Risk Attitudes and Utility Functions

Marek J. Drużdżel

Wydział Informatyki

Politechnika Białostocka

m.druzdzel@pb.edu.pl

http://aragorn.wi.pb.bialystok.pl/~druzdzel/

Overview

- Problems with the expected value
- Risk attitudes
- Utility axioms and the Expected Utility Theory
- Measurement of utility
- Risk tolerance
- Paradoxes of the expected utility theory

The Need for Measuring Utility

“Immeasurables”

Some things are difficult to express in numerical terms.

Imagine that you are a juror. How much is it worth to condemn an innocent man or to release a guilty one?

How do you judge money vs. health or happiness?

http://www.nourishingrelationships.blogspot.com/2012/04/attaining-happiness-without-winning.html

"Immeasurables"

It is very convenient to be able to measure things.

How did we ever manage without thermometers ☺?

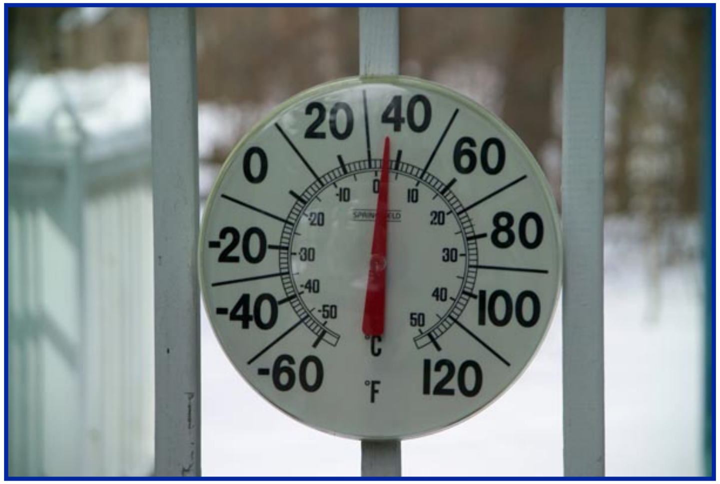

http://auntdisexperimentallife.blogspot.com/2011/07/thermometer-or-thermostat.html

http://www.theguardian.com/society/2012/jun/15/fake-thermometers-seized-meningitis

Problems with Expected Value

Problems with expected value

Even if you can express "immeasurables" in numbers, there are problems with expected value, found quite a while ago (even though probability is quite young).

Bernoulli (17^{th} century) pointed out these problems and the need to have some measure of preferences.

Then there was long nothing, just a qualitative, ordinal notion (note the gymnastics around qualitative notion of utility in economics) and finally a quantitative, cardinal utility in 1940s due to von Neuman & Morgenstern.

Problems with expected value

Who would call a pauper foolish for selling a lottery ticket paying $20,000,000 tomorrow with p=0.5 for $9,000,000 today?

http://magu1988.wrzuta.pl/obraz/9DNFVoalzoa/bezdomny

A “modernized” version of an argument made by Bernoulli in 1738

St. Petersburg's paradox

(also due to Bernoulli)

Imagine a game that involves flipping a coin infinitely many times and that pays progressively more for reaching each step.

If you get just one heads (p=0.5), you get $2, if you get two heads in a row (p=0.25), you get $4, if you get three heads in a row (p=0.125), you get $8, etc.

You can't lose and you can win by playing.

What is the fair price for a ticket to play?

How much are you willing to pay to participate in this game?

http://www.eagleonline.com/you-can-please-some-of-the-people/

St. Petersburg's paradox: Expected value

If you get just one heads (p=0.5), you get $2, if you get two heads in a row (p=0.25), you get $4, if you get three heads in a row (p=0.125), you get $8, etc.

The expected value of this game is:

$$EV = \sum_{i=1}^{\infty} 2^i \left(\tfrac{1}{2}\right)^i = \sum_{i=1}^{\infty} 1 = \infty$$

http://www.eagleonline.com/you-can-please-some-of-the-people/

The expected value of playing this game is infinity!

Now that you know it, what would you be willing to pay to participate in this game?

Bernoulli's solution

The solution proposed by Bernoulli is that, what he called "moral worth" of a quantity, is different from that quantity.

He introduced the law of diminishing marginal utility and proposed the logarithm function as one that satisfies this law. (Just take the logarithm of the value to get the utility.)

Bernoulli, Daniel; Originally published in 1738 in the *Commentaries of the Imperial Academy of Science of Saint Petersburg*

Translated by Dr. Louise Sommer. (January 1954). "Exposition of a New Theory on the Measurement of Risk". Econometrica, 22 (1): 22–36.

