Very Rough Notes: Uncertainty

September 2014 John Kennan

Fair Bets. A fair bet is a gamble that breaks even on average (the expected value of the net payoff is zero). An even money bet on a coin toss is the canonical example. Casinos are not in the business of offering fair bets (since it costs money to operate a casino, a casino that only offered fair bets would have zero net revenue and positive costs, so it would go out of business).

Lotteries. A lottery $L = (p, x)$ (or gamble) is a vector of payoffs $(x_1, x_2, \ldots, x_n)$ with probabilities $(p_1, p_2, \ldots, p_n)$. The payoffs might be consumption bundles or monetary payoffs.

St Petersburg Paradox. The payoff is 2^n, where n is the first time the coin comes up heads.

The expected value is infinite. If the game lasts just two rounds, it is worth $1 + 1$. If it lasts 3 rounds, it is worth $1 + 1 + 1$. And so on.

See Shapley, Journal of Economic Theory, Volume 14, Issue 2, April 1977, Pages 439–442.

Also StPetersburg.ods (LibreOffice spreadsheet [www.libreoffice.org]).

Preferences over Lotteries. We start by assuming that consumers have rational preferences over lotteries, and we assume that the preference ordering $\succsim$ is continuous (in exactly the same way as before). Then we assume something new:

Definition. The preference ordering $\succsim$ on the space of lotteries satisfies the *independence axiom* if

$$L^A \succsim L^B \iff \alpha L^A + (1-\alpha) L \succsim \alpha L^B + (1-\alpha) L$$

for all lotteries L, and for all $\alpha \in (0, 1)$.

The idea here is that the individual ultimately cares only about consequences, and the probabilities of these consequences. The two mixtures of the two simple lotteries yield the lottery L with probability $(1-\alpha)$, and in this contingency the individual is indifferent between the two mixtures. Thus it makes sense to focus on the complement of this set of contingencies, and there L^A dominates L^B.

Definition. The preference ordering $\succsim$ on the space of lotteries has an expected utility representation if there is a function u such that

$$L^A \succsim L^B \iff \sum_{i=1}^{n} p_i^A u_i \geq \sum_{i=1}^{n} p_i^B u_i$$

where $u_i = u(x_i)$

This utility representation is called a von Neumann-Morgenstern utility function.

The existence of a utility representation of $\succsim$ implies that each consequence has an associated utility number, namely the utility of the lottery that gives this consequence with probability 1.

Proposition. *[linearity]*

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

Proposition. *[affine transformations] If U is an expected utility representation of $\succsim$ on $\mathcal{L}$, then $\tilde{U}$ is another EU representation if and only if*

$$\tilde{U}(L) = a + bU(L)$$

where a is a real number, and b is a positive number

The point here is that the expected utility representation is not invariant under monotonic transformations.

Proposition. *[Expected Utility Theorem] If a rational preference ordering $\succsim$ on the lottery space $\mathcal{L}$ satisfies continuity and independence, then it has an expected utility representation.*

Theorem. *[Jensen's Inequality] If the function f is strictly concave, and if X is a random variable on a finite set $\{x_1, x_2, \ldots, x_n\}$ of distinct points with positive probabilities $(p_1, p_2, \ldots, p_n)$ then*

$$Ef(X) < f(EX)$$

Proof. The proof is by induction. First, if $n = 2$ then

$$\begin{aligned} Ef(X) &= p_1 f(x_1) + (1 - p_1) f(x_2) \\ &< f(p_1 x_1 + p_2 x_2) \\ &= f(EX) \end{aligned}$$

where the inequality holds by the definition of strict concavity. Next, if $n = 3$ then

$$\begin{aligned} Ef(X) &= p_1 f(x_1) + p_2 f(x_2) + (1 - p_1 - p_2) f(x_3) \\ &= p_1 f(x_1) + (1 - p_1) \frac{p_2 f(x_2) + (1 - p_1 - p_2) f(x_3)}{1 - p_1} \\ &= p_1 f(x_1) + (1 - p_1)(q_1 f(x_2) + (1 - q_1) f(x_3)) \end{aligned}$$

where

$$q_1 = \frac{p_2}{1 - p_1}$$

and strict concavity implies that

$$q_1 f(x_2) + (1 - q_1) f(x_3) < f(\hat{x})$$

where

$$\hat{x} = q_1 x_2 + (1 - q_1) x_3$$

So

$$Ef(X) < p_1 f(x_1) + (1 - p_1) f(\hat{x})$$

and strict concavity implies that

$$p_1 f(x_1) + (1 - p_1) f(\hat{x}) < f(\bar{x})$$

where

$$\begin{aligned} \bar{x} &= p_1 x_1 + (1 - p_1) \hat{x} \\ &= p_1 x_1 + (1 - p_1)(q_1 x_2 + (1 - q_1) x_3) \\ &= p_1 x_1 + p_2 x_2 + p_3 x_3 \\ &= EX \end{aligned}$$

