Utility

Topics:

1. Decisions Under Uncertainty
   — Certain Equivalents
2. Expected Utility
3. Constant Absolute Risk Aversion
4. Eliciting Utility Functions
5. Choosing Among Lotteries
6. Appendix: Approximating a Certain Equivalent
7. Appendix: Finance.

(See Dixit & Skeath: 2nd ed. pp. 228–230, 300–303;
3rd ed. pp. 258–261, 358–361.)

1. Decisions with Uncertainty

Choose among the four lotteries with unknown probabilities on the branches: uncertainty —

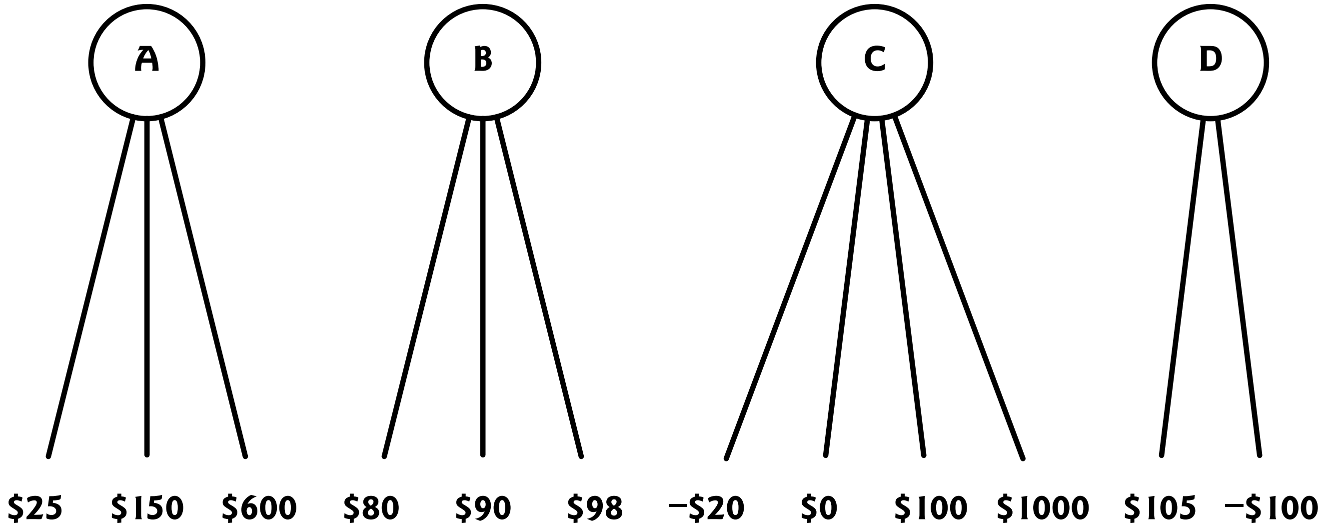

(Write down your answer.)

Five possible answers:

There are several possibilities:

1. ***The extreme pessimist*: choose the lottery with the highest minimum payoff.**

   known as "maxmin" decision making, from *max*imising the *min*imum payoff.

   ∴ would result in choice of lottery B.

2. ***The extreme optimist*: choose the lottery with the highest maximum payoff.**

   known as "maxmax" decision making, from *max*imising the *max*imum payoff.

   ∴ would result in choice of lottery C.

2a. ***Hurwicz:*** **choose the lottery with the highest value of a weighted average of the minimum and maximum, say $\alpha \times$ lottery X's minimum payoff + $(1 - \alpha) \times$ lottery X's maximum payoff.**

- **when $\alpha = 1$, the Extreme Pessimist Rule $\therefore$ B**
- **when $\alpha = 0$, the Extreme Optimist Rule $\therefore$ C**

If $0.899 \leq \alpha \leq 0.901$, then choose lottery A.

If more than two possible payoffs, this rule ignores the intermediate payoffs, which shouldn't happen.

3. **Choose the lottery with the highest *average* payoff.**

has the advantage that it includes *all* payoffs, not just the extreme ones, but it imputes equal probabilities to each payoff's occurrence. (The *Laplace* criterion.)

Moreover, it also assumes a risk-neutral decision maker.

$\therefore$ would result in choice of lottery C.

4. For lotteries with more common structure — say, a matrix R_{ij}, where lottery i and state of the world j result in payoff R_{ij} — we can use the *Savage rule of minimum regret* for the wrong decision:
choose the lottery which minimises the maximum regret, where regret is the difference between the contingent outcome's payoff in the lottery you chose and the highest contingent outcome's payoff.

(See Apocalypse Maybe in the Readings.)

If we have *the probabilities* of the lotteries' outcomes (say, from a smoothly working roulette wheel), a new rule is possible:

5. Choose the lottery with the *highest expected payoff*: weight each outcome with the probability of its occurrence.

Includes all payoffs and the probabilities of their occurrence, but still assumes *a risk-neutral decision maker.*

Risk aversion

Risk aversion: the clear evidence that many people will forgo expected profit to ensure certainty by selling a gamble at a price less than the expected profit.