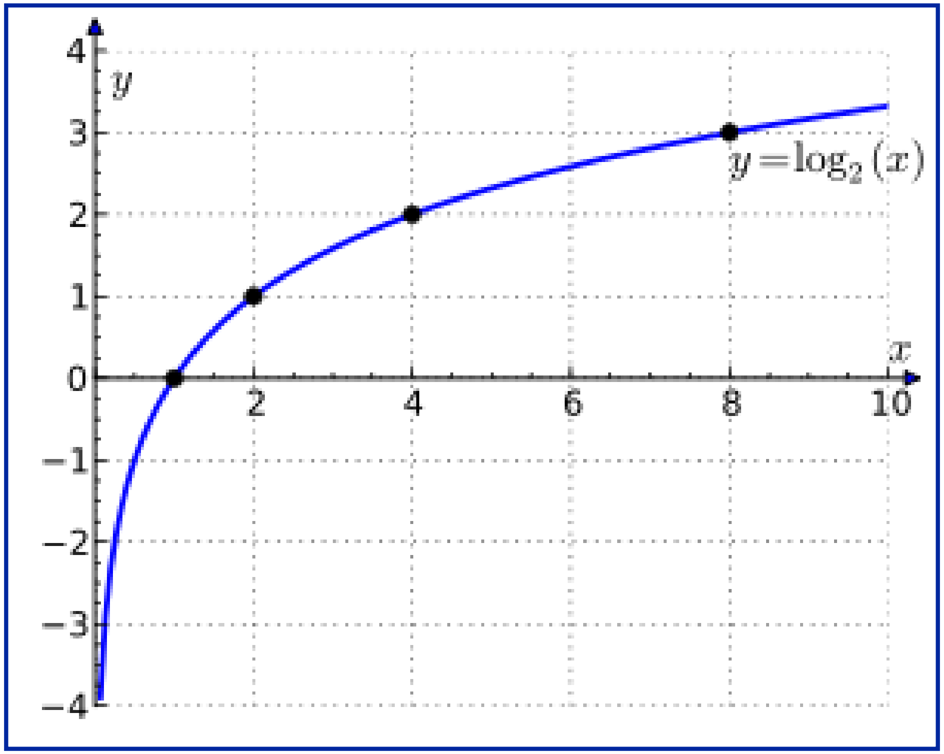

http://en.wikipedia.org/wiki/Logarithm/

The Law of Diminishing Concern ☺

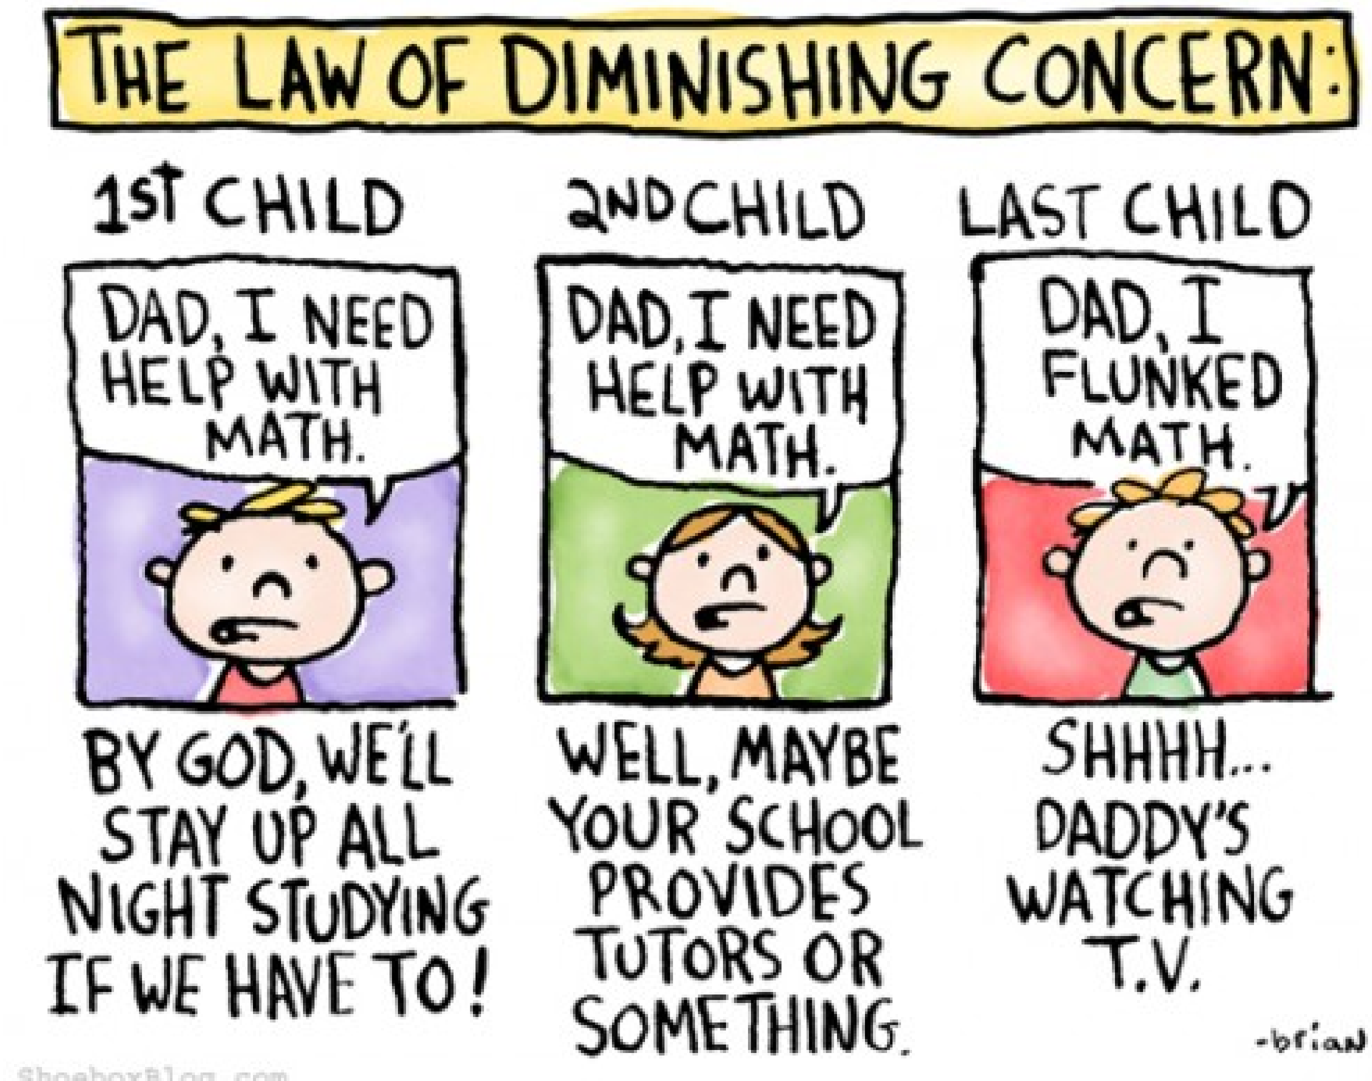

http://en.wikipedia.org/wiki/Logarithm/

Some history

Then there was long nothing, just a qualitative, ordinal notion (note the gymnastics around qualitative notion of utility in economics) and finally a quantitative, cardinal utility in 1940s due to von Neuman & Morgenstern.

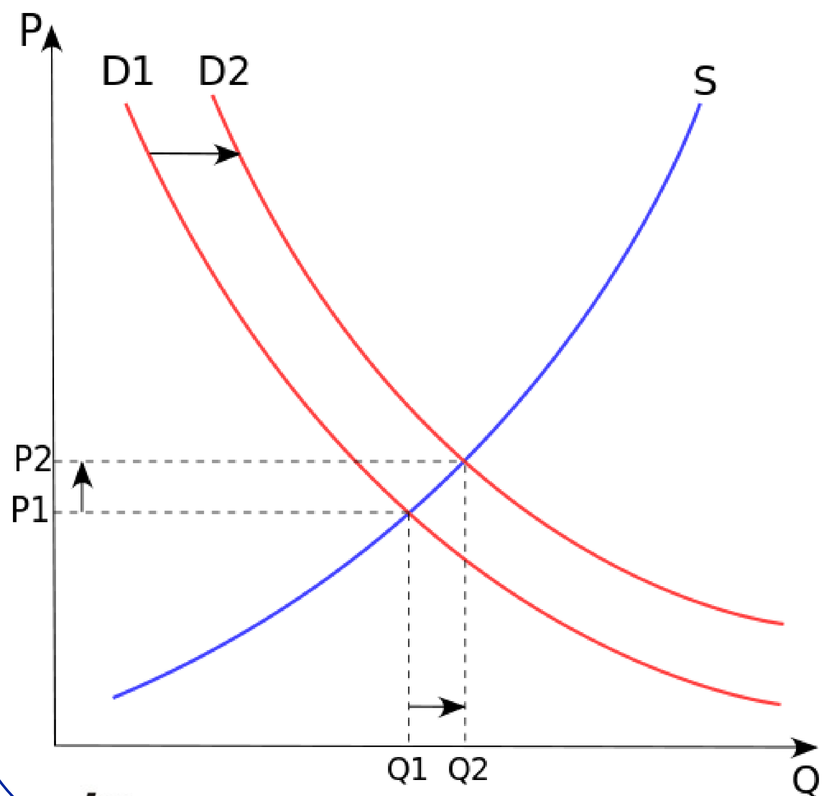

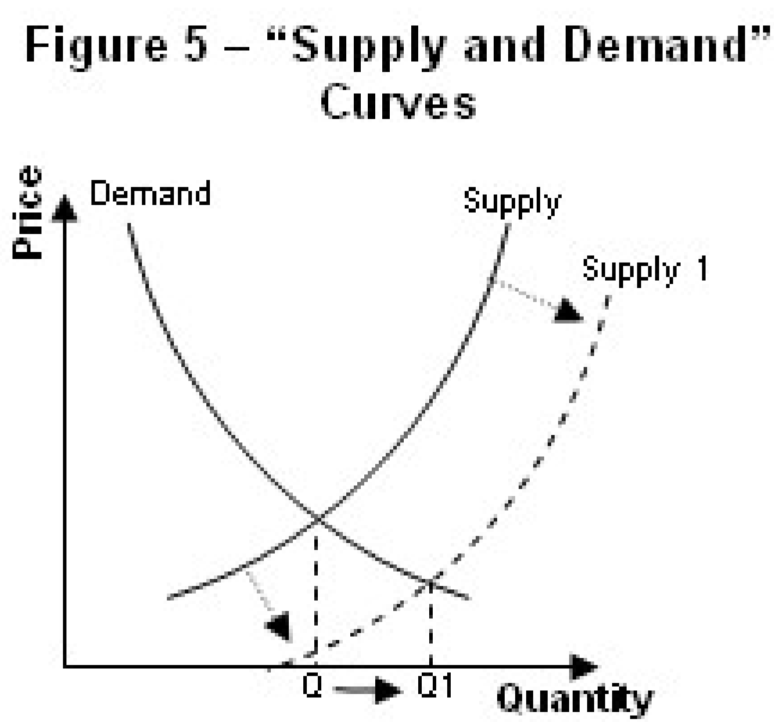

http://en.wikipedia.org/wiki/Supply_and_demand

Risk Attitudes

Risk attitudes

- Three theoretical attitudes: risk neutrality, risk seeking, and risk aversion.
- Easy to understand in terms of the second derivative of the utility function:
  - If each additional dollar is worth more to you than the one before, you are out to win big and you are willing to take risks
  - If the value of each additional dollar is worth less than the last dollar, then you are risk averse.

Risk attitudes: Risk aversion

If the value of each additional dollar is worth less than the last dollar, then you are risk averse.

