Introduction

The Problem of the Decision Analyst

This stylised scenario embodies the core problems of decision analysis:

- You have a client[1].
- The client must choose one action from a set of possibilities.
- This client is uncertain about many things, including:
  - Her priorities.
    Conflicting requirements can be difficult to resolve.
  - What might happen.
    Fundamental uncertainty – things not within her control.
  - How other people may act.
    Other interested parties might influence the outcome.
- You must advise this client on the best course of action.

[1] This may be yourself, but it is useful to separate the two rôles.

A problem of two parts

- Elicitation: Obtain precise answers to several questions:
  - What is the client's problem?
  - what does she believe?
  - What does she want?
- Calculation: Given this information
  - What are its logical implications?
  - What should our client do?

Elicitation $\longrightarrow$ Calculation $\longrightarrow$ Elicitation $\longrightarrow$ Calculation $\longrightarrow \ldots$

What does she really want?

Example (Advising a university undergraduate)

What is their objective?

- Getting the best possible degree?
- Trying to get a particular job after university?
- Learning for its own sake?
- Having as much fun as possible?
- A combination of the above?

Example (A small business owner)

What is their objective?

- Staying in business?
- Making £X of profit in as short a time as possible?
- Making as much profit as possible in time T?
- Eliminating competition?
- Maximising growth?

What does she know?

As well as knowing what our client *wants* we need to know what they *know*:

- What are their options?
- What are the possible consequences of these actions?
- How are the consequences related to the action taken?
- Are any other parties involved? If so, what are their objectives?

Example (Marketing)

- How can we advertise?
- What are the *costs* of different approaches?
- What are the *effects* of these approaches?
- What volume of production is possible?
- What competition do we have?

Example (Insurance)

Insurance against a particular type of loss...

- Probability of the loss occurring is $p \ll 1$.
- Cost of that lost would be, say, £5,000.
- Insurance premium is £10.

Why are both parties happy with this?

Example (A Simple Lottery)

- $\mathbb{P}(\{\text{Win}\}) = 1/10,000$
- $Value(\text{Win}) = £5,000$
- Ticket price £1.

Why is this acceptable? What about simple variations?

Is that *really* what she believes?

It is important to distinguish between that which is *believed* from that which is *hoped*, *feared* or simply asserted.

Example (Economic forecasting)

Recent forecasts of British GDP growth in 2009:

- -0.1% – International Monetary Fund
- -0.75– -1.25% British Government
- -1.1% Organisation for Economic Co-operation and...
- -1.7% Confederation of British Industry
- -2.9% Centre for Economics and Business research

Each organisation has different objectives & knowledge.
Are they necessarily reliable indications of the underlying beliefs of these organisations[2]?

[2]We will put aside the philosophical questions raised by this concept...

Quantification of Subjective Knowledge

Our client has beliefs and some idea about her objective. She probably isn't a mathematician. We have to codify things in a rigorous mathematical framework.
In particular, we must be able to encode:

- Beliefs about what can happen and how likely those things are to happen.
- The cost or reward of particular outcomes.
- In the case of games: What any other interest parties want and how they are likely to react.

Having done this, we must use our mathematical skills to work out how to advise our client.

Some Terminology

Before considering details, we should make sure we agree about terminology.

- In a *decision problem* we have:
  - A (random) source of uncertainty.
  - A collection of possible *actions*.
  - A collection of *outcomes*.

  and we wish to choose the action to obtain a favourable outcome.
- A *game* is a similar problem in which the uncertainty arises from the behaviour of a (rational) opponent.

From Questions to Answers

Now we need to answer some questions:

1. How can be elicit and quantify beliefs?
2. How can we represent their particular problem mathematically?
3. How do we represent her objectives quantitatively?
4. What should we advise our client to do?
5. What can we do if other rational agents are involved?

We will begin by answering question 1: we can use probability.

Probability

Foundations of An Axiomatic Theory of Probability

The *Russian school* of probability is based on axioms.
The abstract specification of probability requires three things:

1. A set of all possible outcomes, Ω.

   The *sample space* containing elementary events.

2. A collection of subsets of Ω, $\mathcal{F}$.

   Outcomes of interest.

3. A function which assigns a probability to our events: $\mathbb{P} : \mathcal{F} \rightarrow [0, 1]$

   The probability itself.

