Hubbard
Decision Research

How to Make Decisions Under Uncertainty

Hubbard Decision Research
2 South 410 Canterbury Ct
Glen Ellyn, Illinois 60137
www.hubbardresearch.com

Reviewing Where We Are

- **Overview of Probabilistic Models and Monte Carlo Methods**
- **Computing Information Value**
- **Understanding and Setting Risk Tolerance**

A General Procedure for Measurement

Model The Current State of Uncertainty

AIE quantifies and then optimizes decisions by focusing measurements where it matters most

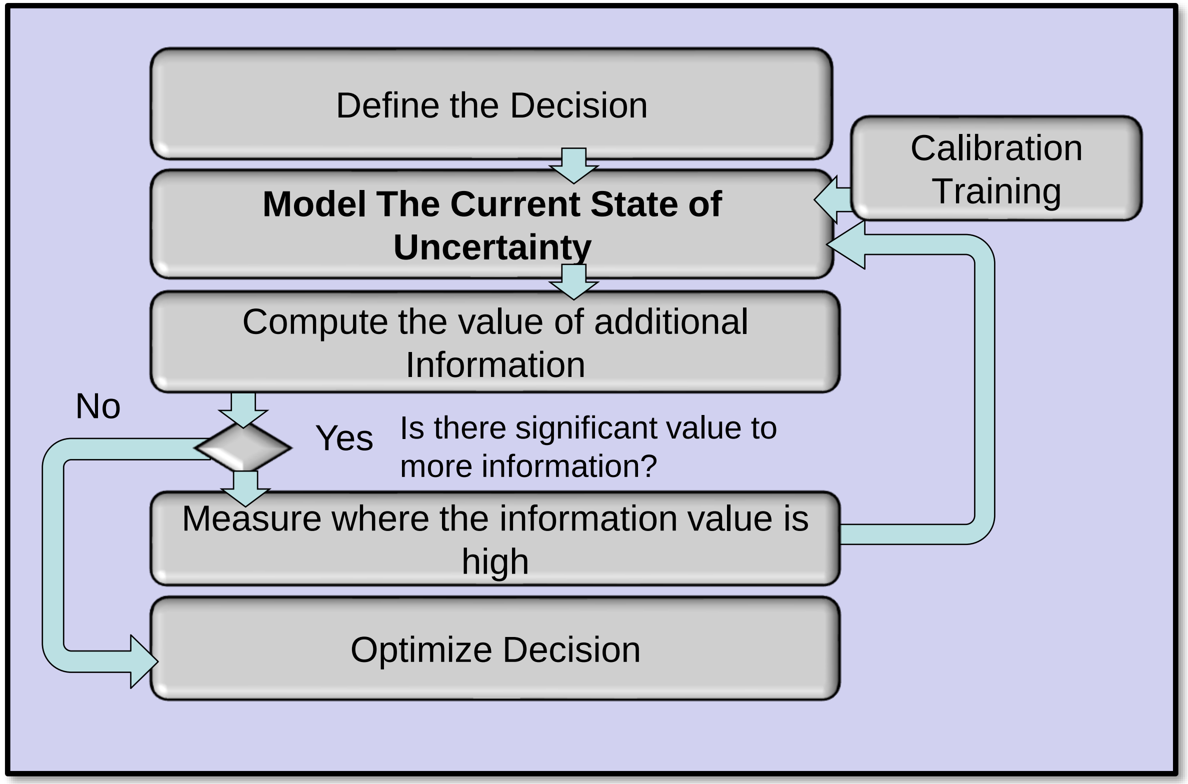

Model The Current State of Uncertainty

Why Model Your Uncertainty

- **Acknowledging Uncertainty** — Uncertainty should be a big factor in making decisions. Explicitly modeling the uncertainty has large impacts on making decisions.
- **Improved Performance** — People who build Monte Carlo models tend to make better forecasts.
- **Avoid The Measurement Inversion** — When people try to reduce their uncertainty, they measure the wrong things.

Model The Current State of Uncertainty

Modeling the UBDM System

1. The clarification discovered that the reason for the measurement was to make a decision about the UBDM system.

2. After we've clarified the decision, we identify variables relevant to the UBDM system:

- $5 million initial investment
- 25% reduction in time spent on document management
- 1,000 staff w/average loaded salary of $90,000 currently spent 15% of time in document management

Model The Current State of Uncertainty

Case A2.1: Conventional Business Case

Unstructured Big Data Management System

Input Variables	Point Est.
Percent reduction in time spent in unstructured data analysis	25%
Number of employees	1,000
Current time spent in unstructured data analysis	15%
Employee loaded cost	$90,000
Initial investment	$5,000,000
Cost of money	6.5%

Do we know these as exact points?

Cash Flow	2020	2021	2022	2023	2024
Total costs	$ 5,000,000				
Total benefits		$ 3,375,000	$ 3,375,000	$ 3,375,000	$ 3,375,000
Net benefit	**$ (5,000,000)**	**$ 3,375,000**	**$ 3,375,000**	**$ 3,375,000**	**$ 3,375,000**

Results	
Net Present Value	**$ 6,161,568**

Single point result

Model The Current State of Uncertainty

Uncertainty in Business Decisions

Risk analysis should be part of your business cases. You don't know all of the assumptions in the business case exactly. We need to be able to quantify that uncertainty.

Ideal Values:
"Points"

Real-world Measures

Normal Distribution

Uniform Distribution

Lognormal Distribution

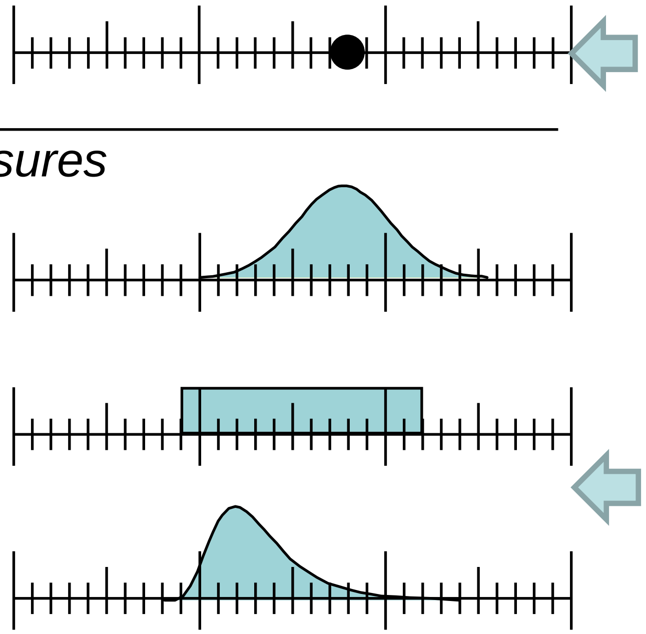

Assumptions: Most values in business cases are represented as exact values – even though exact values are almost never known.

No Assumptions: Most things we DO know are better represented by ranges and probabilities — we don't have to assume anything we don't really know.

Model The Current State of Uncertainty

Constructing a Distribution

Uncertainty about “either/or” events are expressed as “discrete” probabilities (e.g., “35%).

Uncertainty about continuous values can still be thought of as sets of discrete probabilities.