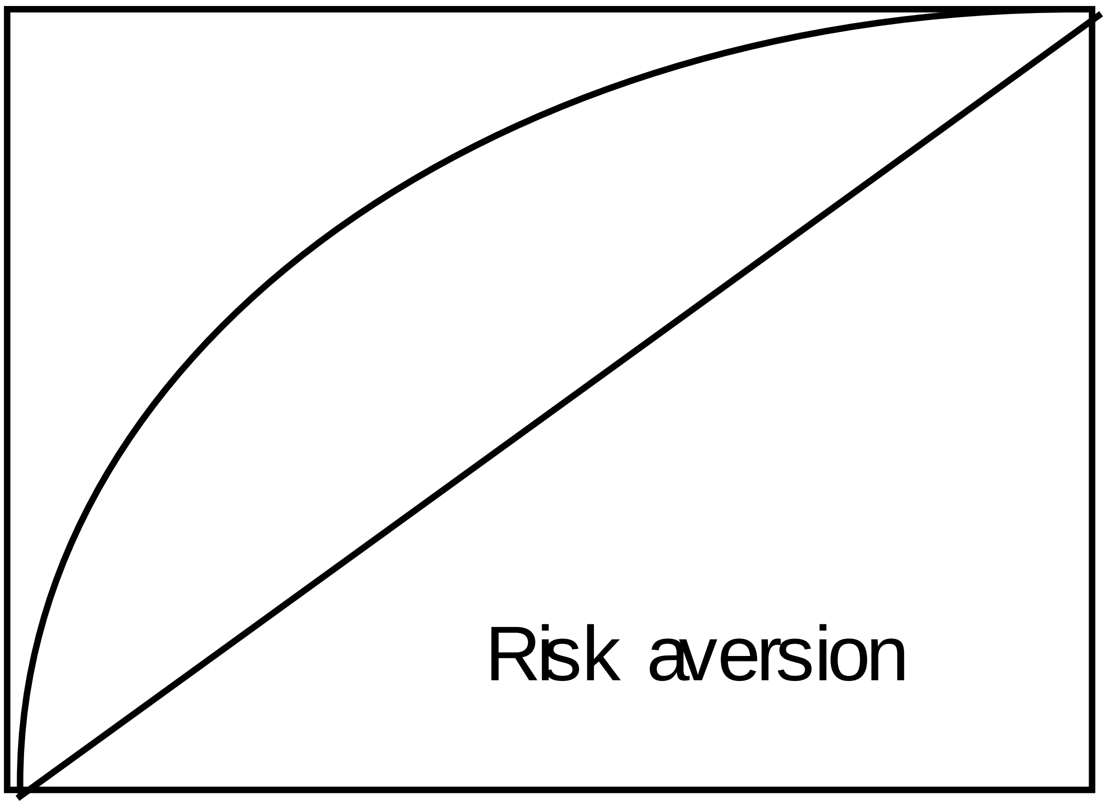

e.g., people buying flood or health insurance

Risk attitudes: Risk aversion

If each additional dollar is worth more to you than the one before, you are out to win big and you are willing to take risks

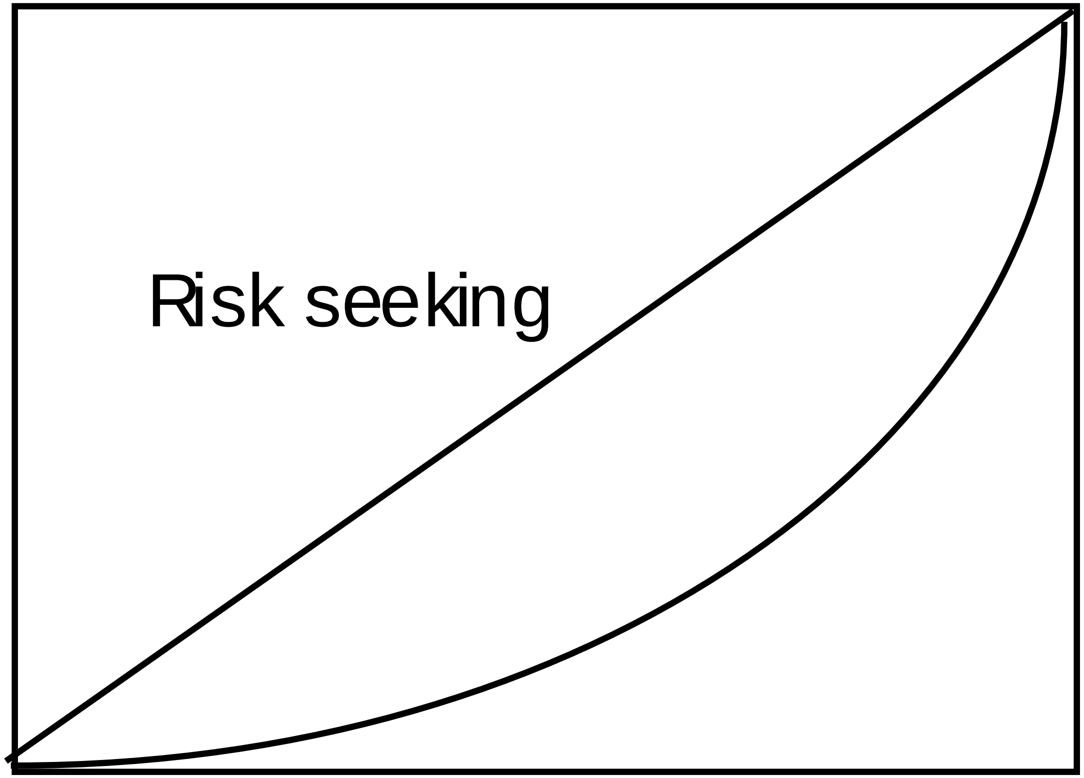

e.g., lottery players

Certainty equivalent

Certainty equivalent of a gamble: How much would you pay for an opportunity to participate in this gamble?

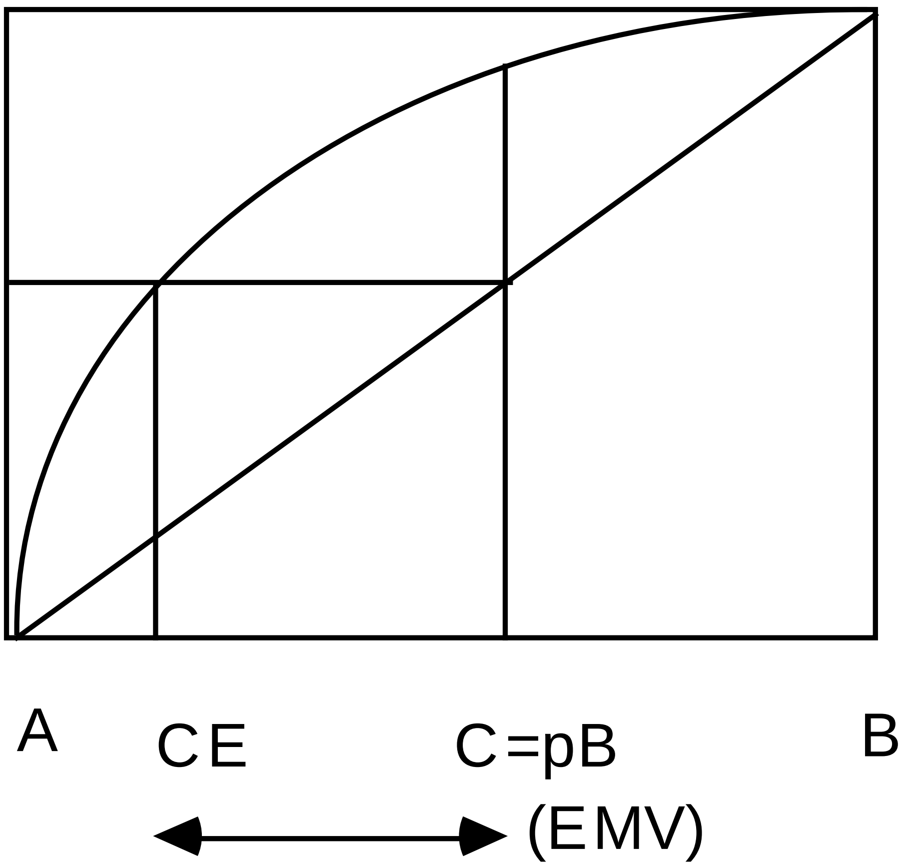

Risk premium can be positive or negative (How will the picture look for somebody who is risk prone ☺?).

Note that each of CE, EMV and RP are not in terms of utility but rather in terms of the quantity that we are measuring.

