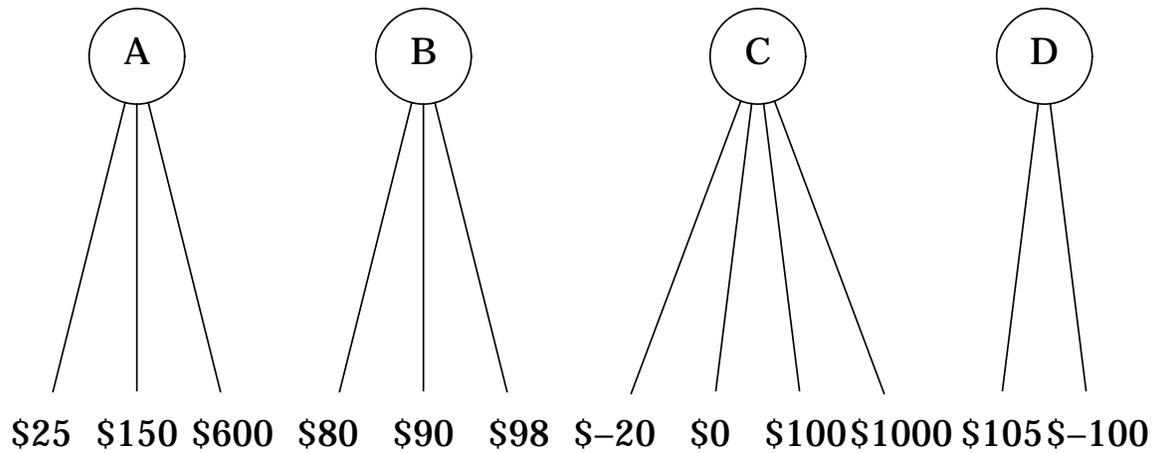


2.11 Utility

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



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 *maximising* the *minimum* 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 *maximising* the *maximum* payoff.

In this case such a rule 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;

— when $\alpha = 0$, the Extreme Optimist Rule.

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.

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:
(see Apocalypse Maybe, in the Package)
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.

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*.

Go with the flow?

Frustrating to search for good decision rules:
the five above can result in quite different choices.

What if a decision maker just states her preferences as she feels them, and acts accordingly, without trying to justify her actions by referring to universally accepted principles?

She may simply consider it more important to satisfy herself than to satisfy axioms.

Inconsistencies

She may be perfectly happy to follow her own preferences until someone points out that her preference ordering contains some obvious contradictions. There are two common sorts:

- inconsistencies with respect to preferences over the possible outcomes, and
- inconsistencies with respect to beliefs about the probabilities of the possible states of the world.

Expected utilities

There exist orderings which do not contain any of these two types of inconsistencies. One:

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

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

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

Wealth Independence?

Acceptance of a sixth property, the Wealth Independence Property (or Delta Property), restricts possible utility functions to be linear (i.e. risk neutral) or exponential:

$$U(X) = a + be^{-\gamma X},$$

where $\gamma = -\frac{U''(X)}{U'(X)}$ is the *risk aversion coefficient*.

- When $\gamma = 0$, the decision maker is risk neutral;
- when γ is positive, risk averse;
- when γ is negative, risk preferring.

2.11.1 Application to Finance

Suppose we consider a lottery on x described by the probability density function $f_x(\cdot)$.

- Its *Certainty Equivalent*, \tilde{x} , must satisfy the equation:

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

- Substituting the exponential form for $u(\cdot)$, we can derive:

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

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

- The Certain Equivalent 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!)

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

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

The Certain Equivalent of a *normal* (or Gaussian) lottery to a constant risk averter with a risk aversion coefficient¹ $\gamma =$

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.)

1. For a normal distribution, the exponential transform of the density function, $f_x^e(\gamma)$, is given by $e^{-\gamma m + \frac{1}{2}\gamma^2\sigma^2}$.

2.12 Risk Aversion

The Delta property: an increase of all prizes in a lottery by an amount Δ increases the Certain Equivalent (the minimum you'd sell the lottery ticket for) by Δ .