Risk aversion is not indicated by the slope of the utility curve: it's the *curvature*: if the utility curve is locally —

1. **linear (say, at a point of inflection), then the decision maker is *locally risk neutral.***
2. **concave (its slope is decreasing — Diminishing Marginal Utility), then the decision maker is *locally risk averse;***
3. **convex (its slope is increasing), then the decision maker is *locally risk preferring.***

See

http://www.gametheory.net/Mike/applets/Risk/risk.html

Risk-averse, risk-neutral, and risk-preferring utility functions.

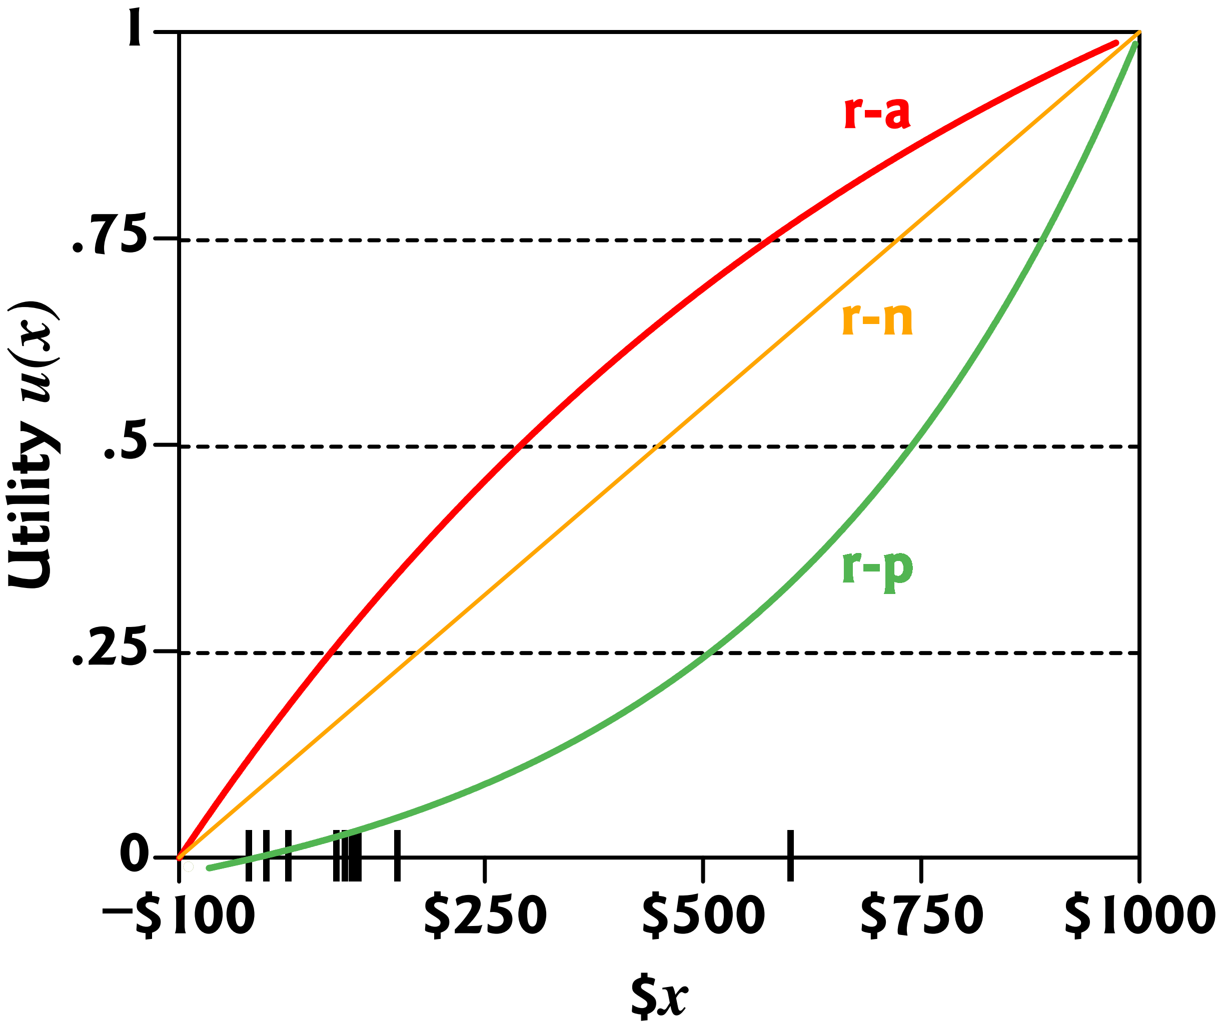

The Certain Equivalent (C.E.) of a lottery.