Risk Premium = EMV – CE

Certainty equivalent

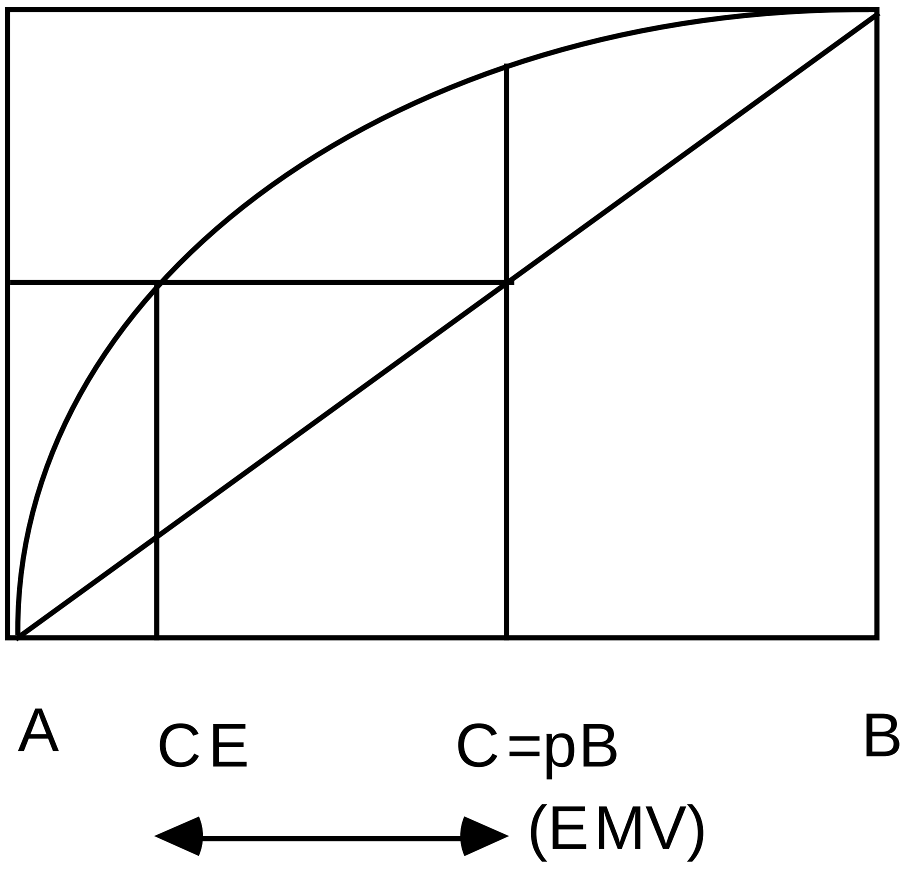

Restatement: If the gamble is worth to you less than the expected value, then CE is to the left of EV.

Restatement: If the utility of a value is higher than the utility of that value when it is only expected, then we are dealing with risk aversion.

Typical utility function for humans

We tend to be:

- Risk neutral for small amounts
- Risk averse in the domain of gains
- Risk seeking in the domain of losses

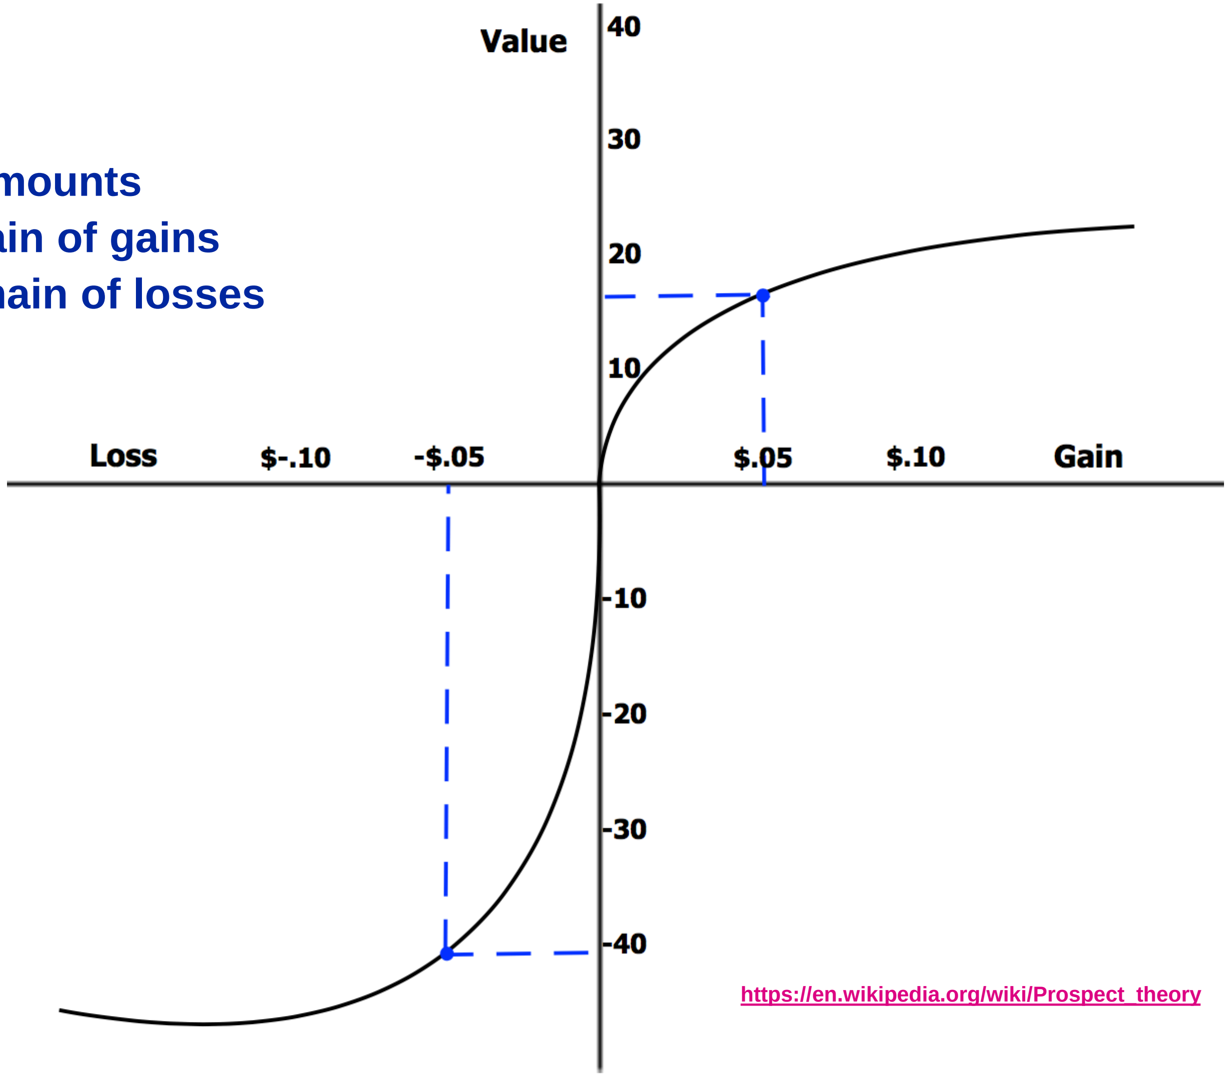

https://en.wikipedia.org/wiki/Prospect_theory

Expected Utility Theory

Axioms of Expected Utility Theory

How does mathematics work?

We start with assumptions (axioms) and then prove theorems.

The theorems will be useful if the assumptions that we have made are reasonable and bear on reality.

Expected Utility Theory (EUT) works the same way.

Axiomatization proposed by a mathematician, John von Neumann and an economist, Oscar Morgenstern in 1940s.

So, let us examine whether the axioms make sense.

Axiom 1: Orderability

(A>B) or (B>A) or (B ~ A)

“You know what you want”

Axiom 2: Transitivity

$$(A>B) \text{ and } (B>C) \Rightarrow (A>C)$$

Axiom 3: Decomposability

$$[p,A;\ 1-p,[q,B;1-q,C]] \sim [p,A;\ (1-p)q,B;\ (1-p)(1-q),C]$$

"Compound lotteries can be reduced to simpler ones by the laws of probability (no fun in gambling)."

Axiom 4: Continuity

$A > B > C \Rightarrow \exists\ p,\ [p, A;\ 1-p, C] \sim B$

"There exists a gamble with odds that will make you indifferent between choosing B for sure and playing it."

Axiom 5: Substitutability

$$A \sim B \Rightarrow [p,A;\ 1\text{-}p,C] \sim [p,B;\ 1\text{-}p,C]$$

"If you are indifferent between two lotteries A and B, then you will be indifferent between more complex lotteries involving something else."

Axiom 6: Monotonicity

$$A > B \Rightarrow (p \geq q \Leftrightarrow [p,A;1-p,B] > [q,A;1-q,B]$$

"All things being equal, you prefer the lottery that gives you a higher probability of getting the more desirable outcome."

Axiom 7: Invariance

“Need only probabilities and utilities.”

Axiom 8: Boundedness

“No outcomes are infinitely good or infinitely bad.”

Expected Utility Theorem

From these eight axioms, von Neumann and Morgenstern prove a theorem that essentially states that we can describe preferences of a person who adheres to the axioms in terms of a numerical utility function U such that:

U(A)>U(B) ⇔ A preferred to B

U(A)=U(B) ⇔ A ~ B

Utility axiomatized as above allows for decision making in a way that is consistent with maximizing expected utility. (In other words, the utility combines like the probabilistic expectation.)