- Suppose you say that your Certain Equivalent 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 Certain Equivalent would now be \$125 and reason consistently in all such situations, then you satisfy the Delta property.

Accepting Wealth Independence

Acceptance of the Delta property has strong consequences:

- The utility curve is restricted to be either linear or an exponential: $u(x)$ must *either* have the form:

$$u(x) = a + b x$$

or the form:

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

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 Certain Equivalent 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*.

Exponential utility functions

Parameterise the exponential utility function as:

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

where $u(0) = 0$ and $u(1) = 1$.

γ is the “risk aversion coefficient”.

Sign of γ	Risk profile
$\gamma = 0$	risk neutral
$\gamma > 0$	risk averse
$\gamma < 0$	risk preferring

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*, $R = 1/\gamma$.

2.12.1 Assessing your Risk Tolerance.

The exponential utility function is given by:

$$U(x) = 1 - e^{-\frac{x}{R}},$$

where R is a parameter that determines how risk-averse the utility function is, the *risk tolerance*.

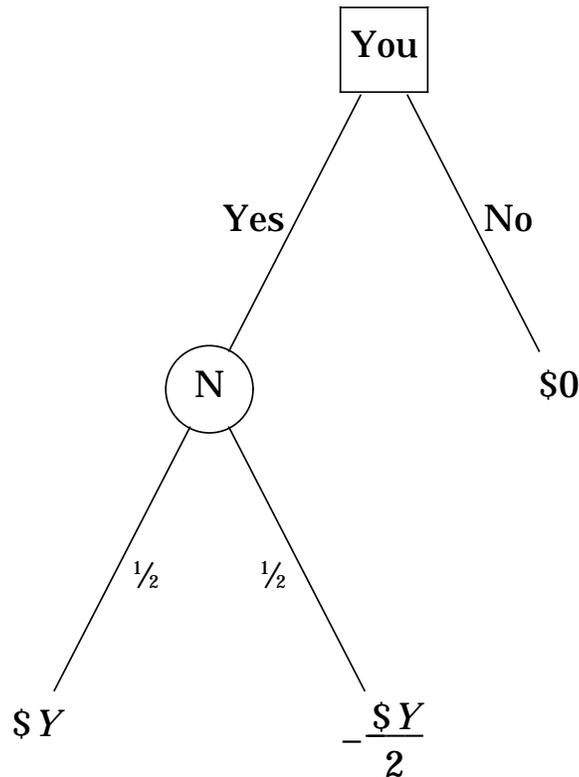
$R = \frac{1}{\gamma}$, where γ is the risk-aversion coefficient.

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 the individual's preferences satisfy the Delta Property of wealth independence.

Consider the gamble:

Win $\$Y$ with probability $\frac{1}{2}$
 Lose $\frac{\$Y}{2}$ with probability $\frac{1}{2}$



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?

This Y is approximately equal to your risk tolerance R in the exponential utility function.

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

2.12.2 Approximation of Certainty Equivalence

Consider the gamble:

Win \$2000 with probability 0.4
 Win \$1000 with probability 0.4
 Win \$500 with probability 0.2

Its mean = \$1300, standard deviation = \$600.

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

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

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

Expected utility of this gamble (with $R = 900$) is 0.7102, and the c.e. is \$1114.71.

This can be approximated with the equation:

$$\text{c.e.} \approx \text{exp. value} - \frac{1}{2} \frac{\text{Variance}}{\text{Risk Tolerance}}$$

which gives \$1100.

Howard (1988) gives reasonable values of determining a company's risk tolerance in terms of sales, net income, or equity. R is about 6.4% of sales, $1.24 \times$ net income, or 15.7% of equity.