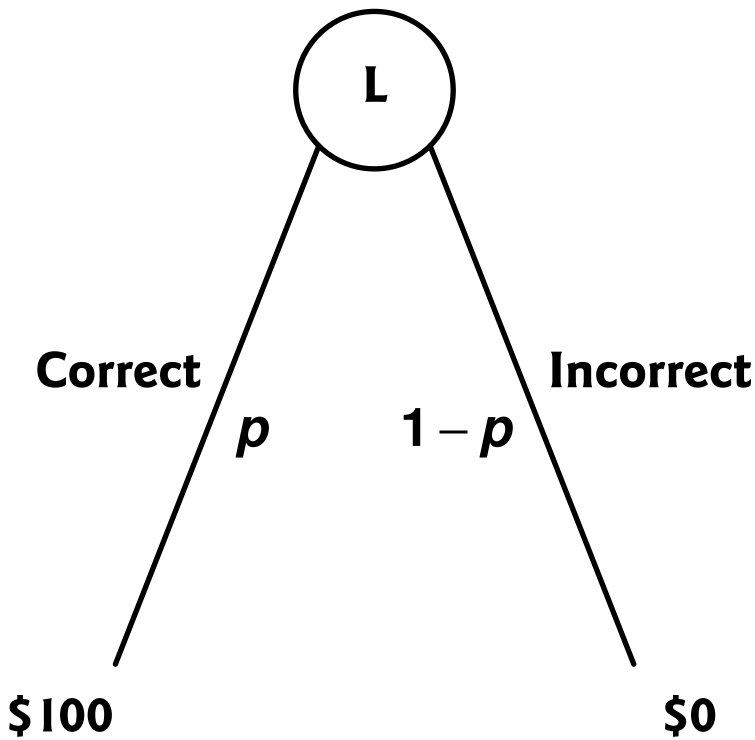

$\sim$ \$_____?

2. Expected Utilities

An ordering which avoids inconsistencies over preferences and likelihoods:

- **Assign utilities to all payoffs,**
  and probabilities to all states of the world,
- **then rank lotteries by their expected utilities.**
- ***The utility of a lottery is its expected utility.***
  (by definition)

$$U(\tilde{x}) = \sum_{i=1}^{n} p_i\, U(x_i), \text{ where } \tilde{x} \text{ is the C.E.}$$

According to Savage, this is the only ordering which satisfies five general conditions (or axioms) we'd like a good decision rule to satisfy:

Completeness and Transitivity,
Continuity,
Substitutability,
Monotonicity, and
Decomposability.

Wealth Independence

Q: What if I gave each of you a \$100 bill together with a ticket to play your preferred lottery of the four at the top of this section. Would your preferences among the lotteries now change?

If not, then your preferences exhibit *wealth independence*.

The Delta property (or Wealth Independence property): an increase of all prizes in a lottery by an amount Δ increases the Certain Equivalent (C.E.) by Δ.

- **Suppose your C.E. for an equiprobable lottery on \$0 and \$100 is \$25.**
- **The lottery owner agrees to pay you an additional \$100 regardless of outcome: your final payoffs will be \$100 and \$200 with equal probability.**
- **If you feel that your C.E. would now be \$125 and reason consistently in all such situations, then you satisfy the Delta property.**

Accepting Wealth Independence

Acceptance of the Wealth Independence (or Delta) property has strong consequences:

➢ **The utility curve is restricted to be either linear or an exponential:**

$u(x)$ must have one of the forms:

$$u(x) = a + b\,x\,,$$
$$\text{or}\quad u(x) = a + b\,e^{-\gamma x}\ ,$$

where a, b, and γ are constants.

➢ **The buying and selling prices of a lottery will be the same for any individual.**

Satisfying the Delta property means that the C.E. of any proposed lottery is independent of the wealth already owned.

This wealth is just a "Δ" that does not affect the preference:

The linear and exponential utility curves are called *wealth-independent,* or constant-absolute-risk-aversion (CARA) functions.

3. Constant Absolute Risk Aversion

Parameterise the exponential utility function as:

$$\text{(A)} \quad u(x) = \frac{1 - e^{-\gamma x}}{1 - e^{-\gamma}},$$

where $u(0) = 0$ and $u(1) = 1$, or as:

$$\text{(B)} \quad U(x) = 1 - e^{-\gamma x},$$

where $U(0) = 0$ and $U(\infty) = 1$.

γ is the *risk aversion coefficient* $\equiv -\frac{u''(x)}{u'(x)}$

Sign of γ	Risk profile	Curvature
$\gamma = 0$	risk neutral	$u''(x) = 0$
$\gamma > 0$	risk averse	$u''(x) < 0$
$\gamma < 0$	risk preferring	$u''(x) > 0$

Acceptance of the Delta property leads to the characterisation of risk preference by a single number, the *risk aversion coefficient.*

The reciprocal of the risk aversion coefficient is known as the *risk tolerance*, $R = \frac{1}{\gamma}$.

Assessing your Risk Tolerance R when your Utility is Wealth-Independent.