Example (Simple Coin-Tossing)

- All possible outcomes might be:

$$\Omega = \{H, T\}.$$

- And we might be interested in all possible subsets of these outcomes:

$$\mathcal{F} = \{\emptyset, \{H\}, \{T\}, \Omega\}.$$

- In which case, under reasonable assumptions:

$$\mathbb{P}(\emptyset) = 0 \qquad \mathbb{P}(\{H\}) = \frac{1}{2}$$
$$\mathbb{P}(\{T\}) = \frac{1}{2} \qquad \mathbb{P}(\{H, T\}) = 1$$

Example (A Tetrahedral (4-faced) Die)

- The possible outcomes are: $\Omega = \{1, 2, 3, 4\}$
- And we might again consider all possible subsets:

$$\begin{aligned}\mathcal{F} = \{ \quad & \emptyset, & \{1\}, && \{2\}, && \{3\}, \\ & \{4\}, & \{1,2\}, && \{1,3\}, && \{1,4\}, \\ & \{2,3\}, & \{2,4\}, && \{3,4\}, && \{1,2,3\}, \\ & \{1,2,4\}, & \{1,3,4\}, && \{2,3,4\}, && \{1,2,3,4\}\}\end{aligned}$$

- In this case, we might think that, for any $A \in \mathcal{F}$:

$$\mathbb{P}(A) = |A|/|\Omega| = \frac{\text{Number of values in } A}{4}$$

Example (The National Lottery)

- $\Omega = \{\text{All unordered sets of 6 numbers from}\{1, \ldots, 49\}\}$
- $\mathcal{F} = \text{All subsets of } \Omega$
- Again, we can construct $\mathbb{P}$ from expected uniformity.

- But there are $\binom{49}{6} = 13983816$ elements of Ω and consequently $2^{13983816} \approx 6 \times 10^{6000000}$ subsets!
- Even this simple discrete problem has produced an object of incomprehensible vastness.
- What would we do if $\Omega = \mathbb{R}$?
- It's often easier not to work with *all* of the subsets of Ω.

Algebras of Sets

Given Ω, $\mathcal{F}$ must satisfy certain conditions.

1. $\Omega \in \mathcal{F}$

   The event "something happening" is in our set.

2. If $A \in \mathcal{F}$, then

$$\Omega \setminus A = \{x \in \Omega : x \notin A\} \in \mathcal{F}$$

   If A happening is in our set then A not happening is too.

3. If $A, B \in \mathcal{F}$ then

$$A \cup B \in \mathcal{F}$$

   If event A and event B are both in our set then an event corresponding to either A or B happening is too.

A set that satisfies these conditions is called an *algebra* (over Ω).

σ-Algebras of Sets

If, in addition to meeting the conditions to be an algebra, $\mathcal{F}$ is such that:

- If $A_1, A_2, \cdots \in \mathcal{F}$ then $\bigcup_{i=1}^{\infty} A_i \in \mathcal{F}$

  If any countable sequence of events is in our set, then the event corresponding to any one of those events happening is too.

then $\mathcal{F}$ is known as a σ-algebra.

Example (Selling a house)

- You wish to sell a house, for at least £250,000.
- On Monday you receive an offer of X.
- You must accept or decline this offer immediately.
- On Tuesday you will receive an offer of Y.
- What should you do?

- $\Omega = \{(x, y) : x, y \geq £100,000\}$
- But, we only care about events of the form:

$$\{(i, j) : i < j\} \text{ and } \{(i, j) : i > j\}$$

- Including some others ensures that we have an algebra:

$$\{(i, j) : i = j\} \quad \{(i, j) : i \neq j\} \quad \{(i, j) : i \leq j\} \quad \{(i, j) : i \geq j\} \quad \emptyset \quad \Omega$$

Atoms

Some events are *indivisible* and somehow fundamental: An event $E \in \mathcal{F}$ is said to be an atom of $\mathcal{F}$ if:

1. $E \neq \emptyset$
2. $\forall A \in \mathcal{F}$:

$$E \cap A = \begin{cases} \emptyset \\ \text{or } E \end{cases}$$

Any element of $\mathcal{F}$ contains all of E or none of E.

If $\mathcal{F}$ is finite then any $A \in \mathcal{F}$, we can write:

$$A = \bigcup_{i=1}^{n} E_i$$

for some finite number, n, and atoms E_i of $\mathcal{F}$.

We can represent any event as a combination of atoms.

Example (Selling a house. . .)

Here, our algebra contained:

$$\begin{array}{cccc} \{(i,j) : i < j\} & \{(i,j) : i > j\} & \{(i,j) : i \neq j\} & \emptyset \\ \{(i,j) : i \leq j\} & \{(i,j) : i \geq j\} & \{(i,j) : i = j\} & \Omega \end{array}$$

Which of these sets are atoms?