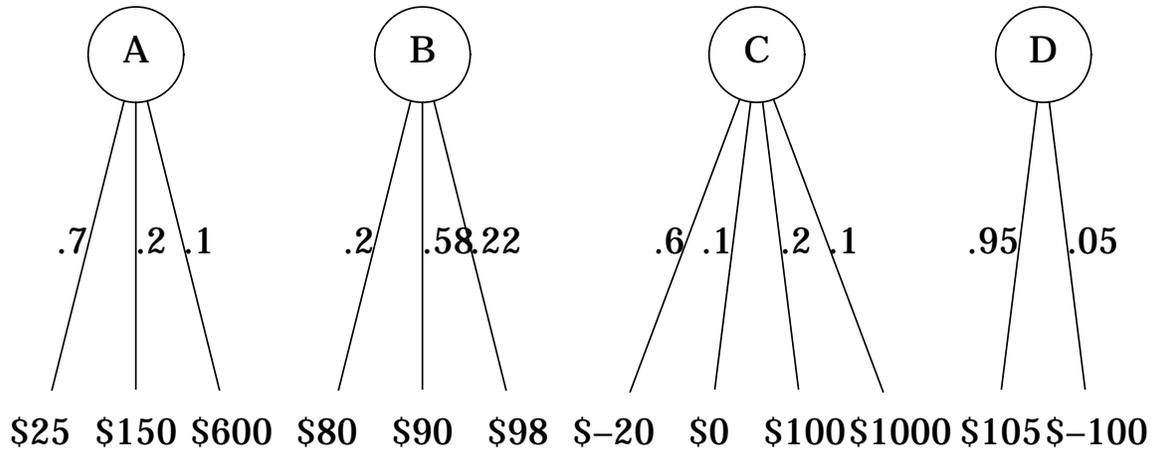
Exponential utility functions exhibit constant risk aversion;

logarithmic utility functions exhibit falling risk aversion — more realistic?

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

2.13 Eliciting Utility Functions

Choose among the four gambles depicted below:



The probabilities are objectively determined:

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

Difficult to choose: difficult to think —

- about probabilities such as .22, and
- about gambles with four possible prizes.

But *if* I subscribe to the three axioms of utility theory (Transitivity, Substitution, Continuity), then I know that my choice *should be* based on: *maximising the expectation of a utility function.*

So I want my choice behaviour among the four gambles to conform to my expected utility maximisation.

I need to discover my own utility function: *how?*

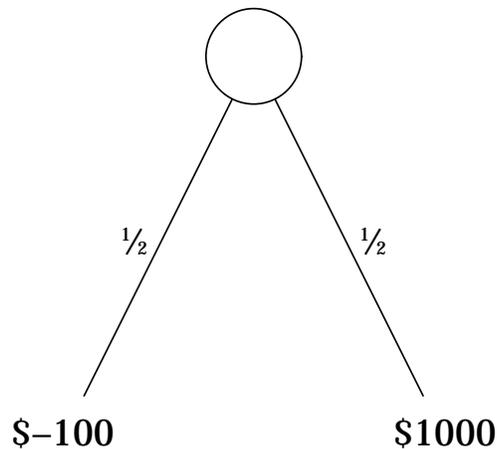
I can assess my utility function:

by making some judgements that are easier than those called for in a direct choice among the four gambles above.

Question 1: What is my Certainty Equivalent for:

- probability $\frac{1}{2}$ of getting \$1000 and
- probability $\frac{1}{2}$ of getting \$-100?

This gamble is selected so that its two prizes span all the prizes in the four gambles from which I must choose, and it gives probability $\frac{1}{2}$ to each:

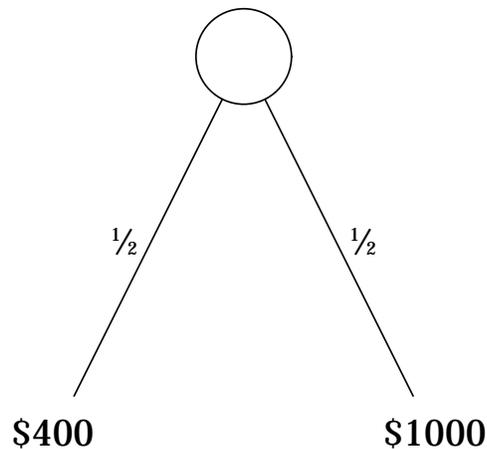


Not a trivial judgement to make, but not so hard, because:

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

Answer 1: indifferent between the gamble above and \$400 for sure.