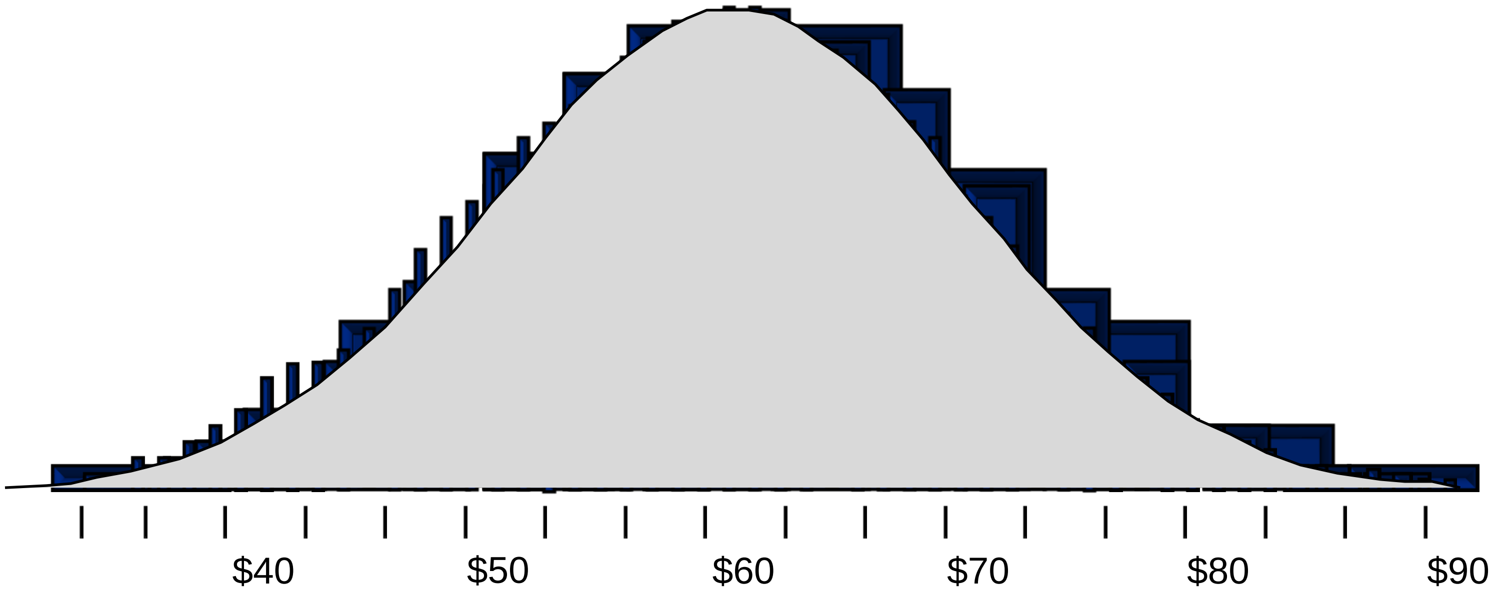

Model The Current State of Uncertainty

Monte Carlo: How to Model Uncertainty in Decisions

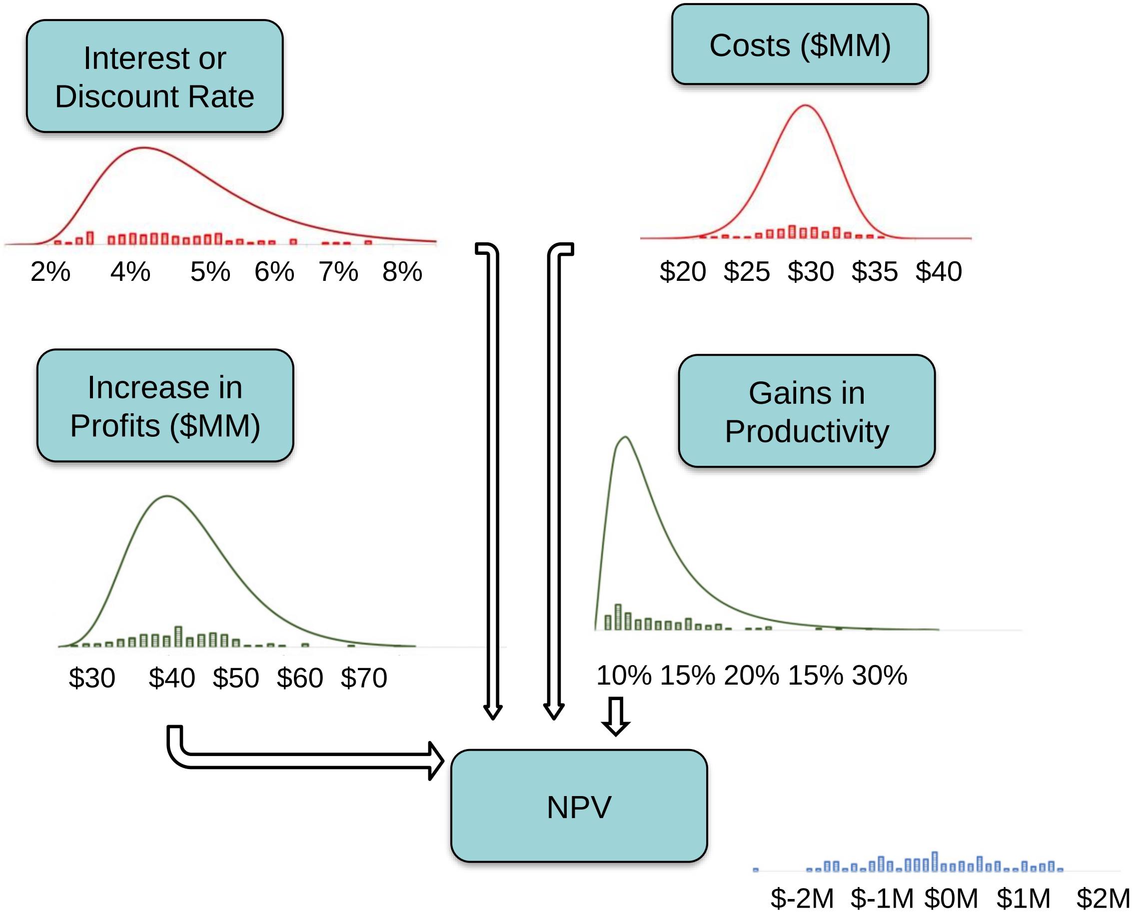

What Published Research Says
(See sources slide for details)

- Psychologists showed that simple decomposition greatly reduces estimation error for estimating the most uncertain variables.
- In the oil industry there is a correlation between the use of quantitative risk analysis methods and financial performance.
- Data at NASA from over 100 space missions showed that Monte Carlo simulations and historical data beat softer methods for estimating cost and schedule

Model The Current State of Uncertainty

Basic Distributions

Each of these examples can be found on: **www.howtomeasureanything.com**

Distributions*	Upper & Lower Bound	Best Estimate
Normal distribution 	Represents the "90% confidence interval"	Always half-way between upper and lower bound
Lognormal distribution 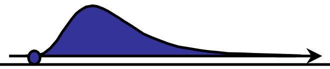	Represents the "90% confidence interval"; the absolute lower bound of a lognormal is always 0	Always a function of the upper and lower bound
Uniform distribution 	Represents the absolute (100% certain) upper and lower bounds	NA
Triangular distribution 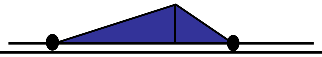	Represents the absolute (100% certain) upper and lower bounds	Represents the mode; the most likely value
Binary distribution 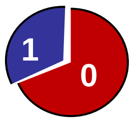	NA	Represents the % chance of the event occurring
Beta distribution 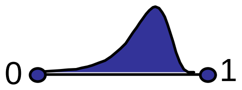	Generates a value between 0 and 1 based on “hits” and “misses”	The mode of a beta is (hits-1)/(hits+misses-2)

*A “●” means a “hard” stop, an “➔” arrow means unbounded

Model The Current State of Uncertainty

Quantifying Your Current Uncertainty

- Decades of studies show that most managers are statistically “overconfident” when assessing their own uncertainty.
- Studies also show that measuring *your own* uncertainty about a quantity is a general skill that can be taught with a ***measurable*** improvement.
- Training can “calibrate” people so that of all the times they say they are 90% confident, they will be right 90% of the time.
- HDR has calibrated over 1,000 people in the last 20 years — 85% of participants reach calibration within a half-day of training.

