So You Want to be a Millionaire?

Christopher Jennison

Department of Mathematical Sciences,

University of Bath, UK

http://people.bath.ac.uk/mascj

BUMS Pub Lecture

Bath

November 2024

Plan for this talk

1. A game of Who Wants to be a Millionaire?

   Selecting the contestant

   Playing the game

   How should you play the game?

2. Optimisation by Dynamic Programming

3. Relation to Sequential Clinical Trials

1. Selecting the contestant

The Countdown Numbers Game

Once you have solved the puzzle:

Shout "Got it" and raise your hand.

You will have to come down to the board to show your working.

We shall wait for 4 people to shout "Got it", so there are some reserves.

Over to Rachel ...

The Millionaire Game

When you see the question, you must decide whether to answer the question or take the current prize.

Prizes

£100	£2,000	£64,000
£200	£4,000	£125,000
£300	£8,000	£250,000
£500	£16,000	£500,000
£1,000	**£32,000**	**£1 million**

Answer correctly: Proceed to the next question

Answer wrongly: Leave with £0, £1,000 or £32,000.

Lifelines

50/50 Ask the audience Phone a friend

The Millionaire Game

Question 1 *for £100*

What is the square root of 169 ?

A 11

B 22

C 13

D π^2

The Millionaire Game

Question 1 *for £100*

What is the square root of 169 ?

C 13

D π^2

The Millionaire Game

Question 1 *for £100*

What is the square root of 169 ?

C 13

The Millionaire Game

Question 2 *for £200*

The ancient Greeks had a method for finding prime numbers known as:

A The sieve of Eratosthenes

B The whisk of Aristotle

C The blender of Pythagoras

D Zenos microwave

The Millionaire Game

Question 2 *for £200*

The ancient Greeks had a method for finding prime numbers known as:

A The sieve of Eratosthenes

C The blender of Pythagoras

The Millionaire Game

Question 2 *for £200*

The ancient Greeks had a method for finding prime numbers known as:

A The sieve of Eratosthenes

The Millionaire Game

Question 3 *for £300*

Baths Centre for Doctoral Training in Statistical Applied Maths is known as:

A TANGo

B RUMBa

C SAMBa

D FANdango

The Millionaire Game

Question 3 *for £300*

Baths Centre for Doctoral Training in Statistical Applied Maths is known as:

A TANGo

C SAMBa

The Millionaire Game

Question 3 *for £300*

Baths Centre for Doctoral Training in Statistical Applied Maths is known as:

C SAMBa

The Millionaire Game

Question 4 *for £500*

There are n people in a room. In order for there to be a probability of a half that there are two people who share the same birthday, n must be at least:

A 365

B 183

C 23

D 5

Question 4 *for £500*

There are n people in a room. In order for there to be a probability of a half that there are two people who share the same birthday, n must be at least:

B 183

C 23

The Millionaire Game

Question 4 *for £500*

There are n people in a room. In order for there to be a probability of a half that there are two people who share the same birthday, n must be at least:

C 23

The Millionaire Game

Question 5 *for £1,000*

The best social event of 2025 will be:

A My birthday party

B Your birthday party

C The Kings garden party

D The BUMS Ball

The Millionaire Game

Question 5 *for £1,000*

The best social event of 2025 will be:

A My birthday party

D The BUMS Ball

The Millionaire Game

Question 5 *for £1,000*

The best social event of 2025 will be:

D The BUMS Ball

The Millionaire Game

Question 6 *for £2,000*

The set of positive prime numbers is:

A Empty

B Finite

C Countably infinite

D Uncountably infinite

The Millionaire Game

Question 6 *for £2,000*

The set of positive prime numbers is:

C Countably infinite

D Uncountably infinite

The Millionaire Game

Question 6 *for £2,000*

The set of positive prime numbers is:

C Countably infinite

The Millionaire Game

Question 7 *for £4,000*

According to Alan Turings theory of morphogenesis, which of the following can an animal NOT have?

A A striped body and a striped tail

B A striped body and a spotty tail

C A spotty body and a striped tail

D A spotty body and a spotty tail

The Millionaire Game

Question 7 *for £4,000*

According to Alan Turings theory of morphogenesis, which of the following can an animal NOT have?

B A striped body and a spotty tail

C A spotty body and a striped tail

The Millionaire Game

Question 7 *for £4,000*

According to Alan Turings theory of morphogenesis, which of the following can an animal NOT have?