Question 2: What is my Certainty Equivalent for the gamble:

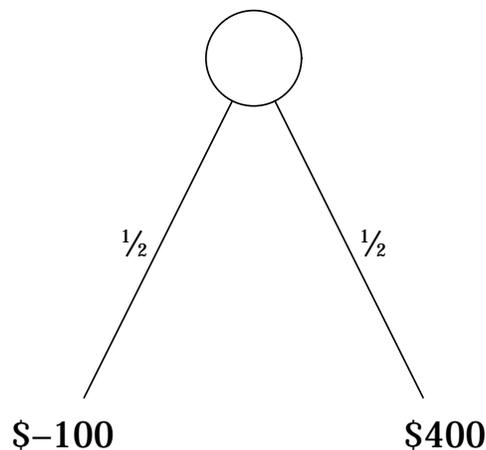


This gamble has:

- top prize equal to the upper prize from the question, and
- bottom prize equal to my previously assessed Certain Equivalent.

Answer 2: Approximately \$675.

Question 3: What is my Certain Equivalent for the gamble:



Answer 3: Approximately \$100.

Why these questions?

If I can answer these questions, then I'll have five points on my utility function in the range \$-100 to \$1000:

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

then the answers reveal that

- my utility = 0.5 at \$400,
- my utility = 0.75 at \$675, and
- my utility = 0.25 at \$100.

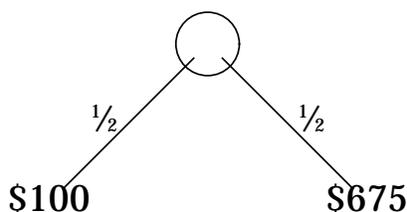
With these five values, I can rough in a pretty good approximation of my utility function and compute expected utilities for the four original gambles, making my choice accordingly.

Even if my approximation is off, it is close to my “true utility” then the choice according to the approximation will be nearly as good as the best gamble using my “true utility”.

I'm making some judgement calls above, and I may not be doing so well.

The data above allow me to run consistency checks, such as:

“What is my Certainty Equivalent for:”



It should be \$400.

Why?

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

and the certain amount of money that has utility 0.5 is \$400.

My assessed Certainty Equivalent for this gamble is approximately \$375, but now I can return to my original assessments and iterate so that I have five consistent values.

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

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

I'm quite ready to believe that I'm better at processing that sort of gamble than I am at the four more complicated gambles with which we started.

Is this benefit coming for free?

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

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

2.14 The Utility Curve

How does eliciting the Certain Equivalents of simple lotteries allow us to construct my 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 the utility of the first of the simple lotteries above (the lottery over $\$-100$ and $\$1000$, c.e. $\$400$) is