Utility

Utility is a peculiar measure with no scale and no zero point.

If U(x) is a utility function, then U'(x)=aU(x)+b is also a utility function, i.e., utility is determined up to a linear transformation

Any other measures that behave like this ☺?

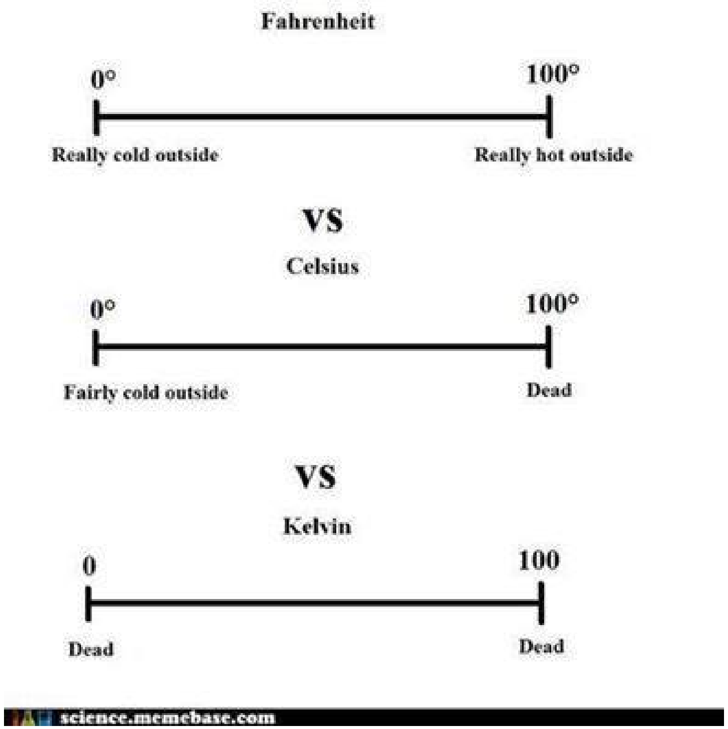

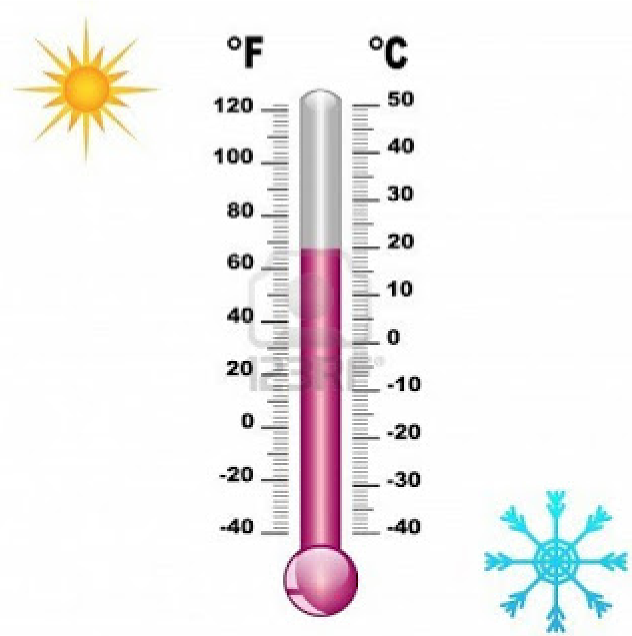

Utility Measurement

Foundations of decision analysis

The foundation of decision analysis (assumption but confirmed by observations):

Humans can provide reliably the structure of a problem and reliable numbers (judgments) but are weak in combining these

Hard to quantify?

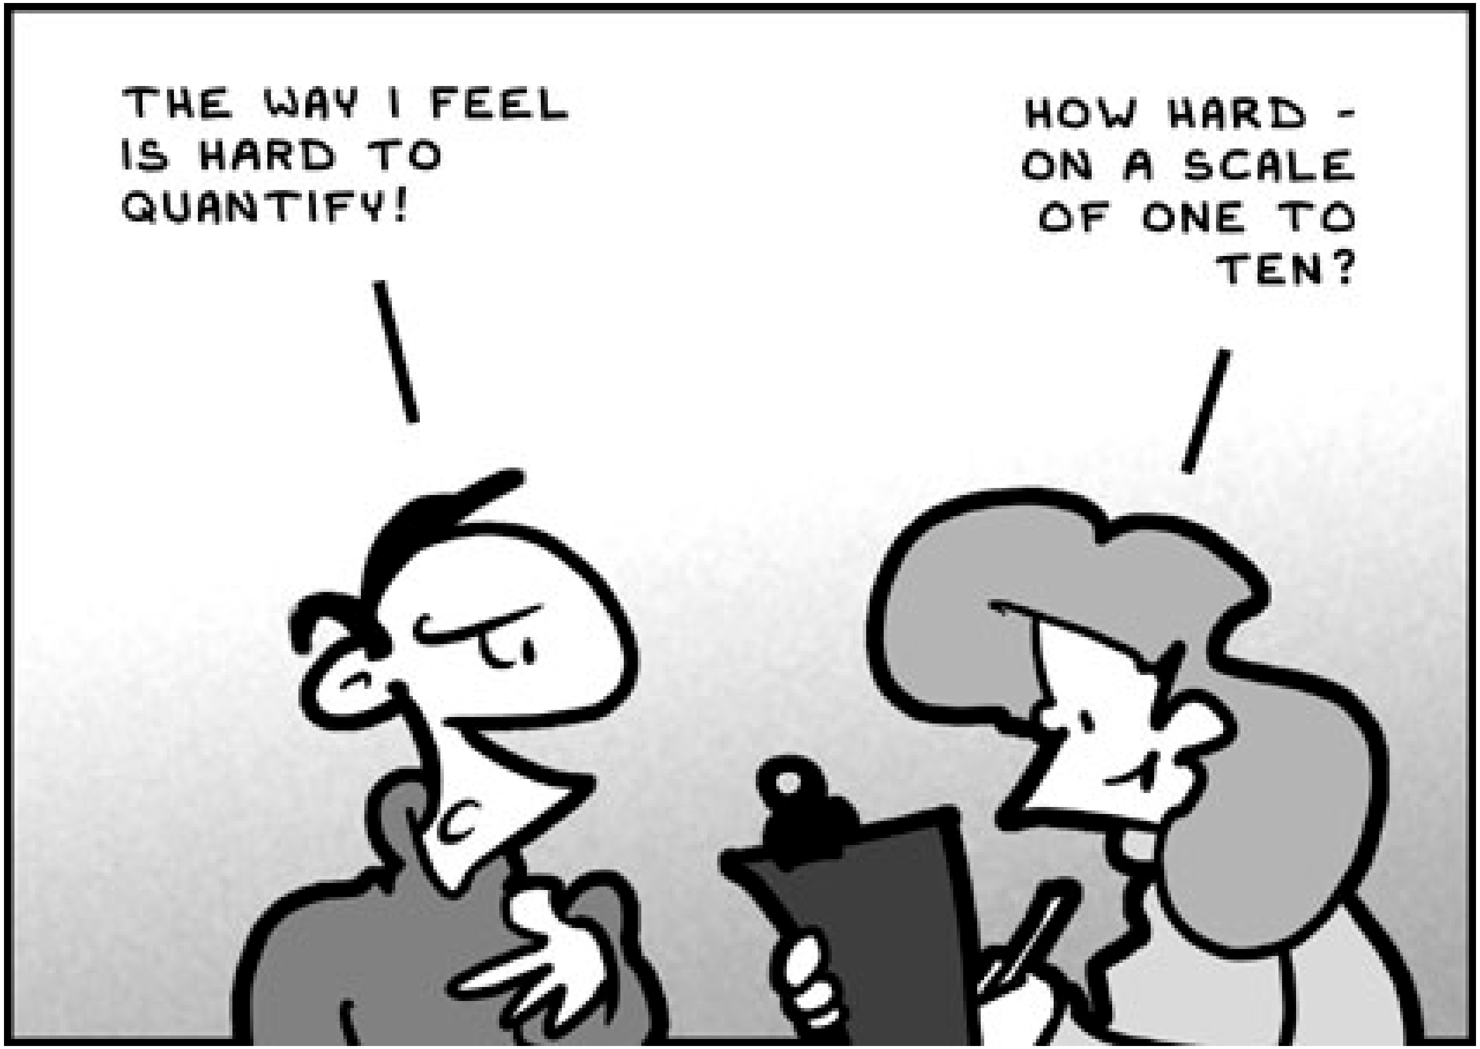

Measurement of utility

We have four variables: p, CE, G, L.

Two are for free and determined by the axioms (these are the lower and the upper bounds of the utility range).