B A striped body and a spotty tail

The Millionaire Game

Question 8 *for £8,000*

The number 1729 is interesting because it is:

A The number of students in the Maths Department at Bath

B The number of pints of beer that Bath Maths students drank last week

C The number of minutes in a day and a quarter

D The smallest number that can be expressed as the sum of two cubes in two different ways

The Millionaire Game

Question 8 *for £8,000*

The number 1729 is interesting because it is:

B The number of pints of beer that Bath Maths students drank last week

D The smallest number that can be expressed as the sum of two cubes in two different ways

The Millionaire Game

Question 8 *for £8,000*

The number 1729 is interesting because it is:

D **The smallest number that can be expressed as the sum of two cubes in two different ways**

The Millionaire Game

Question 9 *for £16,000*

What is $e^{i\pi}$?

A $i + 1$

B ∞

C As small as a number can be without actually being zero

D -1

The Millionaire Game

Question 9 *for £16,000*

What is $e^{i\pi}$?

A $i+1$

D -1

The Millionaire Game

Question 9 *for £16,000*

What is $e^{i\pi}$?

D -1

The Millionaire Game

Question 10 *for £32,000*

Consider the statement

> *If x is the sum of two squares (of integers), then $2x$ is also the sum of two squares (of integers)*

Is this statement:

A Always true

B Sometimes (but not always) true

C True, but only on Mondays

D It may always be true but it is not possible to decide whether that is definitely the case

The Millionaire Game

Question 10 *for £32,000*

Consider the statement

> *If x is the sum of two squares (of integers), then $2x$ is also the sum of two squares (of integers)*

Is this statement:

A Always true

D It may always be true but it is not possible to decide whether that is definitely the case

The Millionaire Game

Question 10 *for £32,000*

Consider the statement

> *If x is the sum of two squares (of integers), then $2x$ is also the sum of two squares (of integers)*

Is this statement:

A Always true

The Millionaire Game

Question 11 *for £64,000*

What is the limit of $(1 - 1/n)^{2n}$ as $n \to \infty$

A 1

B π

C $1/e^2$

D $\sin(\pi/3)$

The Millionaire Game

Question 11 *for £64,000*

What is the limit of $(1 - 1/n)^{2n}$ as $n \to \infty$

A 1

C $1/e^2$

The Millionaire Game

Question 11 *for £64,000*

What is the limit of $(1 - 1/n)^{2n}$ as $n \to \infty$

C $1/e^2$

The Millionaire Game

Question 12 *for £125,000*

The random variables X and Y have correlation zero.

Which statement is true?

A X and Y are definitely independent

B X and Y are definitely not independent

C X and Y may or may not be independent

D None of the above

The Millionaire Game

Question 12 *for £125,000*

The random variables X and Y have correlation zero.

Which statement is true?

A X and Y are definitely independent

C X and Y may or may not be independent

The Millionaire Game

Question 12 *for £125,000*

The random variables X and Y have correlation zero.

Which statement is true?

C X **and** Y **may or may not be independent**

The Millionaire Game

Question 13 *for £250,000*

What is

$$\int_{-1}^{1} -x \sin(1/x)\, dx$$

A π

B 2

C -1

D 0

The Millionaire Game

Question 13 *for £250,000*

What is

$$\int_{-1}^{1} -x \sin(1/x)\, dx$$

A π

D 0

The Millionaire Game

Question 13 *for £250,000*

What is

$$\int_{-1}^{1} -x \sin(1/x)\, dx$$

D 0

The Millionaire Game

Question 14 *for £500,000*

Which of the following is **not** a computer language

A Python

B Adder

C Cobra

D Actually, they are all computer languages

The Millionaire Game

Question 14 *for £500,000*

Which of the following is **not** a computer language

C Adder

D Actually, they are all computer languages

The Millionaire Game

Question 14 *for £500,000*

Which of the following is **not** a computer language

D Actually, they are all computer languages

The Millionaire Game

Question 15 *for £ 1 million*

Fermats Last Theorem, stated by Fermat around 1637, was finally proved in

A 1990

B 1994

C 1998

D 2002

The Millionaire Game

Question 15 *for £ 1 million*

Fermats Last Theorem, stated by Fermat around 1637, was finally proved in

B 1994

D 2002

The Millionaire Game

Question 15 *for £ 1 million*

Fermats Last Theorem, stated by Fermat around 1637, was finally proved in

B 1994

2. Millionaire Game: Optimal Strategy

Suppose you have reached $£16k$, you see the next question, and you have no idea which answer is correct.