The exponential utility function is given by:

$$U(x) = a + b\,e^{-\frac{x}{R}},$$

where R is a parameter that determines how risk-averse the utility function is, the *risk tolerance*, and a and b are constants used to normalise the function.

$R = \frac{1}{\gamma}$, where γ is the risk-aversion coefficient.

Larger values of R make the exponential utility function less curved and so closer to risk neutral, while smaller values of R model greater risk aversion.

As we have seen, the exponential utility function is appropriate if (and only if) the individual's preferences satisfy the Delta Property of wealth independence.

A simple choice to obtain one's CARA Risk Tolerance

$$\begin{cases} \text{Win } \$Y \text{ with probability } ½ \\ \text{Lose } \$\,\frac{1}{2}\,Y \text{ with probability } ½ \end{cases}$$

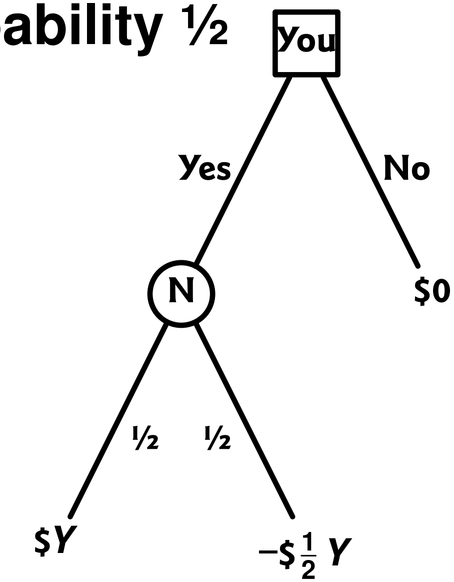

has a C.E. of $0

Q: What is the *maximum* size of Y at which you'd prefer doing nothing to having this lottery: the point at which you'd give the lottery ticket away (i.e. at which it has a C.E. of zero)?

This Y is approximately equal to your risk tolerance R in the exponential (wealth-independent) utility function.

(See Clemen, *Making Hard Decisions*, pp. 379–382.)

Ron Howard's insights ...

Howard (1988) gives reasonable values of determining a company's risk tolerance R in terms of sales, net income, or equity.

According to Howard, a company's risk tolerance R is approximately:

- **6.4% of annual sales,**
- **1.24 $\times$ annual net income, or**
- **15.7% of equity.**

Exponential utility functions exhibit CARA.

Logarithmic utility functions exhibit falling risk aversion (strictly: constant relative risk aversion) — more realistic?

See: Howard R A (1988), Decision analysis: practice and promise, *Management Science*, 34, 679-695.

Two Types of Decision-Makers: Fred and Mary.

Fred's preferences across lotteries don't change when he comes into an inheritance: his utility function exhibits Wealth Independence (or satisfies the Delta Property).

Mary, however, finds that her preferences *do* change with her wealth, even for small amounts. She doesn't exhibit Wealth Independence in her preferences over lotteries.

Fred's utility function is easily determined by eliciting his risk tolerance R through the C.E. of the Y lottery on the previous page, and plugging it into an exponential utility function. Let's say for Fred $Y = R = 900 = 1/\gamma$.

Mary's utility curve can be derived by asking her the C.E. of a series of lotteries, as described below. Each of her answers determines the next lottery she confronts. In general, the lotteries (and so the utility curve elicitted) will be specific to a particular decision of Mary's.

4. Eliciting Utility Functions

Choose among the four lotteries depicted below:

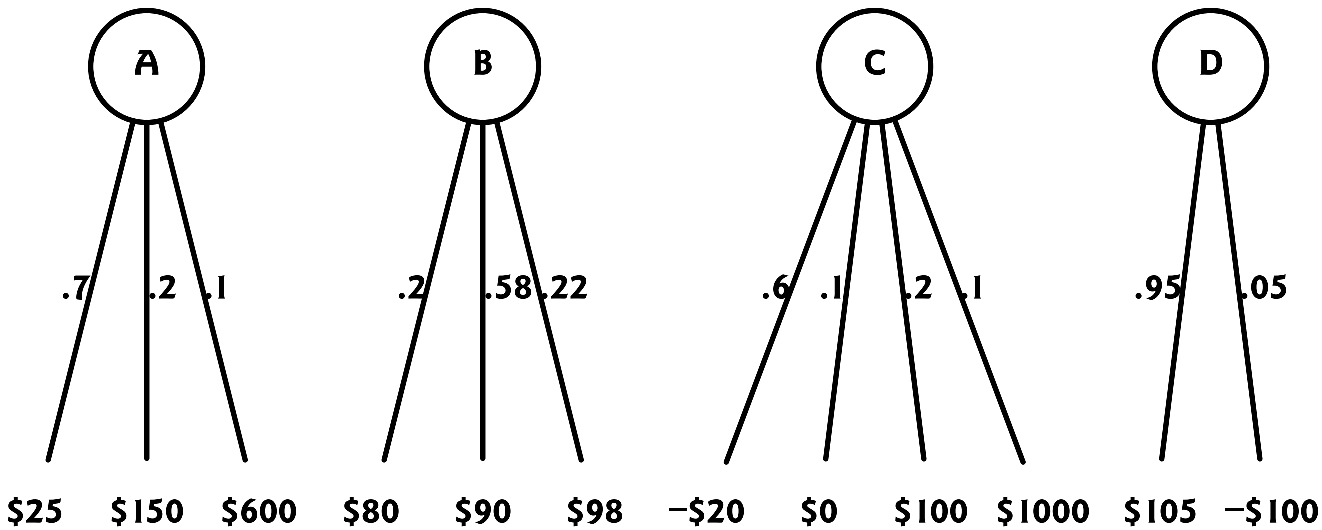

The probabilities are objectively determined:

> **the lotteries are all based on things like the spin of a smooth roulette wheel, etc.**

Choosing among lotteries ...

It is difficult to choose: difficult to think —

- **about probabilities such as 0.22, and**
- **about gambles with four possible prizes.**

But *if* Mary, say, subscribes to the three axioms of utility theory (Transitivity, Substitution, Continuity),

then we know that Mary's choice *should be* based on:

maximising the expectation of a utility function.

So we want Mary's choice behaviour among the four gambles to conform to her expected utility maximisation.

We need to discover Mary's utility function: *how*?

We can assess Mary's utility function:
by making some judgements that are easier than those called for in a direct choice among the four gambles above.

Question 1: **What is Mary's C.E. for the lottery L1:**

$$\begin{cases}\text{probability } \tfrac{1}{2} \text{ of getting \$1000} \\ \text{probability } \tfrac{1}{2} \text{ of owing \$100}\end{cases}$$

This gamble is selected so that its two prizes span all the prizes in the four gambles from which Mary must choose (set $u(\$1000) = 1$ and $u(-\$100) = 0$), and it gives probability ½ to each:

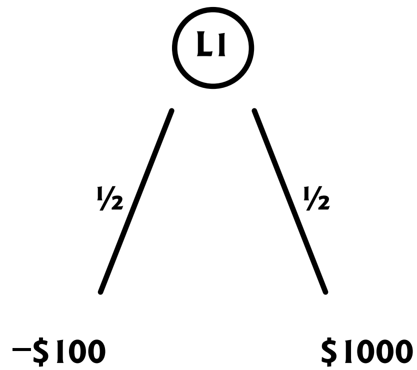

Not a trivial judgement to make, but not so hard, because:

> **comparing a sure thing with a gamble having only two prizes and a simple 50–50 probability structure.**

Mary's Answer 1: **she's indifferent between the gamble above and $400 for sure. So Mary's C.E. for lottery L1 is $400.**

If $u(-\$100) = 0$, and $u(\$1000) = 1$, then $u(\text{L1}) = 0.5 = u(\$400)$, since the utility of a lottery equals its expected utility, by definition.

Question 2: **What is Mary's C.E. for the gamble L2:**

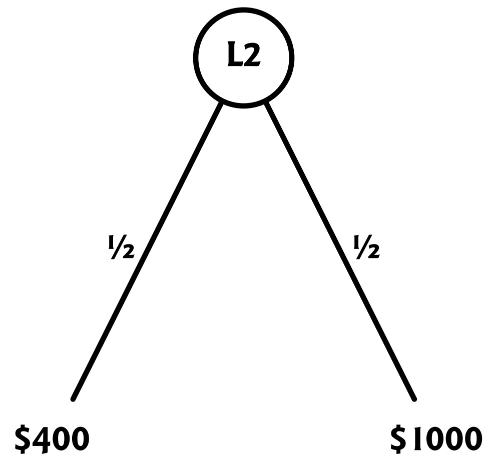

This gamble has:

- **top prize equal to the upper prize from the question, and**
- **bottom prize equal to Mary's previously assessed C.E.**

Mary's Answer 2: **Approximately $675.**

So Mary's Certain Equivalent for lottery L2 is $675.

Question 3: **What is Mary's C.E. for the gamble L3:**

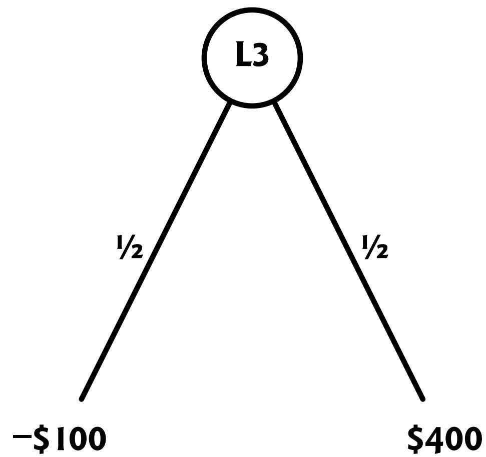

***Mary's Answer 3*: Approximately $100.**

So Mary's Certain Equivalent for lottery L3 is $100.