> “Overconfident professionals sincerely believe they have expertise, act as experts and look like experts. You will have to struggle to remind yourself that they may be in the grip of an illusion.”
>
> Daniel Kahneman, Psychologist, Economics Nobel

Model The Current State of Uncertainty

Calibration Exercise

For the initial calibration test, you have 2 types of questions:

For the questions that ask for a range, provide an upper and lower bound that you are 90% certain contains the correct answer:

- Napoleon Bonaparte was born what year? ______ **to** ______
- What is the average weight of an adult male African elephant (tons)? _____**to**_____

•For the true/false questions, select true or false and then state how confident you are in your response:

- Brazil has a larger population than Spain. **True/False:**_____ **Confidence:**____**%**
- A hockey puck will fit in a golf hole. **True/False:**_____ **Confidence:**____**%**

Reviewing Where We Are

- **Overview of Probabilistic Models and Monte Carlo Methods**
- **Computing Information Value**
- **Understanding and Setting Risk Tolerance**

A General Procedure for Measurement

Compute the Value of Additional Information

AIE quantifies and then optimizes decisions by focusing measurements where it matters most

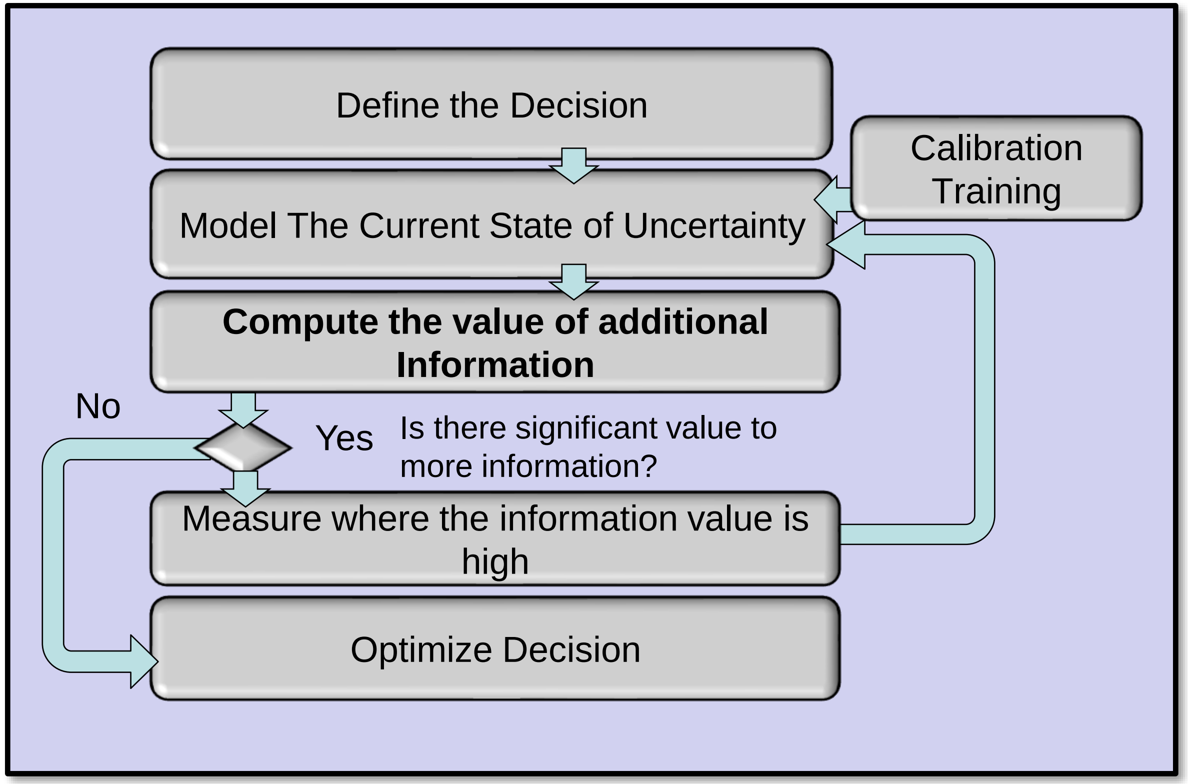

Compute the Value of Additional Information

The Value of Information

The Formula For The Value of Information:

$$EVI = \sum_{i=1}^{k} p(r_i) \max\left[\sum_{j=1}^{z} V_{1,j}\, p(\Theta_j | r_i), \sum_{j=1}^{z} V_{2,j}\, p(\Theta_j | r_i), \ldots \sum_{j=1}^{z} V_{l,j}\, p(\Theta_j | r_i), \right] - EV^*$$

OR, in its simplest form:

"The cost of being wrong times the chance of being wrong"

The formula for the value of information has been around for almost 60 years. It is widely used in many parts of industry and government as part of the "decision analysis" methods — but still mostly unheard of in the parts of business where it might do the most good.

Compute the Value of Additional Information

The EOL Method

The simplest approach computes the change in "Expected Opportunity Loss" (EOL)

Simple Binary Example:

You are about to make an investment in the unstructured big data management system.
- If the new project succeeds, you net $5 million (present value) in benefits.
- If not, you lose (net) $1 million.
- There is a 20% chance of the new system "failing."

Assuming you plan to go forward with the project, answer the following questions.

- What is the opportunity loss in this case?
- What is the EOL?
- What is the value of perfect information?

Compute the Value of Additional Information

The EOL Method Using the Power Tools

Unstructured Big Data Management

- How do we compute information value when values are continuous?

- How would we calculate the value of additional information when we can be “wrong by a little or a lot”?

- How do we compute information values when there are several uncertain variables?

Compute the Value of Additional Information

Information Value w/Ranges

- Estimate a range and distribution of hours/week spent on task
- There is a point below which investment would lose money
- The less time spent below that point, the greater the loss

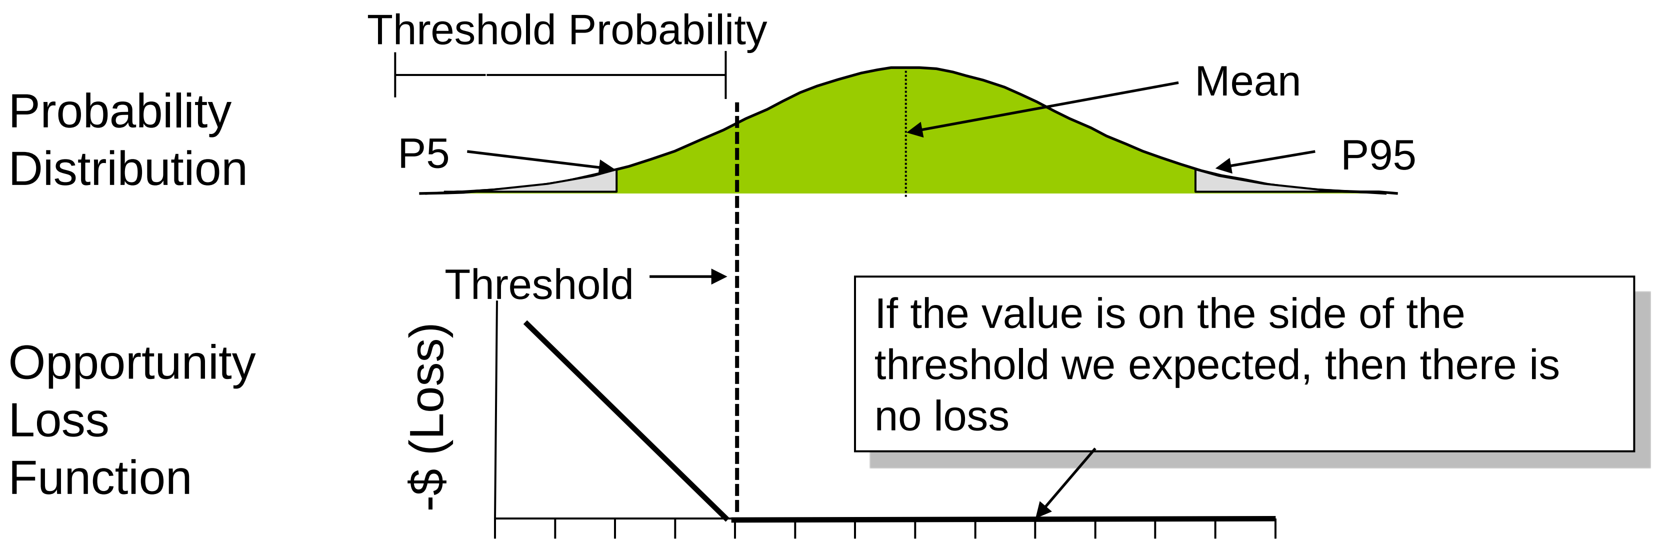

Compute the Value of Additional Information

Normal Distribution Information Value

- The “expected value” of the variable is the mean of the range of possible values.
- A threshold is a point where the value just begins to make some difference in a decision — a breakeven.
- The expected value is on one side of the threshold.
- If the true value is on the opposite side of the threshold from the mean then the best decision would have been different then one based on the mean.
- The “Threshold Probability” is the chance that this variable could have a value that would change the decision.