You use your 50/50 lifeline, so you have two answers to choose from and, as far as you know, both are equally likely to be correct.

Do you take the $£16k$ you have, or do you guess an answer?

Expected gain

Stick:

$$£16,000$$

Guess:

$$0.5 \times £1,000 + 0.5 \times £32,000 \;=\; £16,500$$

Or is the "Guess strategy" worth more than this?

Millionaire Game: Optimal Strategy

Suppose you are at $£16k$ but have already used your 50/50 lifeline.

You see the next question, there are have four answers to choose from, and all seem equally likely to be correct.

Do you take the $£16k$, or do you guess an answer?

Expected gain

Stick: $£16,000$

Guess: $0.75 \times £1,000 + 0.25 \times £\,??? \;=\; £\,???$

If you guess correctly, your $£32k$ becomes your insurance prize.

Even if you have to guess next, the "value" of being at $£32k$ is

$$0.75 \times £32,000 + 0.25 \times £64,000 \;=\; £40,000$$

and the previous formula becomes

$$0.75 \times £1,000 + 0.25 \times £40,000 \;=\; £10,750$$

— but maybe you know the answer to the $£64,000$ question ...

Millionaire Game: Optimal Strategy

When deciding whether it is advantageous to guess an answer, we need to know the "value" of being at the next stage.

This means we should start at the end and work backwards.

Let's do this — but I shall simplify the game to make this easier (for now).

A simplified model for the game

For each new question,

with probability p you know the answer

with probability $1 - p$ all answers appear equally likely

There is a 50/50 lifeline (but no other lifelines).

The goal

To maximise the expected value of the amount we win.

Millionaire Game: Optimal Strategy

Notation

Define the "state" $x = (n, y)$ where $n \in \{1, \ldots, 16\}$, $y \in \{0, 1\}$

$$x = (\text{Number of questions answered so far}) + 1$$

$$y = \begin{cases} 1 & \text{if the 50/50 lifeline is still available} \\ 0 & \text{if not} \end{cases}$$

The value of state $x = (n, y)$

Define the "value" of the current state $x = (n, y)$ (in pounds) as

$$\beta(x) = E(\text{Final prize money} \mid \text{Currently at } x \text{ and play optimally})$$

The Bellman Equation

"Bellman's equation" relates $\beta(n, y)$ to $\beta(n+1, 1)$ and $\beta(n+1, 0)$.

Millionaire Game: Optimal Strategy

Case 1

If you have just answered the final question correctly, so $n = 16$:

$$\beta(16, 1) = \beta(16, 0) = 1,000,000$$

Case 2

At the last step **with no** 50/50 lifeline, so $x = (15, 0)$:

$$\begin{aligned}\beta(15,0) &= p \times 1,000,000 \\ &\quad + (1-p)\max\{500,000,\ 0.75 \times 32,000 + 0.25 \times \beta(16,0)\} \\ &= p \times 1,000,000 + (1-p) \times 500,000\end{aligned}$$

Optimal strategy

When $x = (15, 0)$: Don't guess.

Millionaire Game: Optimal Strategy

Case 3

At the last step **with** a 50/50 lifeline, so $x = (15, 1)$:

$$\begin{aligned}\beta(15,1) &= p \times 1,000,000 \\ &\quad + (1-p)\max\{500,000,\ 0.5 \times 32,000 + 0.5 \times \beta(16,0)\} \\ &= p \times 1,000,000 + (1-p) \times 516,000\end{aligned}$$

Optimal strategy

When $x = (15, 1)$: Use the lifeline, then guess.

Dynamic Programming

And we can carry on working backwards to compute

$$\beta(14,0),\ \beta(14,1),\ \ldots,\ \beta(2,0),\ \beta(2,1),\ \beta(1,1).$$

Millionaire Game: Optimal Strategy

Demonstration

R code for the simplified game

Possible extensions

- Add the other lifelines
- Refine the probability model
- Re-define the goal?

