

so concavity = risk aversion

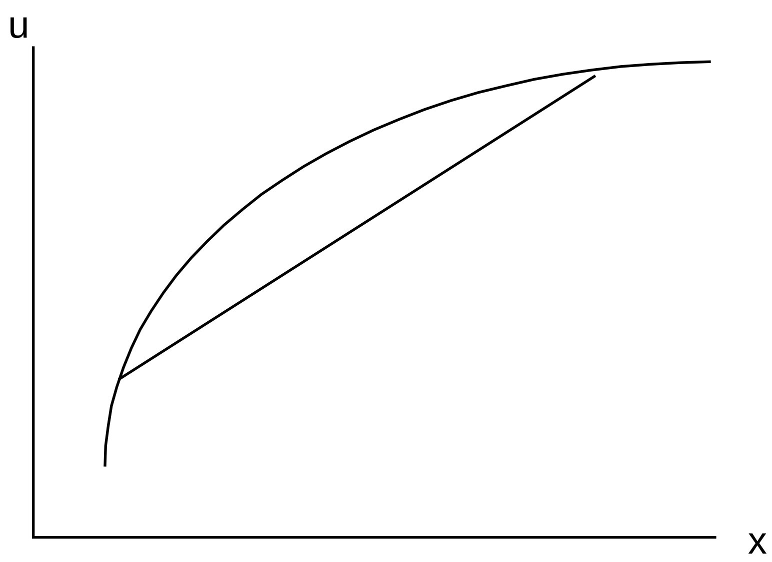

Risk premium

y a random income with $Ey = 0, Ey^2 = 1$

$u(x - p) = Eu(x + \sigma y)$

Taylor series expansion:

$$u(x) - pu'(x) = E\left[u(x) + \sigma u'(x)y + (1/2)\sigma^2 u''(x)y^2\right]$$
$$= u(x) + (1/2)\sigma^2 u''(x)$$

so $p = -\dfrac{u''(x)}{u'(x)}\dfrac{\sigma^2}{2}$

we can also consider the relative risk premium

$u(x - \rho x) = Eu(x + \sigma yx)$

$$\rho = -\frac{u''(x)x}{u'(x)}\frac{\sigma^2}{2}$$

Measures of Risk Aversion

Absolute risk aversion

The coefficient of absolute risk aversion is $-\dfrac{u''(x)}{u'(x)}$

Relative risk aversion

The coefficient of relative risk aversion is $-\dfrac{u''(x)x}{u'(x)}$

Changes in Risk Aversion with Wealth

We ordinarily think of absolute risk aversion as declining with wealth (this is a condition on the third derivative of u).

Constant relative risk aversion

$u(x) = \frac{x^{1-\rho}}{1-\rho}$ also known as "constant elasticity of substitution" or CES

$\rho \geq 0$

$$-\frac{u''(x)x}{u'(x)} = \frac{\rho x^{-\rho-1}x}{x^{-\rho}} = \rho$$

$\rho = 0$ linear, risk neutral

$\rho = 1 \;\; u(x) = \log(x)$

useful for empirical work and growth theory

note that constant relative risk aversion implies declining absolute risk aversion

How risk averse are people?

Equity premium

Mehra and Prescott [1985]; Shiller [1989] data annual 1871-1984

Mean real return on bonds r_b=1.9%; Mean real return on S&P 7.5%

Equity premium $\lambda = .056$

Standard error of real stock return 18.1%, $\sigma = 0.181$.