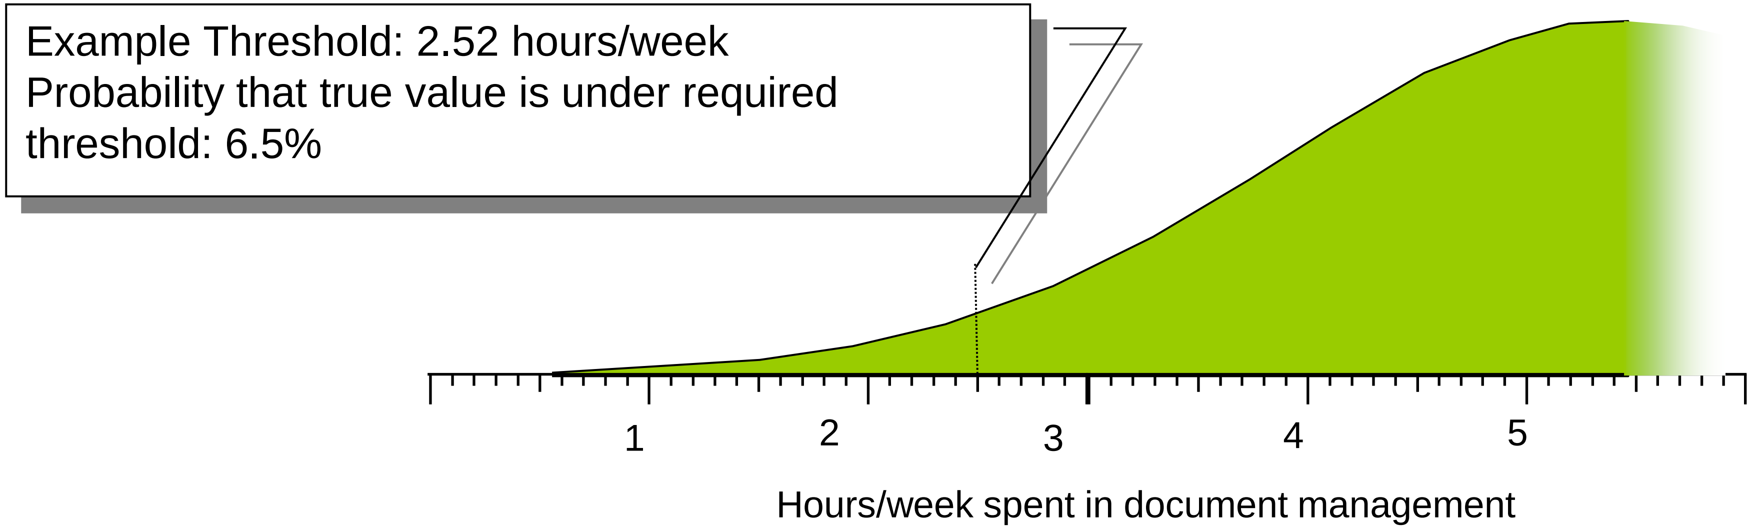

Compute the Value of Additional Information

Normal Distribution Value of Information Analysis (VIA)

- The curve on the other side of the threshold is divided up into hundreds of “slices.”
- Each slice has an assigned quantity (such as a potential productivity improvement) and a probability of occurrence.
- For each assigned quantity, there is an Opportunity Loss.
- Each slice’s Opportunity Loss is multiplied by probability to compute its Expected Opportunity Loss.

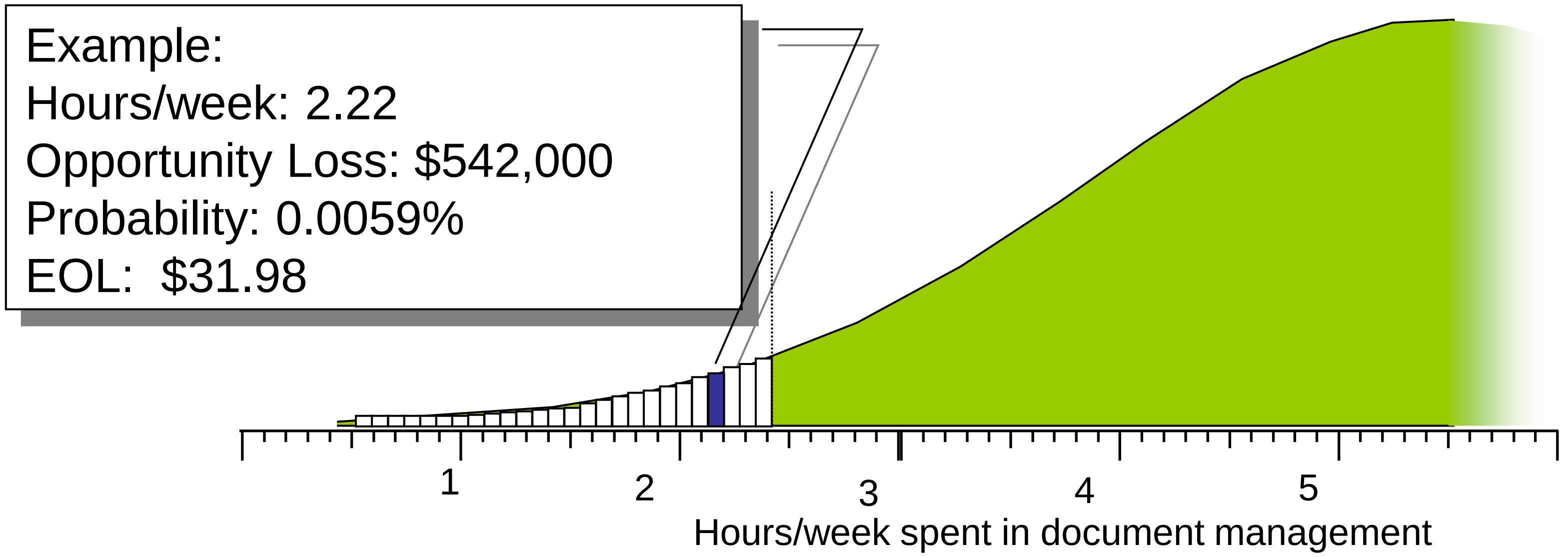

Compute the Value of Additional Information

Normal Distribution VIA (Continued)

- Total EOL for all slices equals the EOL for the variable.
- Since EOL=0 with perfect information, then the Expected Value of Perfect Information (EVPI) = sum(EOLs)
- Even though perfect information is not usually practical, this method gives us an upper bound for the information value, which can be useful by itself.
- Many of the EVPIs in a business case will be zero.
- We do this with a macro in Excel but it can also be estimated.

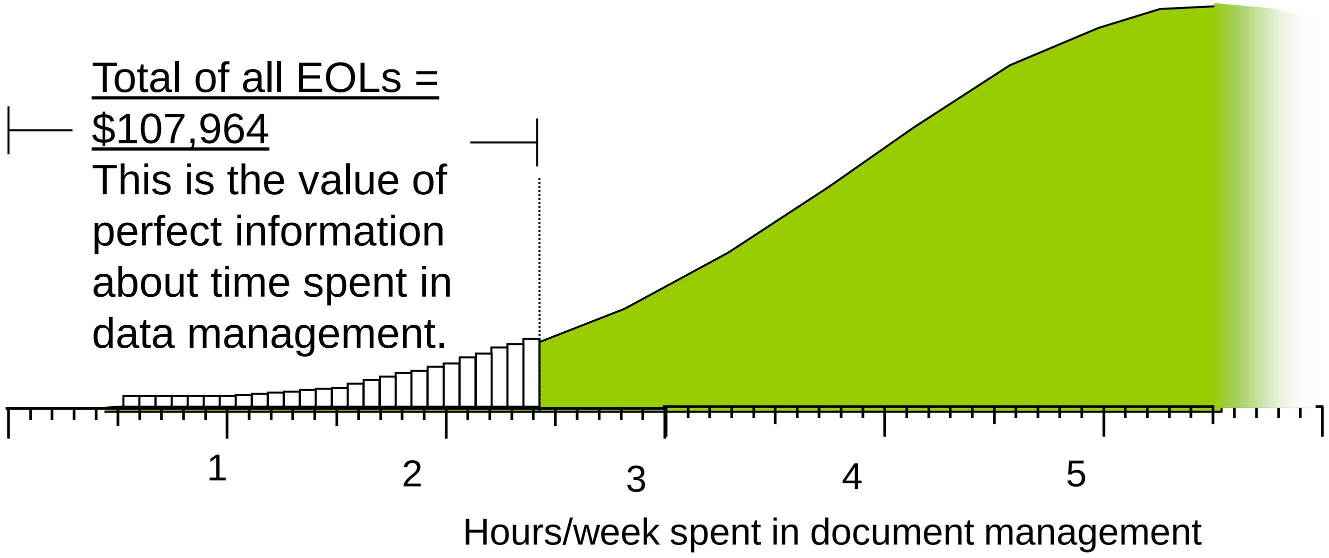

Compute the Value of Additional Information

Case A2.5 Model many uncertain variables in the UBDM

Let’s consider the case where we are uncertain about several variables in a model.