Millionaire Game: Optimal Strategy

A general sequential decision problem

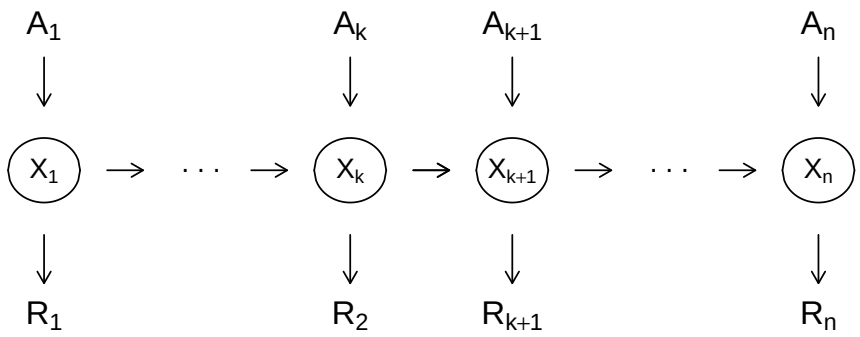

X_k State at stage k

A_k Action at stage k

R_k Reward at stage k

The Bellman equation

$$\beta(X_k) = \max_{A_k}\{E[R_k + \beta(X_{k+1}) \mid X_k, A_k]\}$$

3. Peter Armitage 1924–2024

Last week I gave the "Armitage Lecture" in Cambridge

Peter Armitage introduced the idea of "sequential analysis" to clinical trials.

In my lecture, I described Peter's work and subsequent research in this area, including some of my own work.

A Randomised Clinical Trial

Suppose a clinical trial is run to compare two treatments A and B.

Patients are treated in pairs, with one randomly allocated to Treatment A and the other to Treatment B.

Let X_i denote the difference in observed responses for pair i, where $X_i > 0$ if the patient on Treatment A has the better response.

Suppose $X_i \sim N(\theta, \sigma^2)$, $i = 1, 2, \ldots$.

We shall test H_0: $\theta = 0$ against $\theta \neq 0$.

If H_0 is rejected with $\widehat{\theta} > 0$ we conclude Treatment A is superior.

If H_0 is rejected with $\widehat{\theta} < 0$ we conclude Treatment B is superior.

We design the trial to have two-sided type I error probability 2α and power $1 - \beta$ at $\theta = \pm\theta_1$.

In a sequential design, the trial can be stopped at any point.

Armitage's Restricted Sequential Procedure

In a medical trial one would wish to have both a small average sample size and a small maximum sample size.

To achieve this, Armitage (*Biometrika*, 1957) proposed "Restricted Sequential Procedures" with upper and lower stopping boundaries

$$\sum_{i=1}^{n} X_i = a + b\,n \quad \text{and} \quad \sum_{i=1}^{n} X_i = -a - b\,n$$

and truncation at $n = N$.

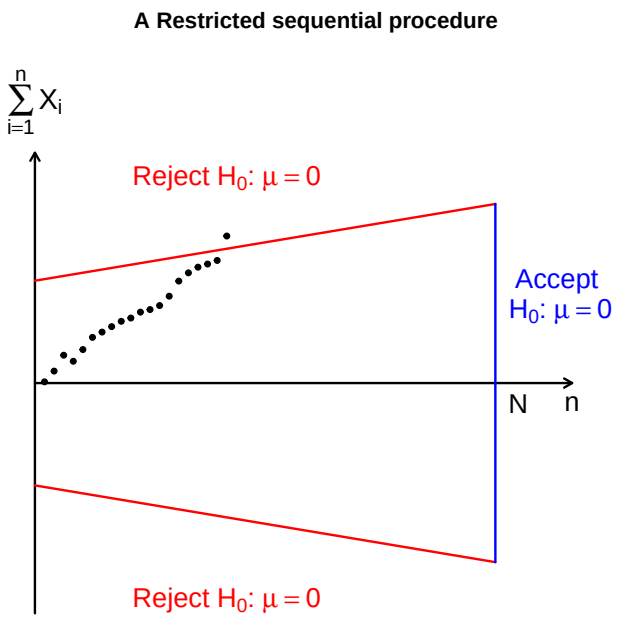

Armitage's Restricted Sequential Procedure

With a normally distributed response, some neat mathematics was needed to make calculations feasible.

Armitage applied a result proved by Bartlett (1946) for diffusion processes and a likelihood ratio argument similar to that used by Wald (1945) to define the "Sequential Probability Ratio Test".