Thus

$$Ef(X) \quad < \quad f(EX)$$

So given that the result is true for $n = 2$, it must also be true for $n = 3$. In exactly the same way, it must also be true for $n = 4$, and so on. □

Insurance. An expected utility maximizer with wealth w buys α units of insurance at price q against a loss D that occurs with probability π, where $q \geq \pi$. Find α.
Expected utility is given by

$$\pi u(w - D + \alpha - \alpha q) + (1 - \pi) u(w - \alpha q)$$

The first-order condition for α is

$$(1 - q) \pi u'(w - D + \alpha - \alpha q) = q(1 - \pi) u'(w - \alpha q)$$

Write this as

$$\frac{u'(w - D + \alpha - \alpha q)}{u'(w - \alpha q)} = \frac{q}{1 - q} \frac{1 - \pi}{\pi}$$

Buying one unit of insurance at price q means getting \$1 with probability π in exchange for \$$q$. So the breakeven price is $q = \pi$. At this price, the right side of the above equation is 1, and this matches the left side if and only if $\alpha = D$. This means that a risk-averse individual always buys full insurance if the price is "fair" – if the insurance company breaks even on average.

If the insurance company charges more than the fair price, then the marginal utility when the event occurs is higher than when it does not (so the individual is hoping that the event does not occur, even though this individual is insured).

Expected marginal utility in the bad state is $\pi u'(w - D + \alpha - \alpha q)$. The price of a contingent claim that pays \$1 in the bad state is q

Expected marginal utility in the good state is $(1 - \pi) u'(w - \alpha q)$. The price of a contingent claim that pays \$1 in the good state is $1 - q$. This means reducing insurance by one unit.

$$\begin{aligned} \Delta C_g &= -q \Delta \alpha \\ \Delta C_b &= (1 - q) \Delta \alpha \end{aligned}$$

Equating expected marginal utility per dollar here means

$$\pi MU_b \Delta C_b = (1 - \pi) MU_g \Delta C_g$$

So

$$\frac{\pi MU_b}{q} = \frac{(1 - \pi) MU_g}{1 - q}$$

Example. Certainty Equivalents

Suppose the utility function is $u(x) = \sqrt{x}$. A lottery that pays 4 with probability $\frac{1}{4}$ and 9 with probability $\frac{3}{4}$ has expected utility $EU = \frac{1}{4} \times 2 + \frac{3}{4} \times 3 = \frac{11}{4}$. The certainty equivalent is a (degenerate) lottery that pays $\left(\frac{11}{4}\right)^2 = \frac{121}{16}$ for sure. The expected value of the lottery is $\frac{31}{4} = \frac{124}{16}$. Someone with this square root utility function could sell this lottery to someone with a linear utility function at a price of $\frac{122}{16}$, and both people would be better off.

Risk Aversion. Risk aversion is how the insurance industry makes a living. Risk-loving behavior is how Las Vegas makes a living.

Definition. Given a utility function u, and wealth level w, the Pratt-Arrow coefficient of absolute risk aversion is $r(w) = -\frac{u''(w)}{u'(w)}$

If this coefficient is constant with respect to x, then $-r = \frac{d\log(u'(x))}{x}$ so $\log(u'(x)) = c_1 - rx$ and $u'(x) = e^{c_1}e^{-rx}$ and $u(x) = -ke^{-rx}$ for some nonnegative constant k, and this constant can be taken as 1.

The idea is to measure the curvature of the utility function in a way that is invariant under a change of scale.

Decreasing absolute risk aversion means that r is decreasing in x.

Definition. Given a utility function u, and wealth level w, the Pratt-Arrow coefficient of *relative* risk aversion is $\varrho(w) = -x\frac{u''(w)}{u'(w)}$

If this coefficient is constant with respect to x, then $-\frac{\varrho}{x} = \frac{d\log(u'(x))}{x}$ so $\log(u'(x)) = c_1 - \varrho\log(x)$ and $u'(x) = e^{c_1}x^{-\varrho}$ and $u(x) = \frac{x^{1-\varrho}-1}{1-\varrho}$.

The interpretation is that wealthy people are approximately risk-neutral, as long as the risks are not big.