We need to fix (preset) the third and then obtain the fourth.

CE method: fix G, L, and p, assess CE

PE method: fix G, L, and CE, assess p

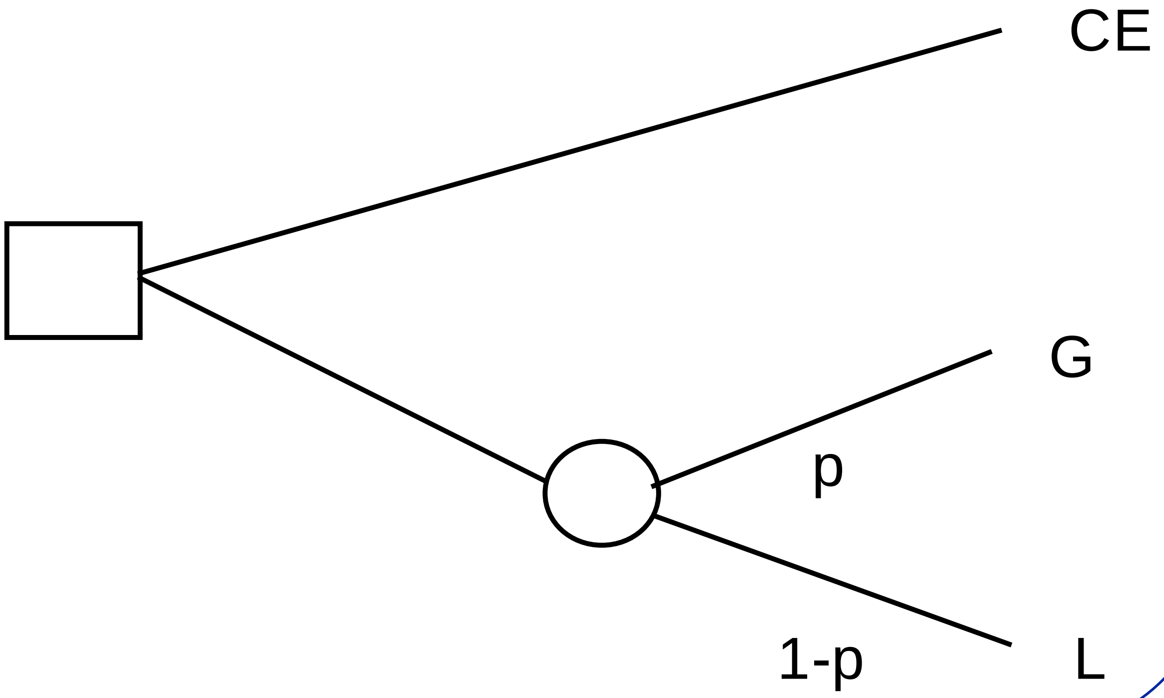

Measurement of utility: Probability equivalent

This is handy if we want to obtain the utility of a given value.

Choose the worst and the best and use them for setting the boundaries of the interval.

This is kind of counterintuitive for us, as for every bad outcome there is always one that is worse.

But don't forget that we are reasoning within a model and it's the worst possible <u>within this model</u>.

Remember about the clarity test!

Measurement of utility: Probability equivalent

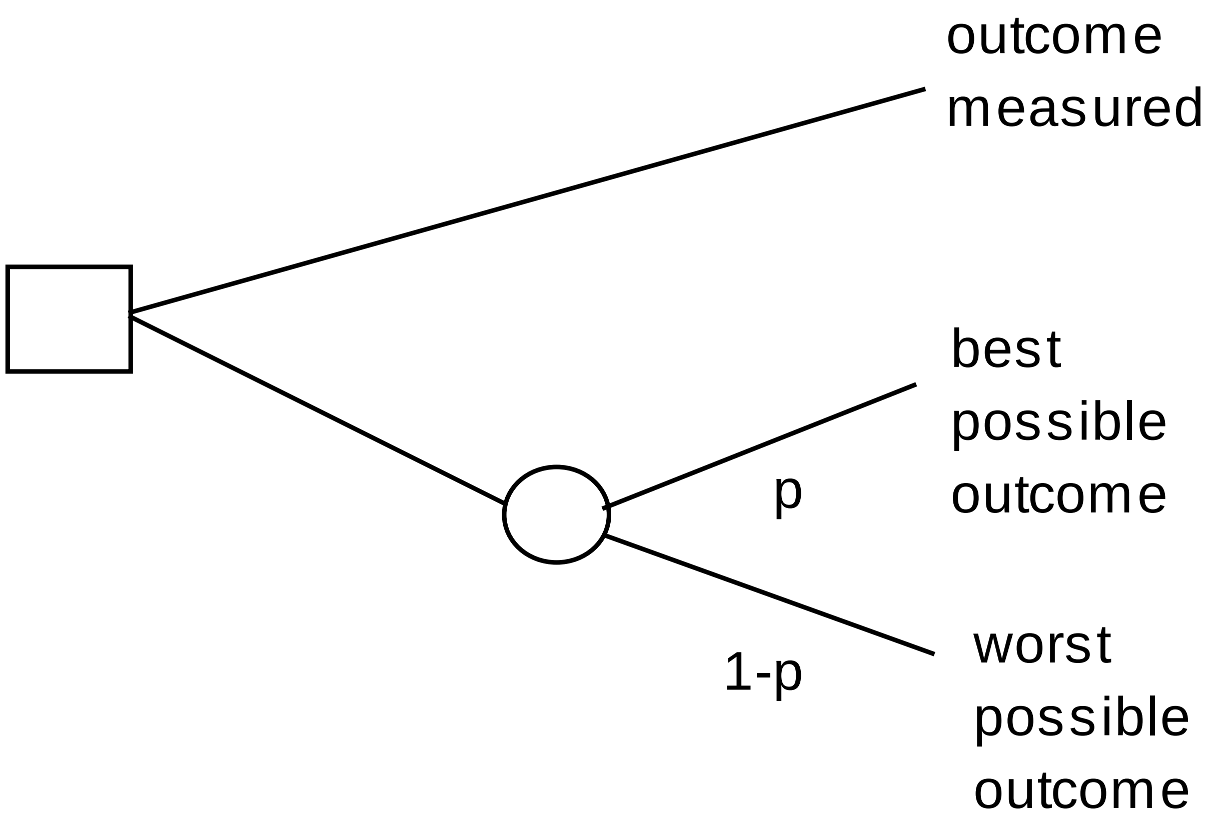

Manipulate p until the decision maker is indifferent between the two choices. Then,

U(Measured) = p U(Best) + (1-p) U(Worst)

U(Measured) = p 100 + (1-p) 0 = p 100

Measurement of utility: Certainty equivalent

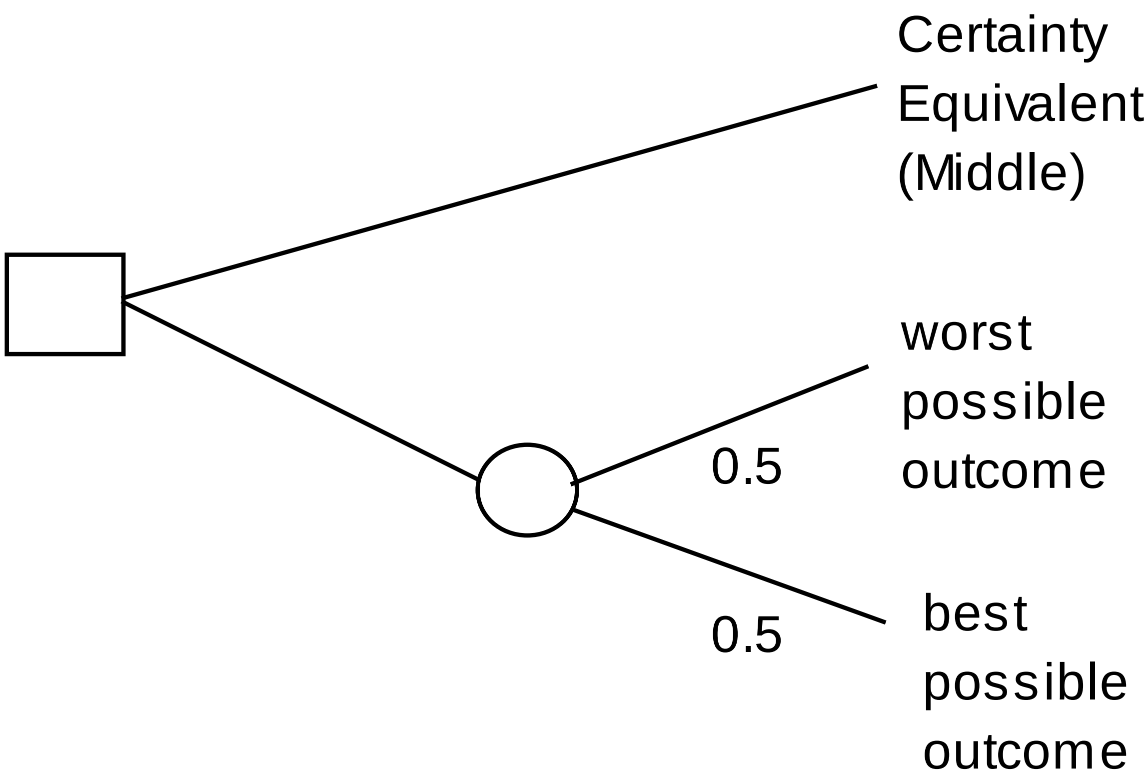

Manipulate CE until the decision maker is indifferent between the two choices. Then,

U(Measured) = 0.5 U(Best) + 0.5 U(Worst)

U(Measured) = 50

Measurement of utility: Comparison of PE and CE

CE leads to more risk-averse responses in gains and risk seeking in losses.