In the UBDM example, which variable would you spend the most time measuring?

What would you be willing to pay to measure that variable assuming you could eliminate uncertainty?

Compute the Value of Additional Information

The Measurement Inversion

In a business case, the economic value of measuring a variable is usually inversely proportional to the measurement attention it typically gets.

Lowest Information Value

Highest Information Value

- Initial cost
- Long-term costs
- Cost-saving benefit other than labor productivity
- Labor productivity
- Revenue enhancement
- Technology adoption rate
- Project completion

Most Measured

Least Measured

Compute the Value of Additional Information

Real Examples of Measurement Inversion

Subject	What they would have measured	What they needed to measure
New Procurement System for Government	Detailed "time and motion" study of procurement process	The price savings from using reverse auctions
Battlefield Fuel Forecasting	Chance of enemy contact, forecasts vehicle maintenance	The difference in mileage between paved and gravel roads
Risks of flooding in mining operations	Drilling test holes all over the mine	How much water the main pumps can handle
Market for new pharmaceutical products	The adoption rate of the new drug in all global regions	The duration of phase 1 testing, chance of a particular clinical outcome
Impact of pesticides regulation	The value of saving endangered species	Whether pesticides regulation ever saves any endangered species
IT security	People who attended training, external threats	Internal theft incidents

Compute the Value of Additional Information

Increasing Value & Cost of Info.

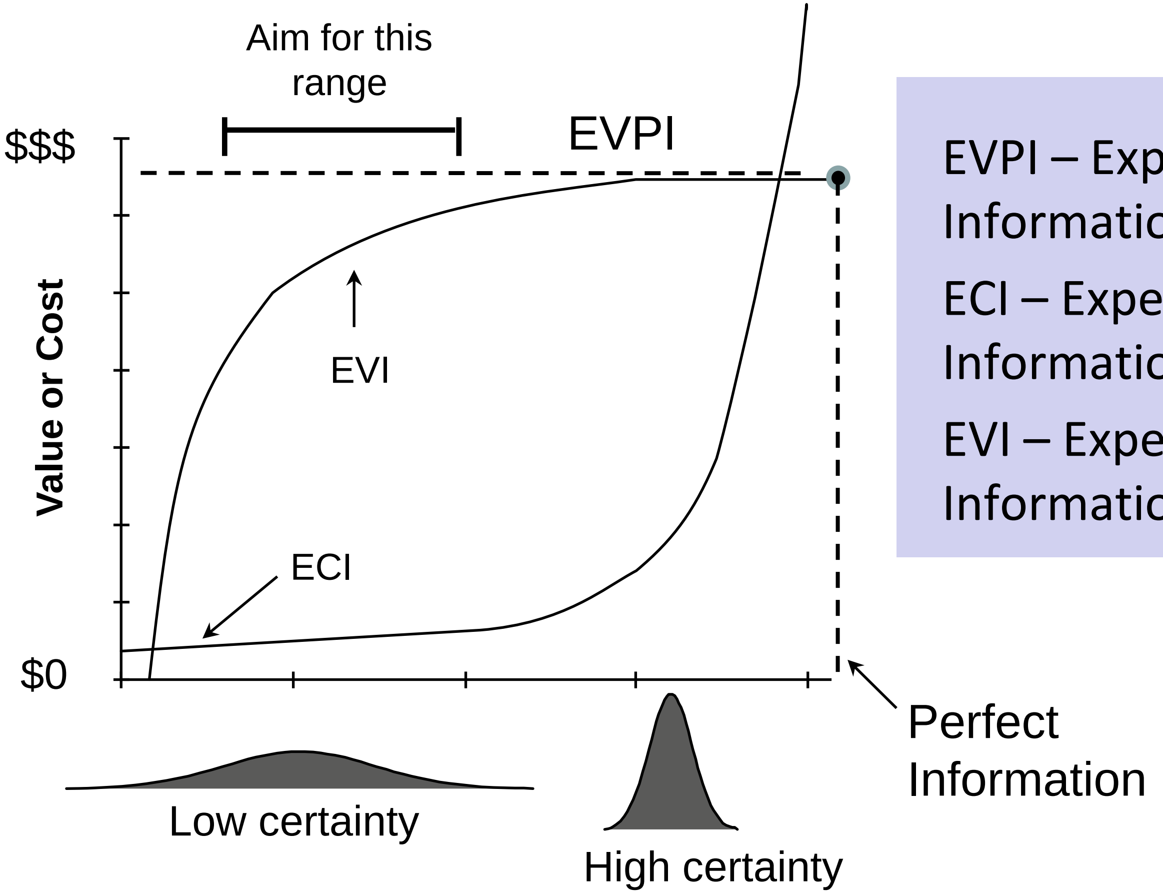

EVPI – Expected Value of Perfect Information

ECI – Expected Cost of Information

EVI – Expected Value of Information

Compute the Value of Additional Information

Next Step: Observations

Now that we know what to measure, we can think of observations that would reduce uncertainty.

The value of the information limits what methods we should use, but we have a variety of methods available.

Take the “Nike Method”: Just Do It — don’t let imagined difficulties like “exception anxiety” get in the way of starting observations.

The next module will go deeper into using observations to reduce uncertainty.

Reviewing Where We Are

- Overview of Probabilistic Models and Monte Carlo Methods
- Computing Information Value
- Understanding and Setting Risk Tolerance

A General Procedure for Measurement

Optimize the Decision

AIE quantifies and then optimizes decisions by focusing measurements where it matters most

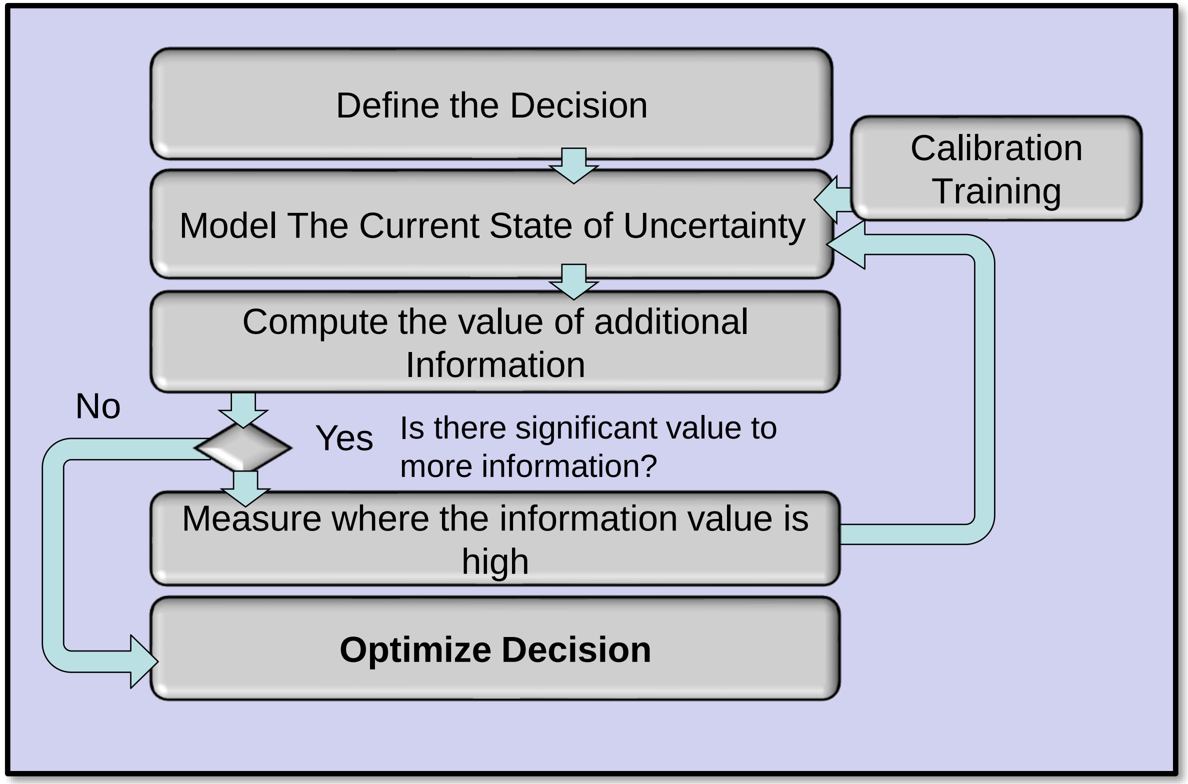

Optimize the Decision

Optimize the Decision in the Context of Uncertainty

- When the inputs to a decision model are uncertain, the output should be uncertain — this is what simulations are for.
- Is this a “good” distribution or a “bad” one? How would you know?