normalized real per capita consumption standard error $s = .035$

let x denote initial wealth

Let α be fraction of portfolio in S&P

calculate consumption

$$u((1-\alpha)x(1+r_b)+\alpha x(1+\bar{r}_s+\sigma y))=$$

$$u(x+xr_b+\alpha\lambda x+\alpha\sigma yx)$$

$$\frac{d}{d\alpha}Eu(x+xr_b+\alpha\lambda x+\alpha\sigma yx)$$

$$=\lambda xEu'+\sigma xEyu'$$

$$=u'\lambda x+\lambda xEu''()[\alpha\sigma yx]+\sigma xEyu'(x+xr_b+\alpha\lambda x+0)+\sigma xEyu''()[\alpha\sigma yx]$$

$$=u'\lambda x+\alpha u''\sigma^2x^2=0$$

$$\rho=\lambda/(\alpha\sigma^2)\approx 1.81\alpha^{-1}$$

$$s^2=\mathrm{var}[((1-\alpha)x+\alpha(1+\lambda)x+\alpha\sigma yx)/x]=\alpha^2\sigma^2$$

or $\alpha^{-1}\approx\sigma/s=5.17$ giving $\rho=8.84$

Risk Aversion in the Laboratory

In laboratory experiments we often observe what appears to be risk averse behavior over small amount of money (typical payment rates are less than $50/hour, and play rarely lasts two hours)

How can people be risk averse over gambles involving such an insignificant fraction of wealth?

Rabin [2000]: Risk aversion in the small leads to impossible results in the large

“Suppose we knew a risk-averse person turns down 50-50 lose $100/gain $105 bets for any lifetime wealth level less than $350,000, but knew nothing about the degree of her risk aversion for wealth levels above $350,000. Then we know that from an initial wealth level of $340,000 the person will turn down a 50-50 bet of losing $4,000 and gaining $635,670.”

Risk Aversion in the Field

There is surprisingly little systematic evidence about how risk averse people are.

One exception: Hans Binswanger [1978] took his grant money to rural India and conducted a series of experiments involving gambles for a significant fraction of annual income.

His findings: risk aversion is high (ρ on the order of 20), and inconsistent with expected utility theory – initial wealth plays a greater role than the theory allows, along much the same lines discussed by Rabin.

Remark: it is easy to see that deviations from the amount that is “expected to be earned” play some role. But it is a long leap from that to a systematic theory.

Intertemporal Preference

Additive Separability

finite sum

$\sum_{t=1}^{T} u_t$

finite average

$(1/T)\sum_{t=1}^{T} u_t$

infinite time average

$$LIM(1/T)\sum_{t=1}^{T} u_t$$

where *LIM* could be liminf, limsup or Banach limit

with liminf and limsup there are two versions of expected utility:

ELIM vs *LIME* former makes sense, latter is actually used

with Banach limit it makes no difference

Impatience

discounted utility

$$\sum_{t=1}^{T} \delta^{t-1} u_t$$

infinite discounted utility

$$\sum_{t=1}^{\infty} \delta^{t-1} u_t$$

average discounted utility

$$(1-\delta)\sum_{t=1}^{\infty} \delta^{t-1} u_t$$

note that average present value of 1 unit of utility per period is 1

The real equity premium puzzle

Utility $u(x) = \dfrac{x^{1-\rho}}{1-\rho}$, $\sum_{t=1}^{\infty} \delta^{t-1} u_t$

Consumption grows at a constant rate $x_t = \gamma^t$

$u'(x) = x^{-\rho}$

marginal rate of substitution $\dfrac{1}{1+r} = \dfrac{\delta u'(x_{t+1})}{u'(x_t)} = \dfrac{\delta \gamma^{-\rho(t+1)}}{\gamma^{-\rho t}} = \delta \gamma^{-\rho}$

1889-1984 from Shiller [1989]

average real US per capita consumption growth rate 1.8%

$\rho = 8.84 \;\; r = 17\%$

Mean real return on bonds 1.9%; Mean real return on S&P 7.5%

http://www.dklevine.com/econ201/interest.xls

How does the market react to good news?