Portfolio Choice. Suppose there are two investments. One is safe and the other is risky. An investor with wealth w puts a in the risky asset and $w - a$ in the safe asset. The mean return on the risky asset is greater than 1, and the safe return is 1. The expected utility is

$$EU = \int_{-\infty}^{\infty} u(az + w - a)\, dF(z)$$

So the first-order condition is

$$\int_{-\infty}^{\infty} (z-1)\, u'(az + w - a)\, dF(z) = 0$$

At $a = 1$ this condition can't hold, because the expected value of Z exceeds 1 (and $u'(w)$ is a constant). So it is always optimal to accept some risk in exchange for a higher return.

Diversification. Suppose wealth is divided over two risky assets, represented by random variables X_1 and X_2, with means μ_1, μ_2 and variances σ_1^2, σ_2^2, and covariance σ_{12}.

$$W = \alpha X_1 + (1-\alpha) X_2$$

The expected return is

$$EW = \alpha\mu_1 + (1-\alpha)\mu_2$$

The variance is

$$\begin{aligned}
\sigma^2 &= E(W - EW)^2 \\
&= E(\alpha(X_1 - \mu_1) + (1-\alpha)(X_2 - \mu_2))^2 \\
&= \alpha^2\sigma_1^2 + (1-\alpha)^2\sigma_2^2 + 2\alpha(1-\alpha)\sigma_{12}
\end{aligned}$$

To minimize variance (e.g. if the two assets have the same mean), the first-order condition is

$$0 = 2\alpha\sigma_1^2 - 2(1-\alpha)\sigma_2^2 + (2-4\alpha)\sigma_{12}$$

so

$$\sigma_2^2 - \sigma_{12} = \alpha\left(\sigma_1^2 + \sigma_2^2 - 2\sigma_{12}\right)$$

and

$$\alpha = \frac{\sigma_2^2 - \sigma_{12}}{\sigma_1^2 - \sigma_{12} + \sigma_2^2 - \sigma_{12}}$$

Then the asset proportions are given by

$$\frac{1-\alpha}{\alpha} = \frac{\sigma_1^2 - \sigma_{12}}{\sigma_2^2 - \sigma_{12}}$$

So if X_1 has higher variance, the minimum variance allocation favors X_2. But putting everything in the asset with lower variance is not optimal.

Option Value.

$$E\max\{u(X), u(Y)\} \geq \max\{Eu(X), Eu(Y)\}$$

To show this, just note that

$$\max\{u(X), u(Y)\} \geq u(X)$$

so

$$E\max\{u(X), u(Y)\} \geq Eu(X)$$

and in the same way

$$E\max\{u(X), u(Y)\} \geq Eu(Y)$$

CAPM. There is a risk-free asset, and a risky index fund. It is assumed that the investor cares only about the mean and the variance of returns. A more general version of this just assumes that the investor is an expected utility maximizer: this is the so-called Consumption CAPM.

CCAPM (Consumption Capital Asset Pricing Model). Consider a two-period problem: consume now or later. There is an asset that can be purchased at price p, in exchange for a random payment X per unit, realized next period. If a consumer with wealth w buys a units of this asset, consumption now is $C_1 = w - pa$, and consumption next period will be $C_2 = aX$. The consumer maximizes the expected sum of utilities over the two periods, with a preference for current consumption

$$\max_a \left(u(C_1) + \rho Eu(C_2)\right)$$

where ρ represents the rate of time preference. Thus the choice problem is

$$\max_a \left(u(w - pa) + \rho Eu(aX)\right)$$

The first-order condition for this problem is

$$-pu'(C_1) + \rho Eu'(C_2)X = 0$$

Rearrange this as

$$p = \rho E \frac{u'(C_2)}{u'(C_1)} X$$

This is the basic theory of asset pricing (see Cochrane).

The insurance policy mentioned above is an asset, and the first-order condition agrees with the asset pricing formula.

In the case of CRRA utility,

$$u(C) = \frac{C^{1-\gamma} - 1}{1-\gamma}$$

the marginal utility is

$$u'(C) = C^{-\gamma}$$

so the stochastic discount factor is

$$m = \rho \left(\frac{C_1}{C_2} \right)^{\gamma}$$

where the stochastic part is C_2. In the case of log utility, this reduces to

$$m = \rho \frac{C_1}{C_2}$$

If it were known that consumption next period will be the same as consumption this period, then the price of the asset is just the expected value of the payoff, discounted by the time preference factor. Otherwise if X is high when C_2 is low, that makes the asset more valuable.