Why these questions?

If Mary can answer these questions, then we'll have five points on Mary's utility function in the range −\$100 to \$1000:

arbitrarily assigning −\$100 utility = 0 and \$1000 utility = 1,

then the answers reveal that

- **Mary's utility of \$400 [u(\$400)] = 0.5,**
- **Mary's utility of \$675 [u(\$675)] = 0.75, and**
- **Mary's utility of \$100 [u(\$100)] = 0.25.**

Plotting Mary's Curve:

With these five values, we can rough in a pretty good approximation of Mary's utility function and compute her expected utilities for the four original gambles, making her choice accordingly.

Even if our approximation is off, it is close to Mary's "true utility" and her choice according to the approximation will be nearly as good as the best gamble using Mary's "true utility".

Mary's making some judgement calls above, and she may not be doing so well.

Checking for consistency ...

The data above allow us to run consistency checks, such as: "What is Mary's C.E. for L4:"

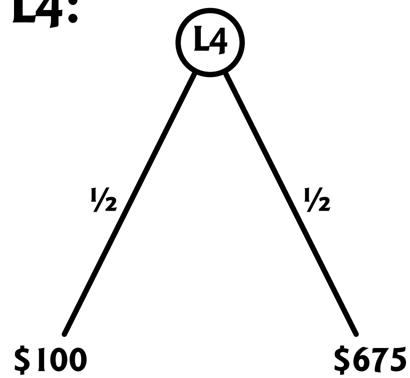

It should be $400. Why?

Because this is a gamble whose prizes have utilities of 0.25 and 0.75 for Mary, so that it has expected utility of 0.5,

and the certain amount of money that has utility 0.5 is $400, for her.

Mary's assessed C.E. for this gamble is approximately $375, but now we can return to Mary's original assessments and iterate so that we have five consistent values.

Why this procedure?

Why is this procedure better than just choosing one of the original gambles?

Because the numerical judgements that we're asking Mary to make are for the easiest conceivable cases that aren't trivial — two-prize lotteries with equally likely outcomes.

Mary is quite ready to believe that she's better at processing that sort of gamble than she is at the four more complicated gambles with which we started.

Is this benefit coming for free?

***No* — we also had to make a qualitative judgement that in this choice situation, the three axioms are good guides for choice behaviour.**

But because we know where the pitfalls in those axioms are, we are confident that, in this case, the axioms are a sound guide to behaviour.

The Utility Curve

How does eliciting the C.Es. of simple lotteries allow us to construct Mary's utility curve?

- **Using the rule that the utility of a lottery is its expected utility,**
- **and setting $u(-\$100) = 0$ and $u(\$1000) = 1$, so that the utility function spans the possible payoffs,**
- **we see that Mary's utility of the first of the simple lotteries above (the lottery over −\$100 and \$1000, C.E \$400) is**
$$\tfrac{1}{2} \times 0 + \tfrac{1}{2} \times 1 = 0.5;$$
- **her utility of the second (over \$400 and \$1000, C.E. \$675) is**
$$\tfrac{1}{2} \times 0.5 + \tfrac{1}{2} \times 1 = 0.75;$$
- **her utility of the third (over −\$100 and \$400, C.E. \$100) is**
$$\tfrac{1}{2} \times 0 + \tfrac{1}{2} \times 0.5 = 0.25;$$
- **and her utility of the fourth (over \$100 and \$675, C.E. \$375) is**
$$\tfrac{1}{2} \times 0.25 + \tfrac{1}{2} \times 0.75 = 0.5.$$

The last three C.E.s (\$675, \$100, and \$375) have been plotted against the lotteries' utilities (0.75, 0.25, and 0.5, resp.) on the following graph, and we've joined the five points with straight lines, to get an approximation for Mary's utility function. *Iterate.*