Value of claims to future consumption relative to current consumption

$$x_1 = 1$$

$$\frac{\sum_{t=2}^{\infty} \delta^{t-1} u'(x_t) x_t}{u'(1)}$$

$$\sum_{t=2}^{\infty} \delta^{t-1} \gamma^{-(t-1)\rho} \gamma^{t-1} = \sum_{t=1}^{\infty} \left[\delta \gamma^{1-\rho}\right]^t = \frac{\delta \gamma^{1-\rho}}{1-\delta \gamma^{1-\rho}}$$

to be finite we need $\delta \gamma^{-\rho} < 1$

$$\frac{\partial}{\partial \gamma} \frac{\delta \gamma^{1-\rho}}{1-\delta \gamma^{1-\rho}} = \frac{\delta(1-\rho)\gamma^{-\rho}}{(1-[\delta \gamma^{-\rho}])^2}$$

$\rho > 1$ this is negative

Hyperbolic Discounting

(based on Villaverde and Mukherji [2001])

Q1: would you like \$10 today or \$15 tomorrow?

Q2: would you like \$10 100 days from now or \$14 101 days from now?

Some people answer prefer \$10 in Q1 and \$14 in Q2. This is inconsistent with (geometric) discounting and a time and risk invariant marginal rate of substitution between days.

Note that (because of asset markets) this makes little sense when expressed in terms of money. So let us suppose that the “paradox” refers to consumption.

One explanation: “hyperbolic discounting” meaning preferences of the form $u(c_1) + \theta \sum_{t=2}^{\infty} \delta^{t-1} u(c_t)$

A more straightforward explanation:

Uncertainty about preferences 100 days from now.

Suppose marginal utility of consumption can take on two values 1 or 2 with equal probability and that the daily subjective discount factor is to a good approximation 1.

Today the value of todays and tomorrow's marginal utility is know with certainty. Hence the subjective interest rate can take on the values of 1, 0 or -½ with probabilities .25, .5 and .25. Expected subjective interest rate is .125 = 1/8. If you are offered 10 today versus 15 tomorrow, you take 10 today with probability .25.

Suppose on the other hand, suppose that preferences 100 days from now are unknown. Ratio of expected utilities is 1, so subjective interest rate is 0. If you are offered 10 in 100 days versus 14 in 101 days you always take 14.

Notice that pigeons have apparently figured this out correctly.

Experiments that have examined demand for commitment and consumption favor the geometric theory.

Interpersonal Preferences

Experimental results

Roth et al [1991]

US $10.00 stake games, round 10

Second and final round of bargaining game:

Player may take x or reject it and get nothing.

The other player gets $10-x

5 of 27 offers with x>0 are rejected

5 of 14 offers with 5>x>0 are rejected

x	*Offers*	*Rejection Probability*
\$2.00	1	100%
\$3.25	2	50%
\$4.00	7	14%
\$4.25	1	0%
\$4.50	2	100%
\$4.75	1	0%
\$5.00	13	0%
total	27	

Dynamic Programming

$\alpha \in A$ action space: finite

$y \in Y$ state space: finite

$\pi(y'|y,\alpha)$ transition probability

period utility $u(\alpha, y)$ with discount factor $0 \le \delta < 1$

$\overline{u} = \max u(\alpha, y) - \min(u(\alpha, y))$

Strategies

finite histories $h = (y_1, y_2, \ldots, y_t)$ with $t(h) = t$ $y(h) = y_t$; $h-1$; $y_1(h)$; $h' \geq h$

H space of all finite histories; this is countable

strategies $\sigma : H \to A$

Σ space of all strategies

all maps from a countable set to a finite set

the product topology $\sigma^n \to \sigma$ means that $\sigma^n(h) \to \sigma(h)$ for every h

Theorem: every sequence in the product topology has a convergent subsequence, so the space of strategies is compact

(proven in any elementary topology textbook) define a strong Markov strategy $\sigma(h) = \sigma(h')$ if $y(h) = y(h')$

a strong Markov strategy is equivalent to a map

$$\sigma : Y \to A$$

recursively define

$$\pi(h|y_1,\sigma) \equiv$$

$$\begin{cases} \pi(y(h)|y(h-1),\sigma(h-1))\pi(h-1|y_1,\sigma) & t(h) > 1 \\ 1 & t(h) = 1 \text{ and } y_1(h) = y_1 \\ 0 & t(h) = 1 \text{ and } y_1(h) \neq y_1 \end{cases}$$

calculate the average present value of the objective function

$$V(y_1,\sigma) \equiv (1-\delta)\sum_{h\in H} \delta^{t(h)-1} u(\sigma(h), y(h))\pi(h|y_1,\sigma)$$