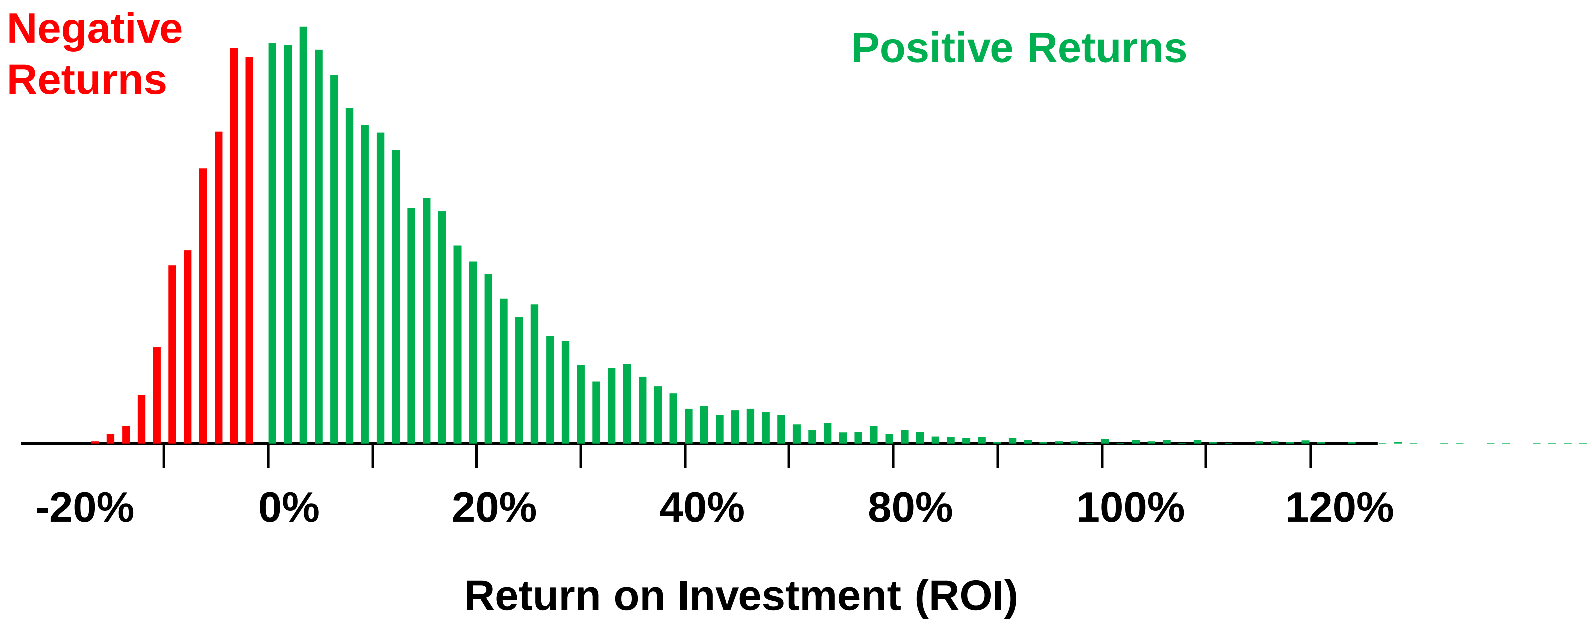

Optimize the Decision

Various Risks & Returns

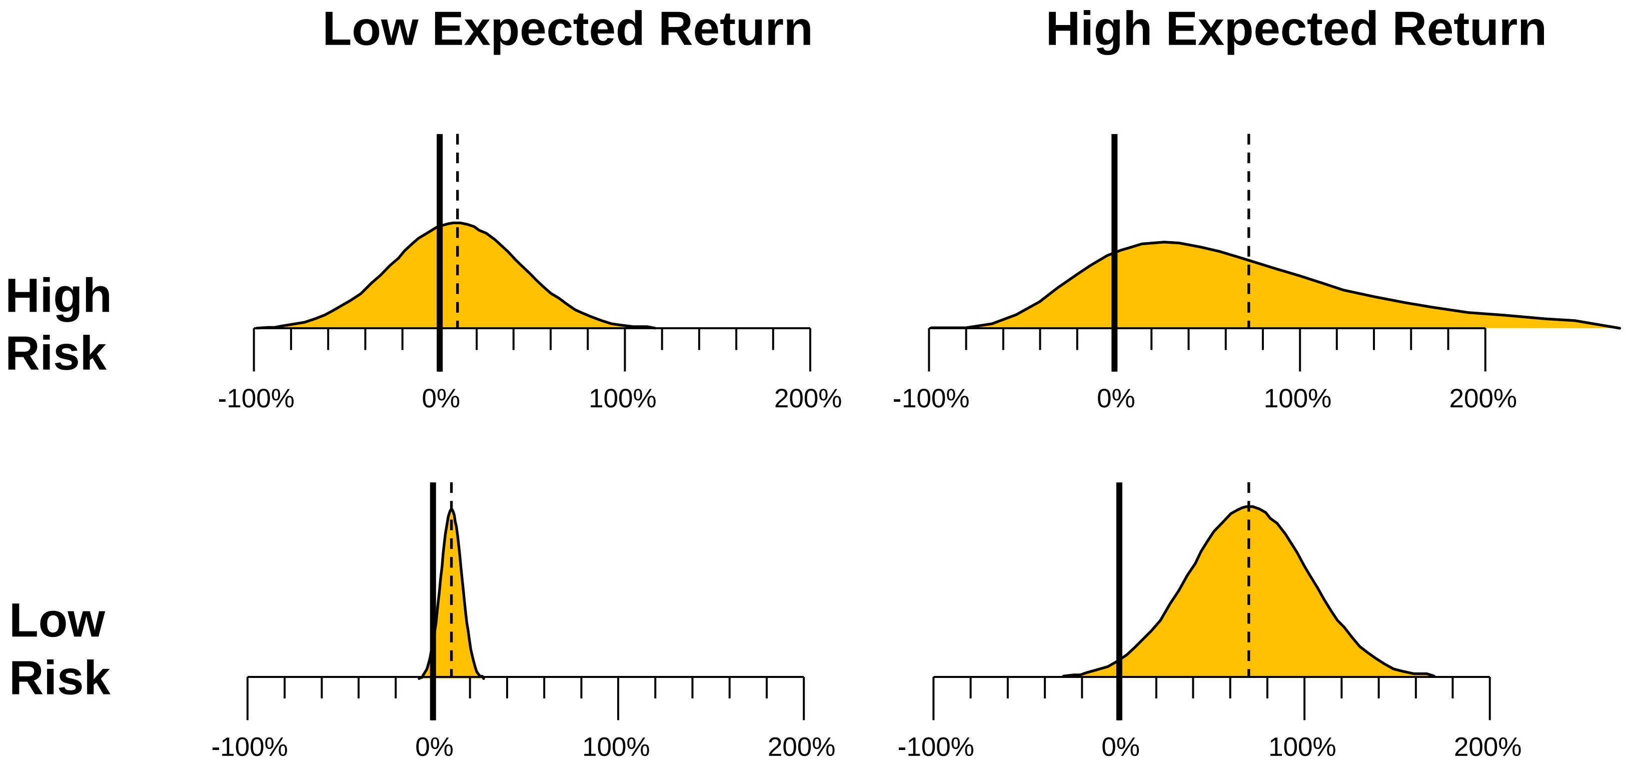

Optimize the Decision

Expert Inconsistency in Estimates & Risks

Studies have shown that estimates are affected by “anchoring” — previous exposure to unrelated numbers affects your estimates.