Utility Curves u = u(x): Mary's, Fred's normalised, and risk-neutral

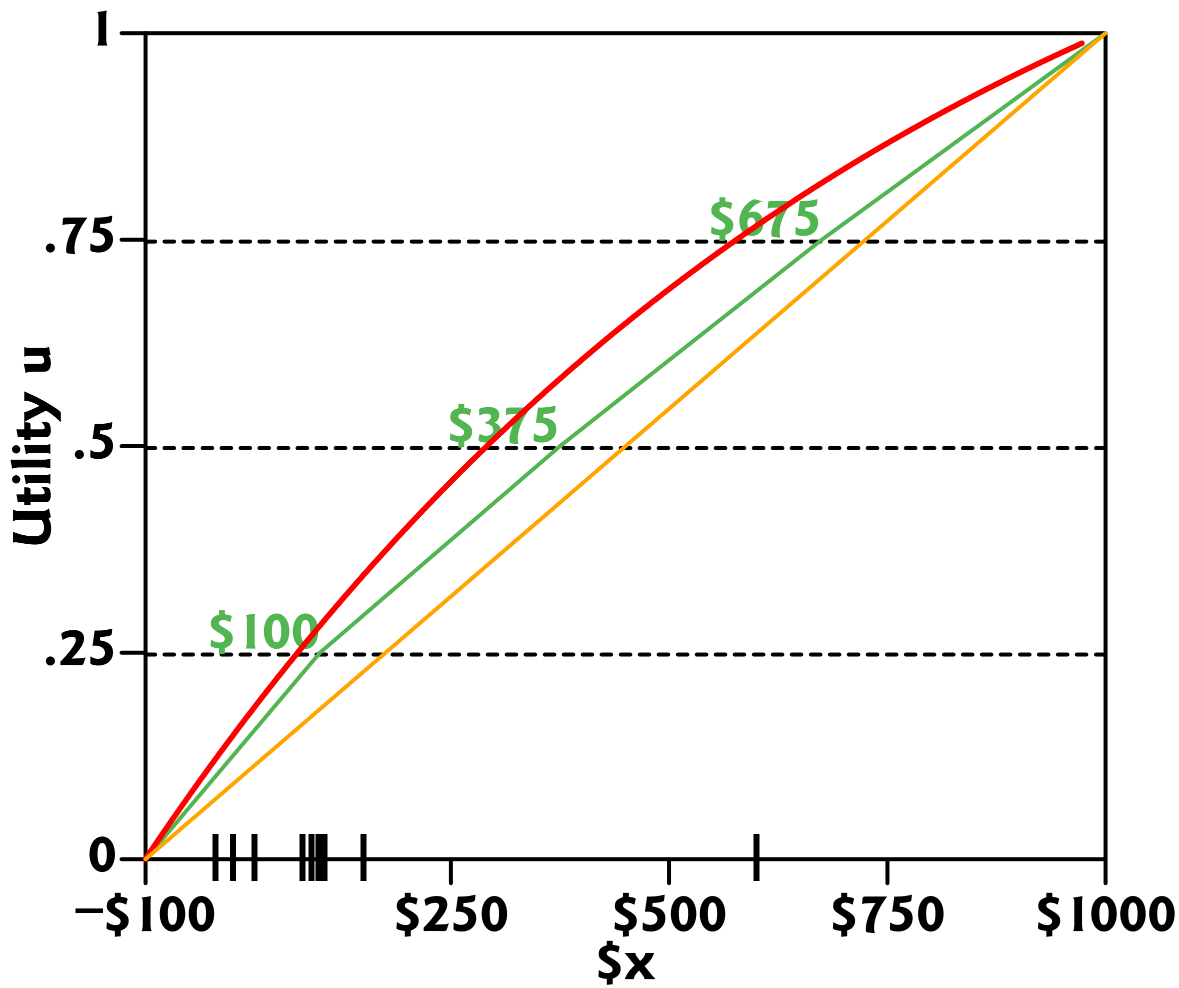

5. The Choice: Using the answers ...

We can use Mary's utility function as plotted, and Fred's utility function $u(x) = 1 - e^{-\frac{x}{900}}$, to calculate their utilities of the four lotteries, A, B, C, and D, of section 4.

Simply a matter of reading off Mary's utilities of the dollar payoffs of the lotteries, and calculating her expected utilities of the four lotteries.

Dollars	Mary's Utility	Fred's Utility
−\$100	0.000	−0.118
−\$20	.100	−0.022
\$0	.125	0.0
\$25	.156	.027
\$80	.225	.085
\$90	.238	.095
\$98	.248	.103
\$100	.250	.105
\$105	.255	.11
\$150	.295	.154
\$600	.69	.487
\$1000	1.0	.671

Mary's expected utilities ...

$\therefore$ **Mary's utility of lottery A is:**

$$0.7 \times 0.156 + 0.2 \times 0.295 + 0.1 \times 0.69 = 0.237$$

➢ **of lottery B:**

$$0.2 \times 0.225 + 0.58 \times 0.238 + 0.22 \times 0.248 = 0.238$$

➢ **of lottery C:**

$$0.6 \times 0.100 + 0.1 \times 0.125 + 0.2 \times 0.250 + 0.1 \times 1 = 0.223$$

➢ **and of lottery D:**

$$0.95 \times 0.255 + 0.05 \times 0 = 0.248$$

So Mary would choose lottery D.

We can see from the plot of her utility function that she's slightly risk averse.

Fred's expected utilities ...

Using Fred's utility function: $u(x) = 1 - e^{-\frac{x}{900}}$

$\therefore$ **Fred's utility of lottery A is:**

$$0.7\times0.027 + 0.2\times0.154 + 0.1\times0.487 = 0.098$$

- **of lottery B:**

$$0.2\times0.085 + 0.58\times0.095 + 0.22\times0.103 = 0.095$$

- **of lottery C:**

$$0.6\times(-0.022) + 0.1\times0 + 0.2\times0.105 + 0.1\times0.671 = 0.086$$

- **and of lottery D:**

$$0.95\times0.11 + 0.05\times(-0.118) = 0.074$$

So Fred would choose lottery A.