- $\{(i,j) : i < j\}$ **is**
- $\{(i,j) : i > j\}$ **is**
- $\{(i,j) : i \neq j\}$ **is not** – it's the union of two atoms
- $\emptyset$ **is not** $\emptyset$ is never an atom
- $\{(i,j) : i = j\}$ **is**
- $\{(i,j) : i \leq j\}$ **is not** – it's the union of two atoms
- $\{(i,j) : i \geq j\}$ **is not** – it's the union of two atoms
- Ω **is not** – it's the union of three atoms

The Axioms of Probability – Finite Spaces

$\mathbb{P} : \mathcal{F} \to \mathbb{R}$ is a probability measure over $(\Omega, \mathcal{F})$ iff:

1. For any $A \in \mathcal{F}$:
$$\mathbb{P}(A) \geq 0$$
All probabilities are positive.
2. $$\mathbb{P}(\Omega) = 1$$
Something certainly happens.
3. For any[3] $A, B \in \mathcal{F}$ such that $A \cap B = \emptyset$:
$$\mathbb{P}(A \cup B) = \mathbb{P}(A) + \mathbb{P}(B)$$
Probabilities are (sub)additive.

[3] This is sufficient if Ω is finite; we need a slightly stronger property in general.

The Axioms of Probability – General Spaces [see ST213]

$\mathbb{P} : \mathcal{F} \to \mathbb{R}$ is a probability measure over $(\Omega, \mathcal{F})$ iff:

1. For any $A \in \mathcal{F}$:

$$\mathbb{P}(A) \geq 0$$

All probabilities are positive.

2.

$$\mathbb{P}(\Omega) = 1$$

Something certainly happens.

3. For any $A_1, A_2, \dots \in \mathcal{F}$ such that $\forall i \neq j : A_i \cap A_j = \emptyset$:

$$\mathbb{P}\left(\bigcup_{i=1}^{\infty} A_i\right) = \sum_{i=1}^{\infty} \mathbb{P}(A_i).$$

Probabilities are countably (sub)additive.

Measures and Masses

- A *measure* tells us "how big" a set is [see MA359/ST213].
- A *probability measure* tells us "how big" an event is in terms of the likelihood that it happens [see ST213/ST318].
- In discrete spaces probability mass functions are often used.

Definition (Probability Mass Function)

If $\mathcal{F}$ is an algebra containing finitely many atoms $E_1, \ldots, E_n$. A *probability mass function*, f, is a function defined for every atom as $f(E_i) = p_i$ with:

- $p_i \in [0, 1]$
- and $\sum_{i=1}^{n} p_i = 1$.

Masses to Measures

- Let $S = \{A_1, \ldots, A_n\}$ be such that:
  - $\forall i \neq j : A_i \cap A_j = \emptyset$

    The elements of S are disjoint.
  - $\cup_{i=1}^{n} A_i = \Omega$

    S *covers* Ω.
- We can construct a finite algebra, $\mathcal{F}$ which contains the 2^n sets obtained as finite unions of elements of S.

  This algebra is *generated* by S.
- The atoms of the generated algebra are the elements of S.
- A mass function f on the elements of S defines a probability measure on $(\Omega, \mathcal{F})$:

$$\mathbb{P}(B) = \sum f(A_i)$$

(the sum runs over those atoms A_i which are contained in B).

So what?

So far we've seen:

- A mathematical framework for dealing with probabilities.
- A way to construct probability measures from the probabilities of every elementary event in a discrete problem.
- A way to construct probability measures from the probability mass function of a complete set of atoms.

But this doesn't tell us:

- What probabilities really mean.
- How to assign probabilities to *real* events...dice aren't everything!
- Why we should use probability to make decisions.

Geometry, Symmetry and Probability

- If probabilities have a geometric interpretation, we can often deduce probabilities from symmetries.

Example (Coin Tossing Again)

- Here, $\Omega = \{H, T\}$ and $\mathcal{F} = \{\emptyset, \{H\}, \{T\}, \{H, T\}\}$
- Axiomatically: $\mathbb{P}(\Omega) = P(\{H, T\}) = 1$.
- The atoms are $\{H\}$ and $\{T\}$.
- Symmetry arguments suggest that $\mathbb{P}(\{H\}) = \mathbb{P}(\{T\})$.

  Implicitly, we are assuming that the symbol on the face of a coin does not influence its final orientation.
- Axiomatically: $\mathbb{P}(\{H, T\}) = \mathbb{P}(\{H\}) + \mathbb{P}(\{T\})$.
- Therefore: $\mathbb{P}(\{H\}) = \mathbb{P}(\{T\}) = 1/2$.