Studies have also shown risk aversion changes due to what should be irrelevant external factors, including:

Factor	Risk Aversion
Being around smiling people	↓
Recalling an event causing fear	↑
Recalling an event causing anger	↓
A recent win in an unrelated decision	↓
A recent loss in an unrelated decision	↑

The result is that experts cannot reproduce even their own judgments consistently — introducing a random variation to their judgment.

Optimize the Decision

A Form of Risk Tolerance

"Expected Utility Theory" is a key concept for assessing risk/return tradeoffs.

It has been shown mathematically that if a person follows certain fairly obvious rules, their aversion to risk can be quantified in a useful way.

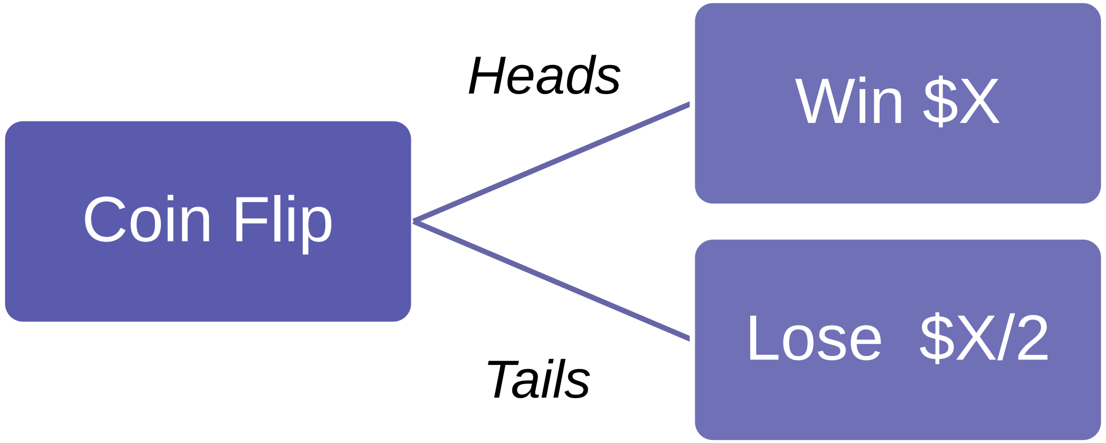

How large would you make X?

Optimize the Decision

The Utility of Outcomes

Most people are "risk averse" to some degree.

They would not consider a 50% chance of winning $10,000 to be exactly equal to a 100% chance of winning $5,000.

This can be modeled as a decreasing "utility" of wins and an increasing "negative utility" of losses.

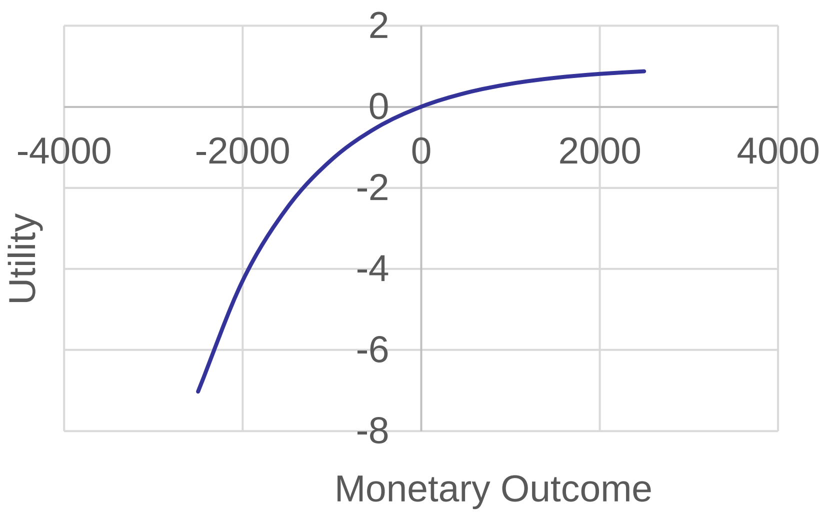

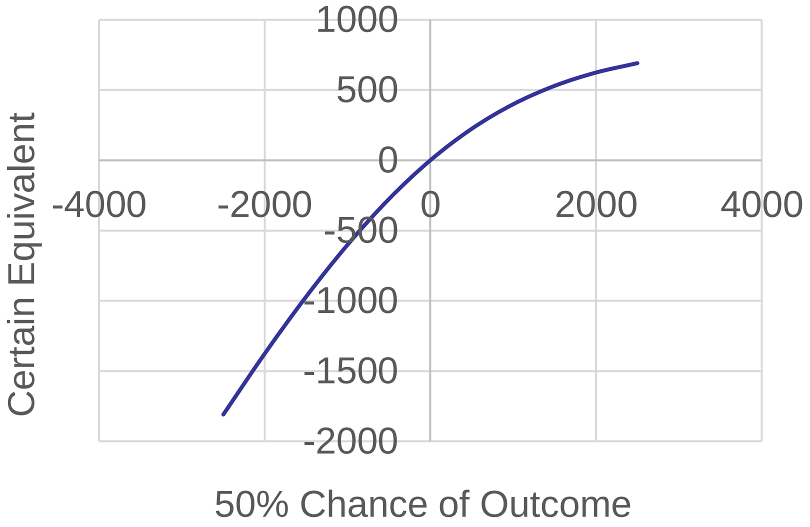

Optimize the Decision

Quantifying Risk Aversion

- The simplest element of Harry Markowitz's Nobel Prize-winning method "Modern Portfolio Theory" is documenting how much risk an investor accepts for a given return.
- The "Investment Boundary" states how much risk an investor is willing to accept for a given return.
- For our purposes, we modified Markowitz's approach a bit.