(Remember: we're only interested in the relative utilities, not the absolute values, and we can't compare Mary's with Fred's utilities directly.)

The risk-neutral decision maker ...

The risk-neutral decision maker would choose the lottery with the highest expected dollar payoff.

➢ **The expected dollar payoff of lottery A is:**

$$0.7\times\$25 + 0.2\times\$150 + 0.1\times\$600 = \$107.50$$

➢ **of lottery B:**

$$0.2\times\$80 + 0.58\times\$90 + 0.22\times\$98 = \$89.76$$

➢ **of lottery C:**

$$0.6\times-\$20 + 0.1\times\$0 + 0.2\times\$100 + 0.1\times\$1000 = \$108.00$$

➢ **and of lottery D:**

$$0.95\times\$105 + 0.05\times-\$100 = \$94.75$$

So the risk-neutral player would choose lottery C (or perhaps lottery A).

6. App: Approximating a Certain Equivalent

Fred (whose Risk Tolerance R = 900 from the R.T. lottery on p.15) is considering the lottery L:

$$\begin{cases} \text{Win \$2,000 with probability 0.4} \\ \text{Win \$1,000 with probability 0.4} \\ \text{Win \$500 with probability 0.2} \end{cases}$$

Its mean m = \$1,300, and standard deviation σ = \$600.

$$\text{Variance} = \sum_{i=1}^{n} [x_i - m]^2 \,\text{Prob.}(X = x_i) = \$600^2 = \$360{,}000,$$

$$\text{where the mean } m = \sum_{i=1}^{n} x_i \,\text{Prob.}(X = x_i).$$

(Standard deviation σ = square root of the variance.)

Approximating a C.E. (cont.)

Since Fred's $R = 900$, then use the utility function:

$$U(x) = 1 - e^{-\frac{x}{900}}.$$

Check normality OK: $U(\infty) = 1$ and $U(\$0) = 0$. ✓

- ∴ **$U(\$2{,}000) = 0.8916$, $U(\$1{,}000) = 0.6708$, and $U(\$500) = 0.4262$.**
- ∴ $U(L) = 0.8916 \times 0.4 + 0.6708 \times 0.4 + 0.4262 \times 0.2 = 0.7102$
- ∴ **Fred's expected utility of this lottery is 0.7102, and**
- ∴ **his C.E. is \$1,114.71, since $U(\$1{,}114.71) = 0.7102$:**
- ☞ **remember: the utility of a lottery is its expected utility, by definition.**

The C.E. can be approximated:

$$\text{C.E.} \approx \text{mean} - \frac{1}{2} \times \frac{\text{Variance}}{\text{Risk Tolerance}}$$

$$\text{C.E.} \approx \$1{,}300 - \frac{1}{2} \times \frac{360{,}000}{900} \approx \$1{,}100.$$

(Exact with a normal distribution.)

7. Appendix: Application to Finance

Consider a lottery on x described by the probability density function $f_x(.)$.

➢ **Its *C.E.* , $\tilde{x}$, must satisfy the equation:**

$$u(\tilde{x}) = \int f_x(x_0)\, u(x_0)\, dx_0.$$

Why? By the definition of utility, the utility of a lottery [$u(\tilde{x})$] equals its expected utility.

➢ **Substituting the exponential form $u(x) = 1 - e^{-\gamma x}$, we can derive:**

$$\tilde{x} = -\frac{1}{\gamma} \ln \overline{e^{-\gamma x}} = -\frac{1}{\gamma} \ln f_x^e(\gamma),$$

where $f_x^e(.)$ represents the exponential transform of the density function $f_x(.)$, and where $\overline{e^{-\gamma x}}$ is the mean of the function $e^{-\gamma x}$ for the lottery.

➢ **The C.E. of any lottery is therefore the negative reciprocal of the risk aversion coefficient times the natural logarithm of the exponential transform of the variable evaluated at the risk aversion coefficient.**
(So there!)

Finance (cont.)

➢ As γ approaches zero, this expression approaches $\bar{x}$: the C.E. of any lottery to a risk-indifferent individual is the expected value, $\bar{x}$.

$$\text{as } \gamma \to 0,\ \tilde{x} \to \bar{x}.$$

A constant-risk-averse decision maker (with a risk-aversion coefficient γ) is facing a *normal* (or Gaussian) lottery. (For a normal distribution, the exponential transform of the density function, $f_x^e(\gamma)$, is given by $e^{-\gamma m + \frac{1}{2}\gamma^2\sigma^2}$.)

Then his C.E. to this lottery =

the mean minus a half γ times the variance, or

$$\tilde{x} = m - \frac{1}{2}\gamma\sigma^2$$

Hence a risk-averse individual will prefer the lottery with the lower variance σ^2, when both have the same expected value, or mean m. (See Finance.)