In PE, p=0.5 is the best as people exhibit probability distortions at more extreme probabilities (certainty effect). One possible answer to this problem is to use the following lotteries:

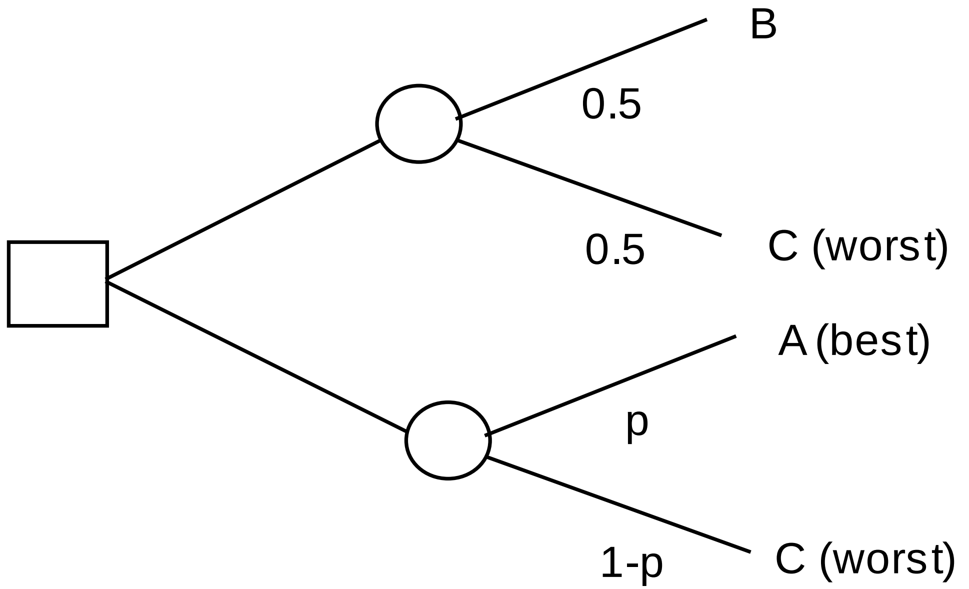

This is known as The McCord-De Neufville utility assessment procedure.

Risk Tolerance

Measurement of utility: Risk tolerance

Measurement of utility: Risk tolerance

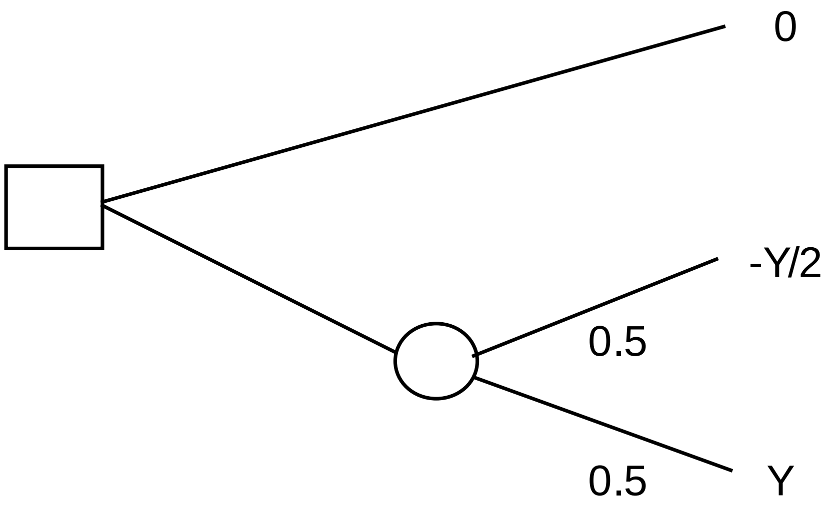

–50	–500	–2,500	–5,000	–10,000
100	1,000	5,000	10,000	20,000

Max Y for which the DM is indifferent =(def)= Risk Tolerance = R

Use the following form of the exponential utility function:

u(x) = 1 - e^(-x/R)

Measurement of utility: Risk tolerance

This is a poor-man's utility function and one can argue whether it models well a DM's preferences.

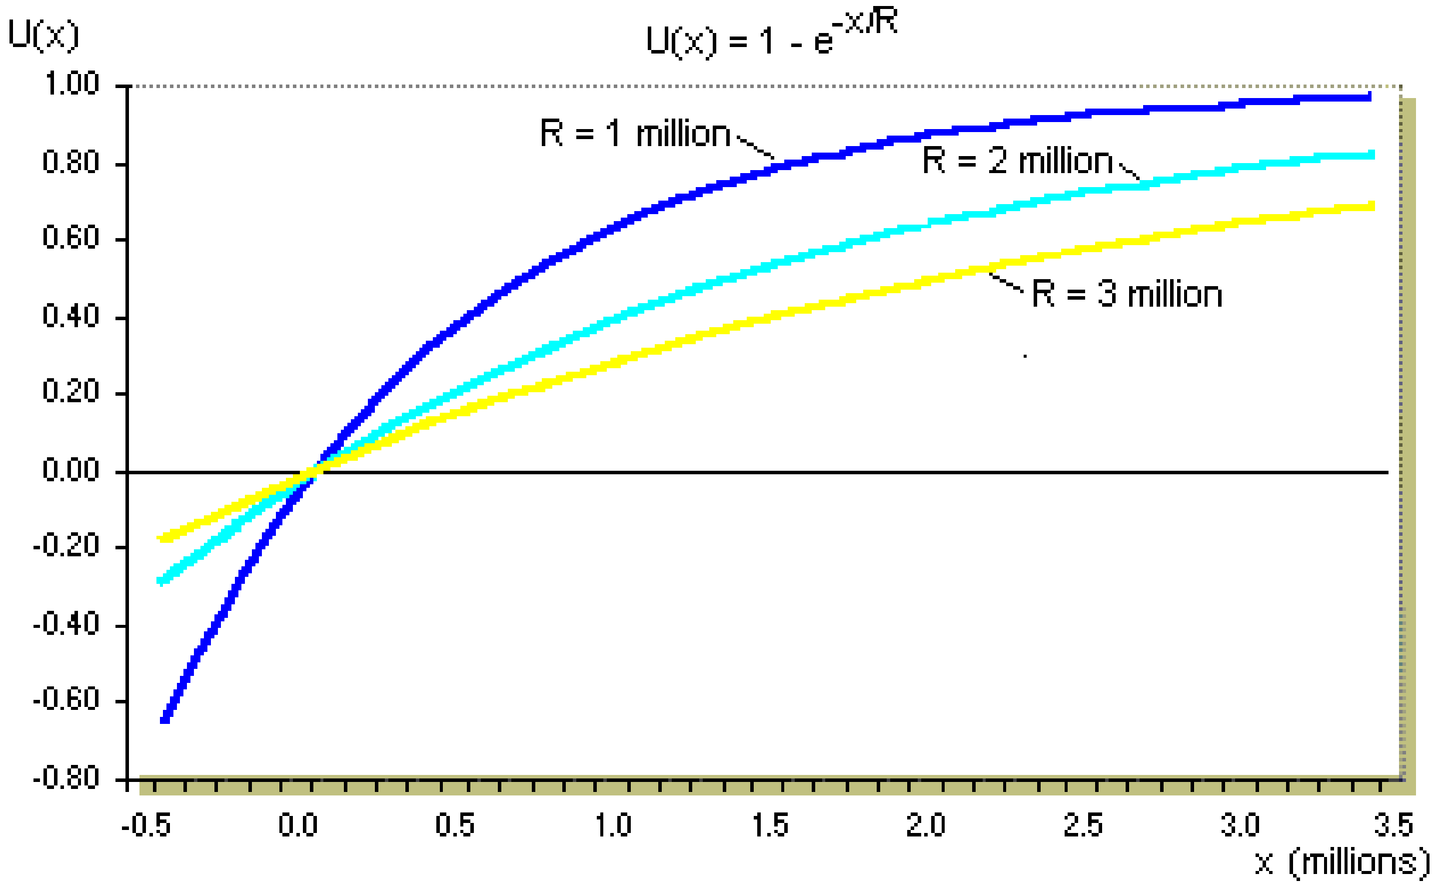

Measurement of utility: Risk tolerance

It is certainly useful as a first-cut approximation in cases when we want to model risk aversion.