Dynamic Programming Problem

(*) maximize $V(y_1, \sigma)$ subject to $\sigma \in \Sigma$

a value function is a map $v: Y \to \Re$ bounded by $\bar{u}$

note that in this setting, it is simply a finite vector v_y

Lemma: a solution to (*) exists

Definition: ***the*** value function

$$v(y_1) \equiv \max_{\sigma \in \Sigma} V(y_1, \sigma)$$

Proof: the maximum exists because in the product topology on Σ $V(y_1, \sigma)$ is continuous in σ and Σ is compact why is V continuous?

suppose $\sigma^n \to \sigma$

$$
\begin{aligned}
V(y_1,\sigma^n) &= (1-\delta)\sum\nolimits_{h\in H}\delta^{t(h)-1}u(\sigma^n(h),y(h))\pi(h|y_1,\sigma^n) \\
&= (1-\delta)\sum\nolimits_{t(h)<T}\delta^{t(h)-1}u(\sigma^n(h),y(h))\pi(h|y_1,\sigma^n) \\
&\quad +(1-\delta)\sum\nolimits_{t(h)\geq T}\delta^{t(h)-1}u(\sigma^n(h),y(h))\pi(h|y_1,\sigma^n) \\
&\to (1-\delta)\sum\nolimits_{t(h)<T}\delta^{t(h)-1}u(\sigma(h),y(h))\pi(h|y_1,\sigma) + O(\delta^T\overline{u})
\end{aligned}
$$

so as $T\to\infty$ we have $V(y_1,\sigma^n)\to V(y_1,\sigma)$

Bellman equation

we define a map $T : \Re^Y \to \Re^Y$ by $w' = T(w)$ if

$$w'(y_1) = \max_{\alpha \in A}(1-\delta)u(\alpha, y_1) + \delta \sum_{y'_1 \in S} \pi(y'_1 | y_1, \alpha) w(y'_1)$$

Lemma: the value function is a fixed point of the Bellman equation $T(v) = v$

in other words the most you can get next period is also given by the value function

$$v(y_1) = \max_{\alpha \in A}(1-\delta)u(\alpha, y_1) + \delta \sum_{y'_1 \in S} \pi(y'_1 | y_1, \alpha) v(y'_1)$$

Lemma: the Bellman equation is a contraction mapping

$$\|T(w)-T(w')\| \le \delta\|w-w'\|$$

Proof:

key observation $\|\max_\alpha f(\alpha) - \max_\alpha g(\alpha)\| \le \max_\alpha \|f(\alpha) - g(\alpha)\|$

$$\left\|\max_{\alpha\in A}(1-\delta)u(\alpha,y_1)+\delta\sum_{y'_1\in S}\pi(y'_1|y_1,\alpha)w(y'_1)-\max_{\alpha\in A}(1-\delta)u(\alpha,y_1)+\delta\sum_{y'_1\in S}\pi(y'_1|y_1,\alpha)w'(y'_1)\right\|$$

$$\le \max_{\alpha\in A}\left\|(1-\delta)u(\alpha,y_1)+\delta\sum_{y'_1\in S}\pi(y'_1|y_1,\alpha)w(y'_1)-(1-\delta)u(\alpha,y_1)+\delta\sum_{y'_1\in S}\pi(y'_1|y_1,\alpha)w'(y'_1)\right\|$$

$$= \max_{\alpha\in A}\left\|\delta\sum_{y'_1\in S}\pi(y'_1|y_1,\alpha)w(y'_1)-\delta\sum_{y'_1\in S}\pi(y'_1|y_1,\alpha)w'(y'_1)\right\|$$