Example (Tetrahedral Dice Again)

- Here, $\Omega = \{1, 2, 3, 4\}$ and $\mathcal{F}$ is the set of all subsets of Ω.
- The atoms in this case are $\{1\}, \{2\}, \{3\}$ and $\{4\}$.
- Physical symmetry suggests that:

$$\mathbb{P}(\{1\}) = \mathbb{P}(\{2\}) = \mathbb{P}(\{3\}) = \mathbb{P}(\{4\})$$

- Axiomatically, $1 = \mathbb{P}(\{1, 2, 3, 4\}) = \sum_{i=1}^{4} \mathbb{P}(\{i\}) = 4\mathbb{P}(\{1\})$.
- And we again end up with the expected result $\mathbb{P}(\{i\}) = 1/4$ for all $i \in \Omega$.

Example (Lotteries Again)

- $\Omega = \{\text{All unordered sets of 6 numbers from}\{1, \ldots, 49\}\}$
- $\mathcal{F} = \text{All subsets of } \Omega$
- Atoms are once again the sets containing a single element of Ω.

This is usual when $|\Omega| < \infty$...

- As $|\Omega| = 13983816$, we have that many atoms.
- Each atom corresponds to drawing one unique subset of 6 balls.
- We might assume that each subset has equal probability... in which case:

$$\mathbb{P}(\{< i_1, i_2, i_3, i_4, i_5, i_6 >\}) = 1/13983816$$

for any valid set of numbers $< i_1, \ldots, i_6 >$.

Complete Spatial Randomness and π

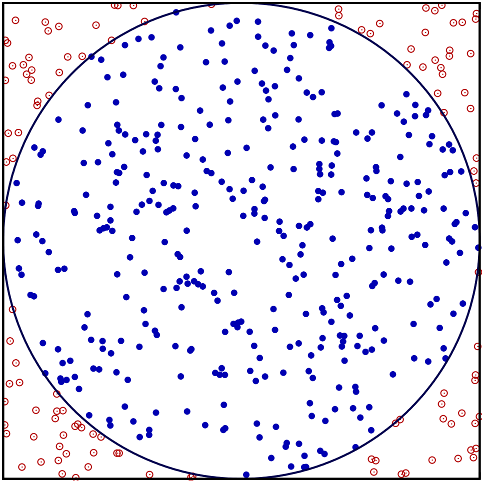

- Let (X, Y) be uniform over the centred unit square.
- Define

$$E = \left\{ (x, y) : x^2 + y^2 \leq \frac{1}{4} \right\}$$

- Now

$$\begin{aligned} \mathbb{P}((X, Y) \in E) =& A_{\text{circle}} / A_{\text{square}} \\ =& \pi \times (1/2)^2 / 1^2 \\ =& \pi/4 \end{aligned}$$

Balls in Urns

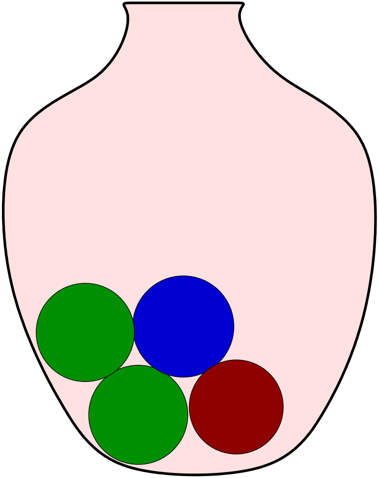

- Let $\mathcal{I}$ be (discrete) a set of colours.
- An urn contains n_i balls of colour i.
- The probability that a drawn ball has colour i is:

$$\frac{n_i}{\sum_{j \in \mathcal{I}} n_j}$$

We assume that the colour of the ball does not influence its probability of selection.

Spinners

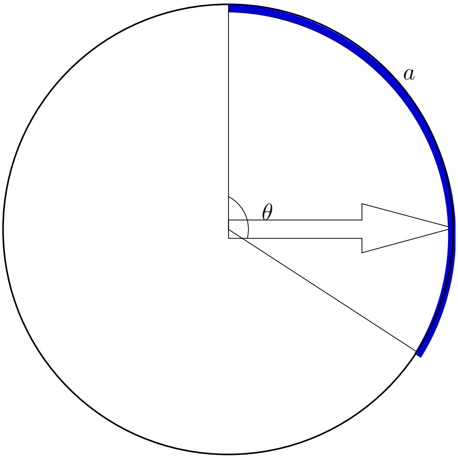

- $\mathbb{P}[\text{Stops in purple}] = a$
- Really a statement about physics.
- What do we mean by probability?