A quick sensitivity analysis can determine a critical risk tolerance, and the decision maker can be asked, via a simple assessment question whether his/her risk tolerance exceeds the critical value.

If the choice is clear, then there is no need for further preference modeling.

If the choice is not clear, it may be a good idea to assess a utility function more carefully.

Paradoxes of the Expected Utility Theory

Paradoxes of the Expected Utility Theory

Starting from 1950s throughout now the axioms of expected utility and "paradoxical" behavior relative to the axioms have generated many debates.

Behavioral research has focused on these paradoxes, those situations in which reasonable and thoughtful people behave in ways inconsistent with the axioms.

The paradox exists because careful explanation of the inconsistency often does not lead such people to modify their choices.

Paradoxes: Allais paradox

Maurice Allais, proposed it around 1953

You have two decisions to make:

Decision 1

A: Win $1M with probability 1.

B: Win $2M with probability 0.10.

Win $1M with probability 0.89.

Win $0 with probability 0.01.

Decision 2:

C: Win $1M with probability 0.11.

Win $0 with probability 0.89.

D: Win $2M with probability 0.10.

Win $0 with probability 0.90.

Paradoxes: Allais paradox

82% of the people choose A>B, while 83% choose D>C.

How does this violate the axioms?

Let U(0)=0, U(2M)=1.

From the preference of A over B we have

$$EU(A) > EU(B)$$

$$U(1M) > 0.1\ U(2M) + 0.89\ U(1M) + 0.01\ U(0)$$

i.e., $0.11\ U(1M) > 0.1\ U(2M)$

But from the preference of D over C we have exactly the opposite

$$EU(C) < EU(D)$$

$$0.11\ U(1M) + 0.89\ U(0) < 0.1\ U(2M) + 0.9\ U(0)$$

i.e., $0.11\ U(1M) < 0.1\ U(2M)$

Paradoxes: Ellsberg paradox

A barrel contains a mixture of 90 red, blue, and yellow balls. Thirty of the balls are red, and the remaining 60 are a mixture of blue and yellow, but the proportion of blue and yellow is unknown. A single ball will be taken randomly from the barrel.

Suppose, you are offered the choice between the gambles A and B:

A: Win $1000 if a red ball is chosen.

B: Win $1000 if a blue ball is chosen.

and then C and D:

C: Win $1000 if a red or a yellow ball is chosen.

D: Win $1000 if a blue or a yellow ball is chosen.

Which do you prefer?

Paradoxes: Ellsberg paradox

Most people prefer A over B and then D over C.

How does this violate the axioms?

Let p be the proportion of yellow balls in the urn.

Choice A over B gives us:

$$1/3\ U(1000) > p\ U(1000).$$

The choice of D over C gives us:

$$(1/3 + 2/3 - p)\ U(1000) < 2/3\ U(1000),$$

i.e., $1/3\ U(1000) < p\ U(1000)$**,**

which is exactly the opposite.

Paradoxes of the Expected Utility Theory

- Much work has been conducted in the direction of twisting the axioms and proposing alternative theories of utility.
- At the very fundamental level we would like to choose a set of axioms that are compelling; they make sense as guiding principles for decision making.
- On the basis of these axioms, then, we derive a decision rule. The decision rule that provides the basis for addressing complex decisions, the choice for which may not be obvious. Expected utility theory, based on the above axioms, has been the standard for over fifty years.
- In fact, the axioms of expected utility theory provide the basis for decomposing hard decisions into a structure that consists of decisions, uncertain events, and outcomes that can be valued independently of the "gambles" in which they may occur.
- Thus, the entire decision analysis approach has been dictated, at least implicitly, by the axioms.

Paradoxes of the Expected Utility Theory

- At first glance, the axioms of expected utility do seem compelling.
- Deeper inspection of the axioms, though, has led a number of scholars to question whether the axioms are as compelling as they might seem.
- The sure-thing principle, in particular, has been called into question, as has the transitivity axiom.
- In spite of these rumblings at the foundations of decision analysis, though, no compelling set of axioms, accompanied by a decision rule and implicit procedure for decomposing large problems, has emerged.

Paradoxes of the Expected Utility Theory

In fact, a fundamental question still exists: Should we change the axioms to make our decision rule consistent with the way people actually do behave?

Or do we leave the axioms and decision rule as they are because we believe that in their current form they provide the best possible guidance for addressing hard decisions?

Lacking answers to these basic questions, axiomatic research and the development of generalized utility models may continue to generate interesting results but without a clear notion of how the results relate to practical decision-analysis applications.

Pascal's wager

Pascal's wager: Should we believe in God or not?

	God exists	God does not exist
believe	eternal salvation	forgo some earthly pleasures in your life
doubt	eternal damnation	enjoy some earthly pleasures in your life

Pascal's wager

Pascal's wager: Decision tree

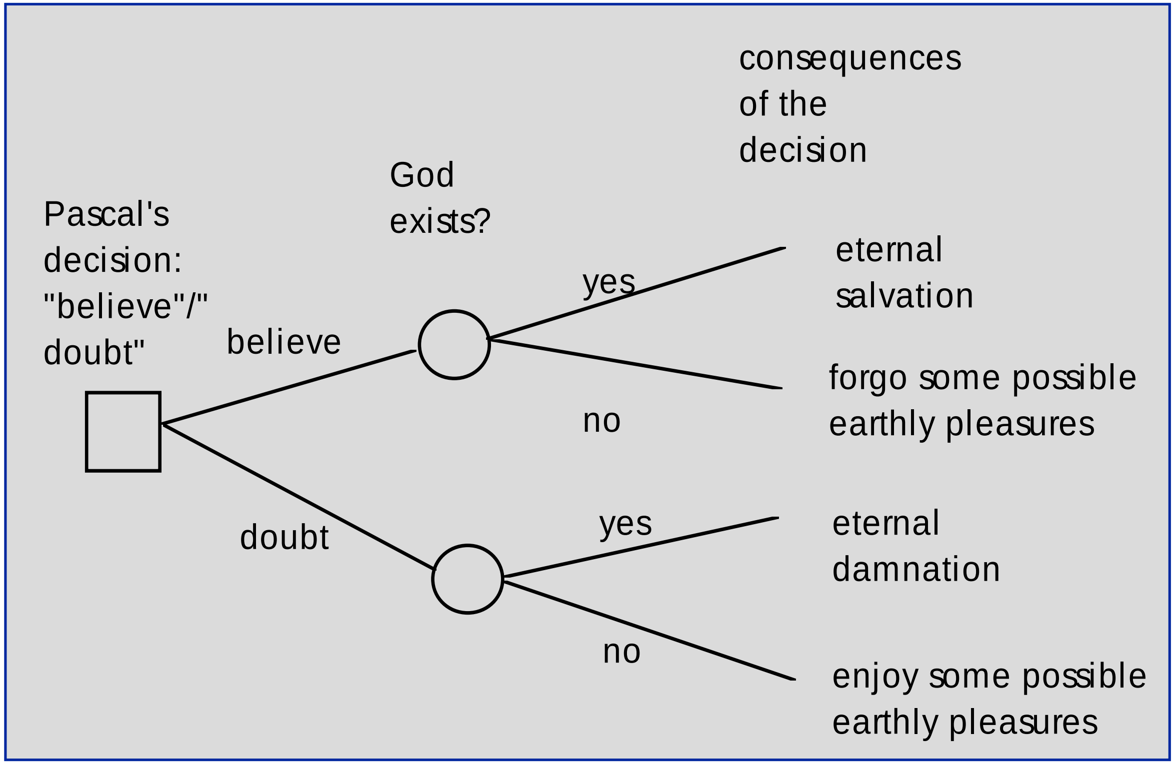

EU(believe) = p ∞ + (1-p)(-ε) = ∞
EU(doubt) = p (- ∞) + (1-p) ε = -∞

The only rational thing is to believe ☺!

Pascal's wager

Where does Pascal's Wager depart from the Expected Utility Theory ☺?