$$\le \delta\max_{\alpha\in A}\sum_{y'_1\in S}\pi(y'_1|y_1,\alpha)\mid w(y'_1)-w'(y'_1)\mid$$

$$\le \delta\max_{\alpha\in A}\sum_{y'_1\in S}\pi(y'_1|y_1,\alpha)\|w-w'\|$$

$$= \delta\|w-w'\|$$

Corollary: the Bellman equation has a unique solution

Proof: Let w be another solution

$$\|v - w\| = \|T(v) - T(w)\| \leq \delta \|v - w\| \Rightarrow \|v - w\| = 0$$

Conclusion: the unique solution to the Bellman equation is the value function

since the value function is *a* solution to the Bellman equation, and the solution is unique

Lemma: there is a strong Markov optimum and it may be found from the Bellman equation

Proof:

define a strong Markov plan by

$$\sigma(y_1) \in \arg\max_{\alpha \in A}(1-\delta)u(\alpha, y_1) + \delta \sum_{y'_1 \in S(y_1)} \pi(y'_1 | y_1, \alpha) v(y'_1)$$

work the value function forward recursively to find

$$\begin{aligned} v(y_1(h)) = {} & (1-\delta)\sum\nolimits_{t(h)<T} \delta^{t(h)-1} \pi(y(h) | y_1(h), \sigma) u(\sigma(y(h)), y(h)) \\ & +(1-\delta)\sum\nolimits_{t(h)=T} \delta^{T} \pi(y(h) | y_1(h), \sigma) v(h) \end{aligned}$$

and observe that v is bounded by $\bar{u}$ so that the final terms disappears asymptotically

Application – job search

three states: unemployed (u), have a bad job (b), have a good job (g)

the only choice: whether or not to quit a bad job and become unemployed

pr(g|g) = 1 (good job is absorbing)

pr(g|u) = a > b = pr(g|b, not quit) (chance of getting a good job)

pr(b|u) = c (chance of getting a bad job when unemployed)

u(g) = d

u(b) = 1

u(u) = 0

procedure: find the value function

$$v(g) = d$$

$$v(u) = (1-\delta)0 + \delta(av(g) + cv(b) + (1-a-c)v(u))$$

$$v(b) = \max \begin{cases} (1-\delta) + \delta(bv(g) + (1-b)v(b)) \\ (1-\delta) + \delta v(u) \end{cases}$$

step 0: substitute out v(g)

$$v(u) = (1-\delta)0 + \delta(ad + cv(b) + (1-a-c)v(u))$$

$$v(b) = \max \begin{cases} (1-\delta) + \delta(bd + (1-b)v(b)) \\ (1-\delta) + \delta v(u) \end{cases}$$

case 1: optimum is to quit a bad job

$$v(b) = (1-\delta) + \delta v(u)$$

substitute:

$$v(u) = \delta(ad + c((1-\delta) + \delta v(u)) + (1-a-c)v(u))$$

$$(1 - \delta(1-a-c) - \delta^2 c)v(u) = \delta ad + \delta(1-\delta)c$$

$$v(u) = \frac{\delta ad + \delta(1-\delta)c}{(1 - \delta(1-a-c) - \delta^2 c)}$$

verify the Bellman equation:

$$
\begin{aligned}
(1-\delta)+\delta v(u) &\geq (1-\delta)+\delta(bd+(1-b)v(b)) \\
&= (1-\delta)+\delta(bd+(1-b)((1-\delta)+\delta v(u)))
\end{aligned}
$$

$$v(u) \geq \frac{bd+(1-b)(1-\delta)}{1-\delta(1-b)}$$

$$\frac{\delta ad+\delta(1-\delta)c}{(1-\delta(1-a-c)-\delta^2 c)} \geq \frac{bd+(1-b)(1-\delta)}{1-\delta(1-b)}$$

for example when $b=0, c=0, a=1$, $\delta d \geq 1$