$$\frac{1}{2} \times 0 + \frac{1}{2} \times 1 = 0.5;$$

- the utility of the second (over $\$400$ and $\$1000$, c.e. $\$675$) is

$$\frac{1}{2} \times 0.5 + \frac{1}{2} \times 1 = 0.75;$$

- the utility of the third (over $\$-100$ and $\$400$, c.e. $\$100$) is

$$\frac{1}{2} \times 0 + \frac{1}{2} \times 0.5 = 0.25;$$

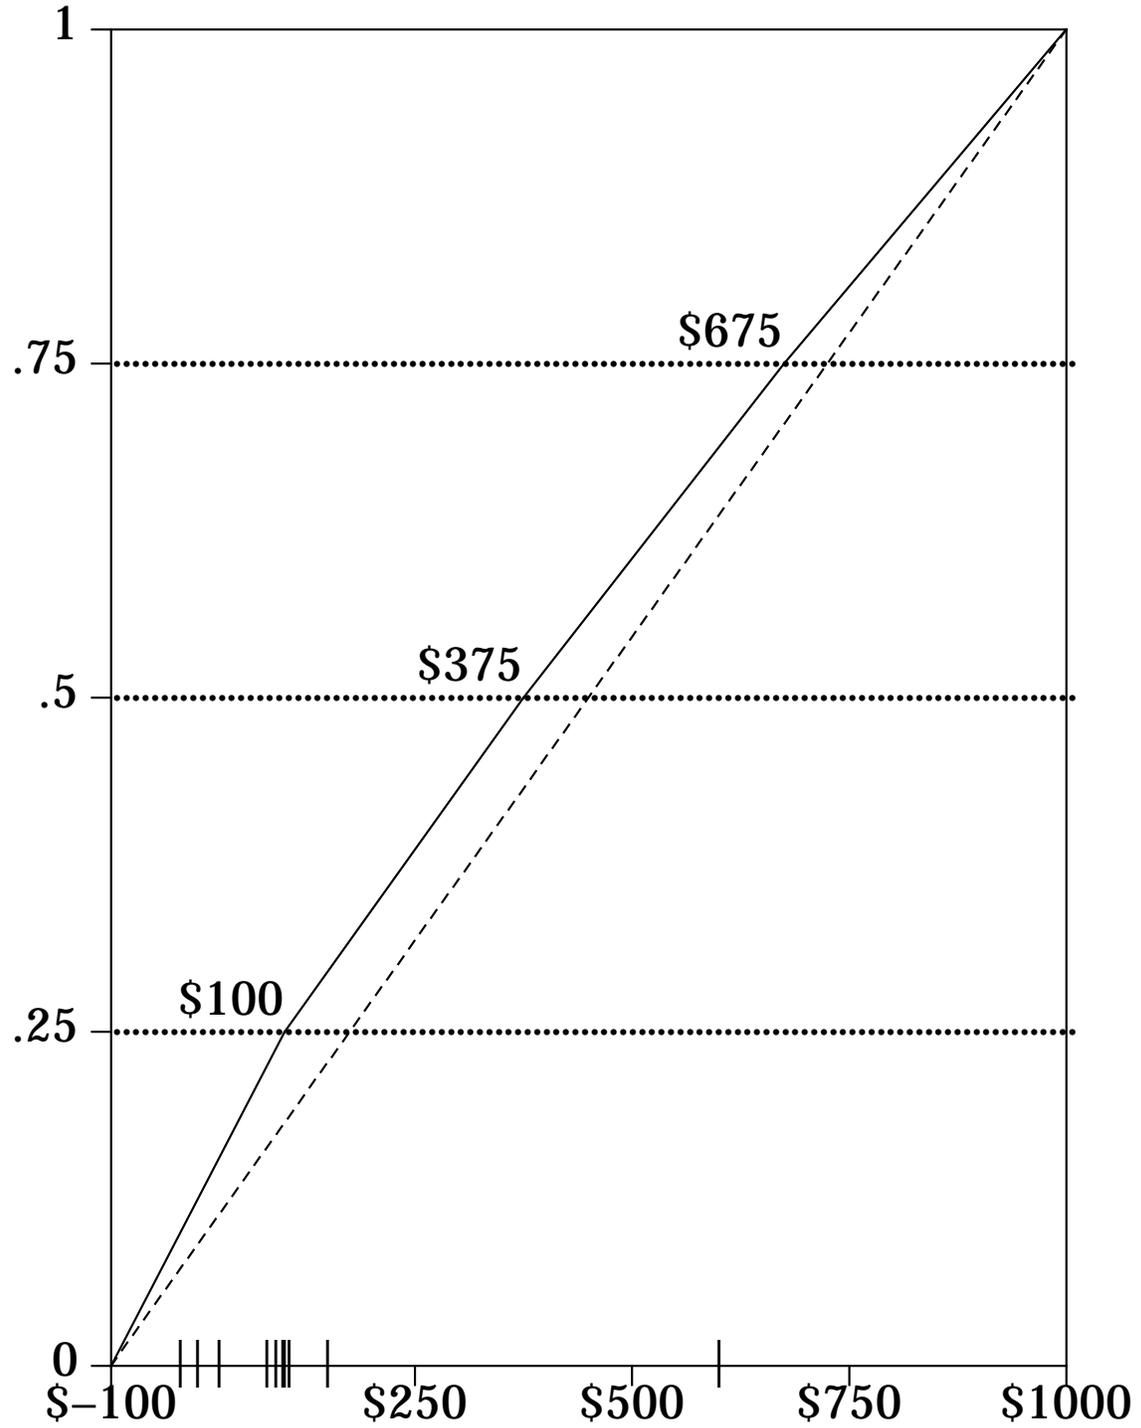
- and the utility of the fourth (over $\$100$ and $\$675$, c.e. $\$375$) is

$$\frac{1}{2} \times 0.25 + \frac{1}{2} \times 0.75 = 0.5.$$

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 I've joined the five points with straight lines, to get an approximation for my utility function.