Armitage set

$$a = \frac{\sigma^2}{\theta_1} \log\left(\frac{1-\beta}{\alpha}\right), \quad b = \frac{\theta_1}{2}$$

and found N satisfying

$$\beta = \Phi\left(\frac{a}{\sigma\sqrt{N}} - \frac{b\sqrt{N}}{\sigma}\right) - \left(\frac{1-\beta}{\alpha}\right)\Phi\left(\frac{-a}{\sigma\sqrt{N}} - \frac{b\sqrt{N}}{\sigma}\right).$$

Calculation is needed, using tables of the standard normal CDF Φ.

Calculating machines

The Braunsviga Calculating Machine

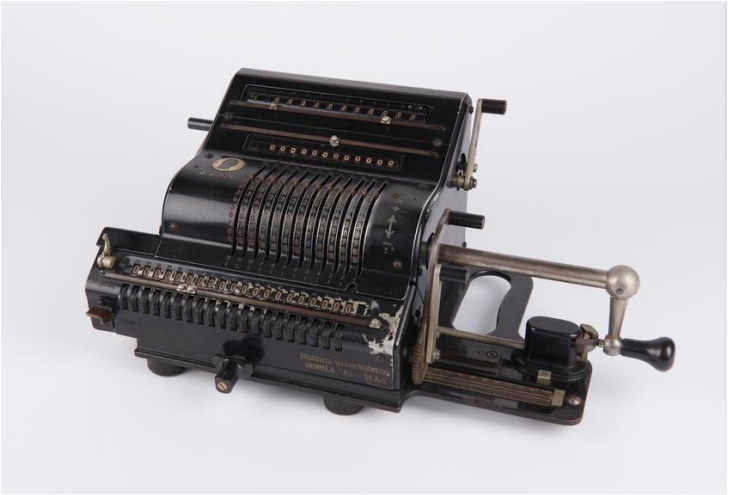

Cambridge research students used to spend time each week doing the calculations needed to produce statistical tables.

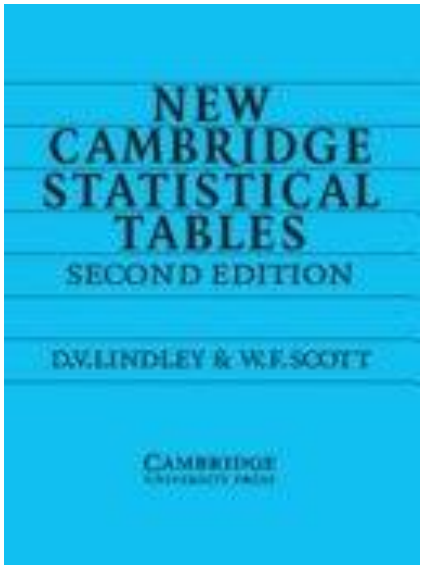

Armitage's Restricted Sequential Procedure

Schneiderman & Armitage (*Biometrika*, 1962) inserted a "wedge" in the continuation region to facilitate early stopping for $\theta \approx 0$.

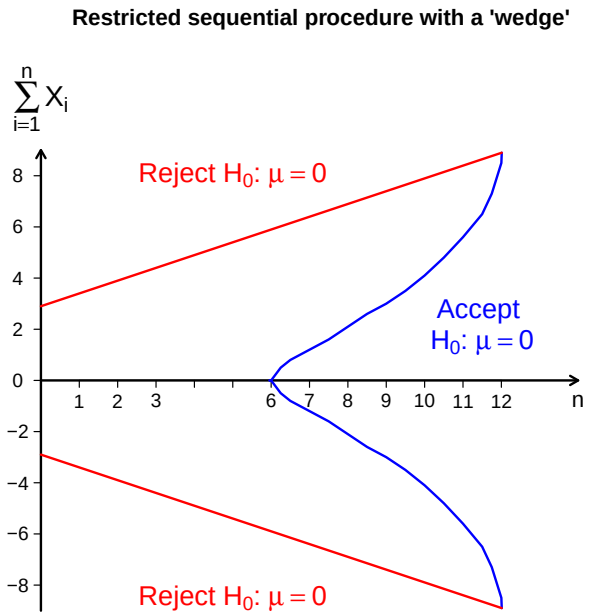

The boundary for the "wedge" was derived analytically but this still required significant computation.

Average sample sizes were computed by Monte Carlo simulation with 100 replicates

Calculations were carried out using the NIH's IBM 650 computer.