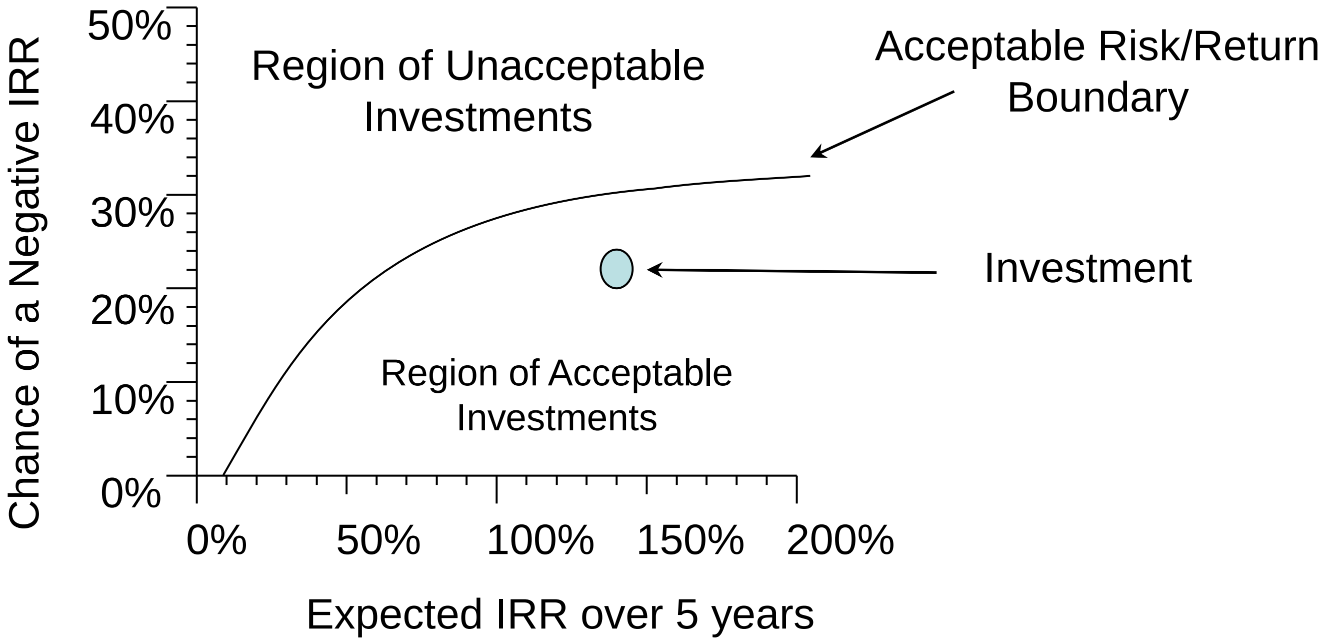

Optimize the Decision

Example of Risk Effects

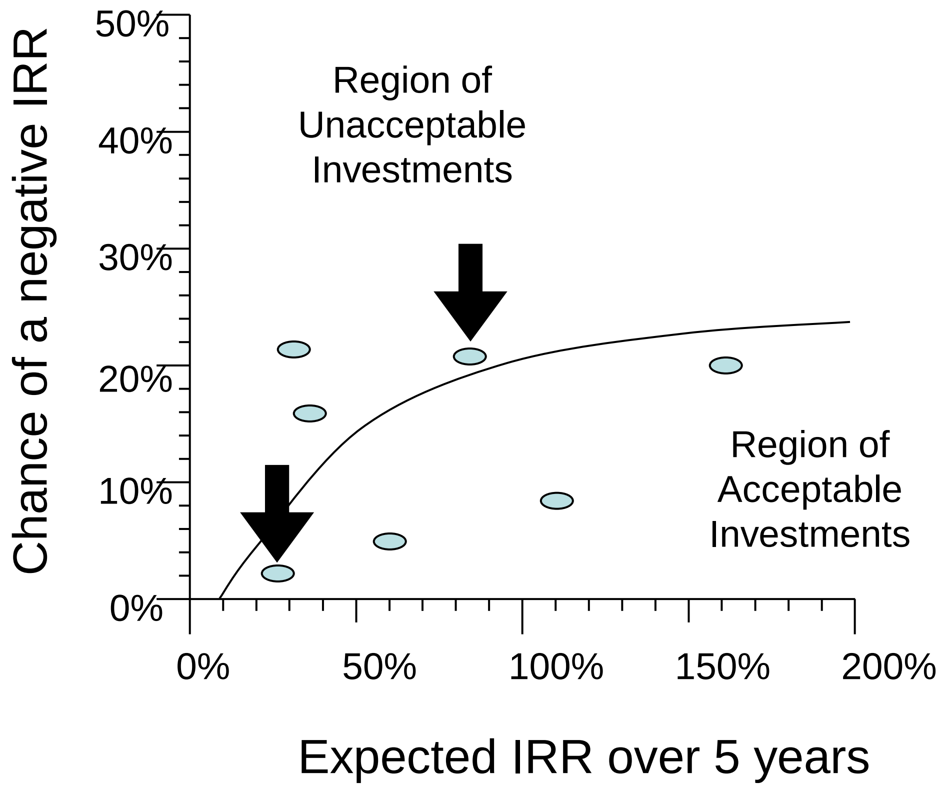

- These are real IT investments of $2M-$3M plotted against a client's investment boundary.
- The 27% ROI investment is actually preferred to the 83% ROI investment.

Summary

How to Make Decisions Under Uncertainty Review

1. Probabilistic models allow us to realistically account for our uncertainty.
2. Information values let us focus our time and money on measuring the most important variables
3. Quantifying risk aversion can help us standardize our investment decisions.

Summary

Review Question

1. What is the 5-step, generic process for any approach to measurement?
2. What method do we use to generate the chance of various outcomes based on uncertain inputs?
3. What has a high information value? (Select all that apply)
    - Any quantity about project costs
    - Any variable that “straddles” the threshold for a big, risky decision
    - Most uncertain quantities in most big decisions
    - A variable that is uncertain and makes a big difference in a big bet
4. Select all the true statements:
    - EVPI is the minimum amount you should be willing to spend on any measurement.
    - You should measure to the point where EVI is maximized.
    - Small, iterative measurements often have the highest net information value.
    - The value of information rises slowly at first and then increases rapidly.
5. Why do we need to quantify and document risk aversion?

Questions?

Contact:

Doug Hubbard

Hubbard Decision Research

dwhubbard@hubbardresearch.com

www.hubbardresearch.com

630 858 2788

Supplementary Material

Hubbard Decision Research
2 South 410 Canterbury Ct
Glen Ellyn, Illinois 60137
www.hubbardresearch.com

“Misunderestimating”

Experts and even some popular quantitative models consistently underestimate the chance of catastrophic outcomes.

Mathematical models often use simplified probability distributions that that don’t even match history — more realistic methods exist (such as “power law distributions”).

In addition to overconfidence, people tend to become more risk tolerant as time passes without a catastrophe.

People even apparently misinterpret near misses as evidence that disasters are less common.

Power Law Distributions look like straight lines on log/log charts

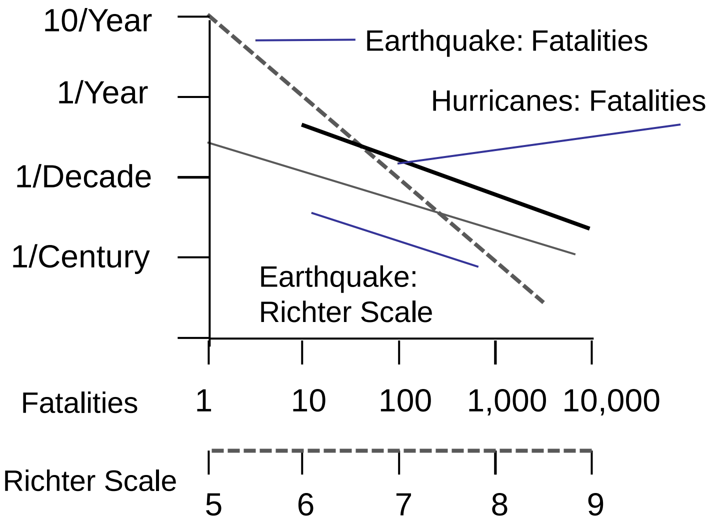

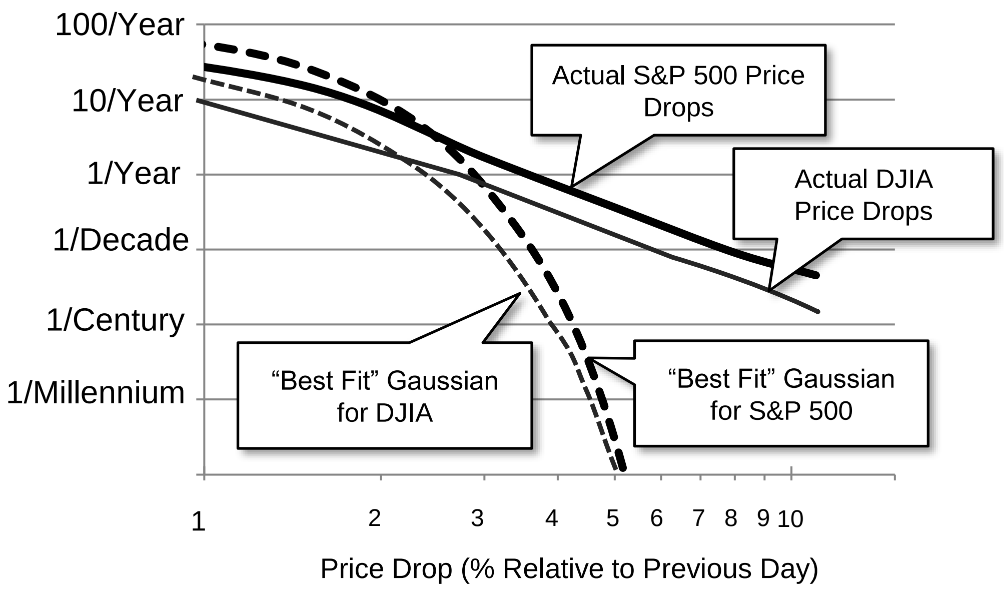

Misunderestimating…by a Lot

August 2007, the Chief Financial Officer of Goldman Sachs, David Viniar, was quoted in the Financial Times saying:

“We are seeing things that were 25-standard deviation moves, several days in a row”

A 25-standard deviation event is a number with a probability so tiny it is smaller than…

- …one divided by the national debt…
- …of every country in the world…
- …measured in yen…
- …divided again by the number of individual bacteria on Earth…
- …divided again by all the atoms in the observable universe.
- In fact, this number would be TOO big…