Iterate.

Utility



The Utility Curve

We can use the utility function as plotted to calculate the utilities of the four lotteries, A, B, C, and D, of the last lecture.

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

Dollars	Utility
\$-100	0.000
\$-20	.100
\$0	.125
\$25	.156
\$80	.225
\$90	.238
\$98	.248
\$100	.250
\$105	.255
\$150	.295
\$600	.69
\$1000	1.0

- The expected 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 I would choose lottery D.

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).

2.15 The Decision-Analysis Procedure

We can break the procedure into:

1. a deterministic phase—values, trade-offs, time-discounting with no uncertainty—and
2. a probabilistic phase—uncertainty, risk aversion, and lotteries.

1. The Deterministic Phase

- define the decision: *“What decision must be made?”*
- identify the alternatives by listing all possible eventualities: *“What courses of action are open to us?”*
- assign values to outcomes: *“How to determine which outcomes are good and which are bad?”*
- select state variables: *“If you had a crystal ball, then what numerical questions would you ask it about the outcome in order to specify your profit measure?”*
- how are the state variables related to each other: a profit function or cost function
- the rate of time preference, the discount rate: *“How does profit received in the future compare in value with profit received today?”*

then the following analysis:

- a. determine *dominance* in order to eliminate alternatives, if possible
- b. determine *sensitivities* in order to identify crucial state variables
Influence Diagrams? — See Clemen & MDM

2. The Probabilistic Phase

- *encode uncertainty* on crucial state variables; assign probabilities to the possible levels of each state variable.
Analysis: develop a *profit lottery*.
- *encode risk preference*; develop a utility function on \$, in practice, having calculated the expected returns, using the values and probabilities, subtract from each an appropriate risk premium.
Analysis: select the best alternative, the one with the highest risk-adjusted return.

3. The Post-Mortem Phase

Analysis:

- a. determine the *value of eliminating uncertainty* in crucial state variables: the maximum value of (im)perfect information
- b. develop the most economical information-gathering program: reductions in which uncertainties are most valuable?

Three reasons for using Decision Analysis in your decision making:

- Decision Analysis (DA) requires you to *think through the structure* of the decision to be made: the structure includes the possible outcomes, your alternative courses of action, and the areas of uncertainty.
- DA will help you make the most *prudent decisions* in cases where you don't know with any certainty the outcomes of the different courses of action open to you.

There can be no guarantees that the best decision will invariably give the best outcome—"the best-laid plans o' mice an men gang aft agley," as the poet Robbie Burns reminds us—so we should distinguish between *good decisions* and *good outcomes*.

Prudent decisions will in the long run give you better odds of good outcomes than will ad hoc decisions.

- By using the DA method, you'll be able to *put a maximum value on additional information* to reduce the uncertainty in the decision-making process, before knowing what it is.

This will put limits on the amounts you should spend in the process of making specific decisions.

2.16 Implementation

DA is increasingly used in business.

- In the U.S. it has been used by Du Pont, Ford, Pillsbury, General Mills, and General Electric.
- The Pillsbury Company used it to determine whether the company should switch from a box to a bag as a package for a certain grocery product.
- General Electric used it to increase by a factor of twenty the R & D budget for a new product.
- Major consulting firms have reported numerous applications for their clients.

To increase the use of DA:

- The CEO must see decision analysis as useful and necessary. His or her support is a must.
- The key executives involved must understand what DA can do for them.
- Chose a problem and run a trial to show the usefulness of the method.
- Develop inside specialists so that staff expertise is available.
- Keep DA relevant and meaningful: don't get lost in the techniques, which should never dominate the analysis — only the decision makers should.
(Software helps — see Treeage's DATA in the student lab.)