IBM 650 computer

The IBM 650

The IBM 650's
Magnetic memory drum

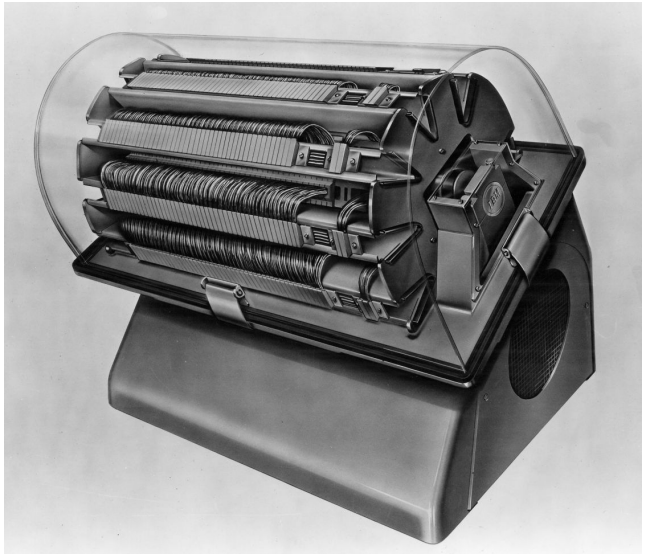

"The average time for accessing data or programming was 2.4 milliseconds, less than the time it takes a fruit fly to flap its wings"

Compare current terminology: A petaflop is one quadrillion (10^{15}) floating-point operations per second.

Group Sequential Tests

Interest moved on to designs with a small number of analyses.

These "Group Sequential" methods suited a multi-centre trial with an Independent Data Monitoring Committee and a separate party (e.g., a Clinical Research Organisation) cleaning and analysing the accruing data.

In a comparison of a new treatment to a control, it was noted that the hypothesis testing formulation should be one-sided.

If θ is the improvement from using the new treatment (the "treatment effect"), we should test H_0: $\theta \leq 0$ against $\theta > 0$.

If $\theta \leq 0$, it would be unethical to randomise patients in order to learn whether $\theta = 0$ or $\theta < 0$.

"Error Spending" designs were proposed to handle unequal group sizes or, more generally, unequal increments in information.

Optimal Stopping Boundaries

One can ask

What is the best choice of stopping boundary?

What is the best choice of error spending function?

What is needed for these questions to be well-posed?

Consider a test of H_0: $\theta \leq 0$ against $\theta > 0$ with type I error probability α and power $1-\beta$ at $\theta = \delta$.

A fixed sample size study needs information

$$\mathcal{I}_{fix} = \frac{\{\Phi^{-1}(1-\alpha) + \Phi^{-1}(1-\beta)\}^2}{\delta^2},$$

where Φ is the standard normal CDF.

Here, "Information" is the reciprocal of the variance of $\widehat{\theta}$ and is (roughly) proportional to sample size in many clinical trial settings.

Optimal Stopping Boundaries

A group sequential test (GST) with K analyses will require a maximum information possible level $\mathcal{I}_K$, greater than $\mathcal{I}_{fix}$.

We call $R = \mathcal{I}_K / \mathcal{I}_{fix}$ the *inflation factor* of a group sequential test.

We can seek a GST that minimises expected information $\mathbb{E}_\theta(\mathcal{I})$ under certain values of the treatment effect, θ, with a given number of analyses K and inflation factor R.

We may aim to minimise

$$\sum_i w_i \, \mathbb{E}_{\theta_i}(\mathcal{I})$$

for selected treatment effects θ_i and weights w_i.

Alternatively, we may minimise

$$\int f(\theta) \, \mathbb{E}_\theta(\mathcal{I}) \, d\theta,$$

where f is, say, a normal density.

Computing Optimal Group Sequential Tests

In optimising a GST, Eales & Jennison (*Biometrika*, 1992) and Barber & Jennison (*Biometrika*, 2002) create a Bayes sequential decision problem, placing a prior on θ and defining costs for sampling and for making incorrect decisions.

Such a problem can be solved rapidly using the numerical integration methods of Armitage, McPherson & Rowe (1969), combined with **Dynamic Programming**.

One then searches for the combination of prior and costs such that the solution to the (unconstrained) Bayes decision problem has the specified frequentist error rates α at $\theta = 0$ and β at $\theta = \delta$.

The resulting design solves both the Bayes decision problem and the original frequentist problem.

Benefits of Group Sequential Testing

One-sided GSTs with binding futility boundaries, minimising $\{\mathbb{E}_0(\mathcal{I}) + \mathbb{E}_\delta(\mathcal{I})\}/2$ for K equally sized groups, $\alpha = 0.025$, $1 - \beta = 0.9$ and $\mathcal{I}_{max} = R\mathcal{I}_{fix}$.

Minimum values of $\{\mathbb{E}_0(\mathcal{I}) + \mathbb{E}_\delta(\mathcal{I})\}/2$, as a percentage of $\mathcal{I}_{fix}$

	R					
K	1.01	1.05	1.1	1.2	1.3	*Minimum over R*
2	80.8	74.7	**73.2**	73.7	75.8	73.0 at R=1.13
3	76.2	69.3	66.6	**65.1**	65.2	65.0 at R=1.23
5	72.2	65.2	62.2	59.8	**59.0**	58.8 at R=1.38
10	69.2	62.2	59.0	56.3	**55.1**	54.2 at R=1.6
20	67.8	60.6	57.5	54.6	**53.3**	51.7 at R=1.8

Note: $\mathbb{E}(\mathcal{I}) \searrow$ as $K \nearrow$ but with diminishing returns,

$\mathbb{E}(\mathcal{I}) \searrow$ as $R \nearrow$ up to a point.

Efficient Error Spending GSTs

In their book, *Group Sequential Methods with Applications to Clinical Trials*, Jennison & Turnbull (1999) suggest using the "ρ-family" of one-sided error spending tests.

A target information level $\mathcal{I}_{max}$ is specified and the type I and type II error probabilities "spent" up to analysis k are, respectively,

$$f(\mathcal{I}) = \min\{(\mathcal{I}/\mathcal{I}_{\max})^2, 1\}\,\alpha \quad \text{and} \quad g(\mathcal{I}) = \min\{(\mathcal{I}/\mathcal{I}_{\max})^2, 1\}\,\beta.$$

The value of ρ governs the rate at which error probability is spent, with $\rho = 1$ producing Pocock-type boundaries and $\rho = 3$ producing O'Brien & Fleming-type boundaries.

The choice of ρ determines the inflation factor R and $\mathcal{I}_{max}$ is R times the information needed for a fixed sample test.

Barber & Jennison (2003) show this family yields tests that are close to optimal for a variety of measures of $\mathbb{E}_\theta(\mathcal{I})$.

Efficient Error Spending GSTs

Plots of $\int f(\theta) E_\theta(\mathcal{I}_T) d\theta$ as a percentage of fixed sample $\mathcal{I}_{fix}$ vs inflation factor R for tests with 5 analyses, $\alpha = 0.025$ and $\beta = 0.1$.

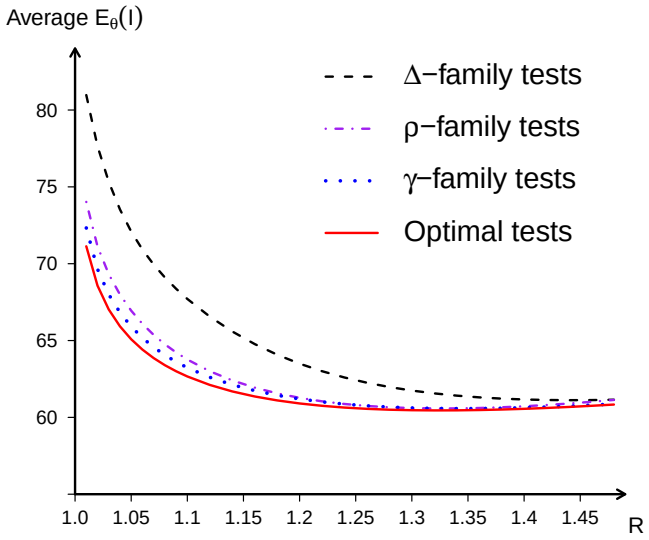

Here, $f(\theta)$ is the density of a $N(\delta, \delta^2/4)$ distribution and $\mathcal{I}_{fix}$ the information required for a fixed sample size test.

The Δ family of parametric boundaries is that proposed by Wang & Tsiatis (*Biometrics*, 1987).

The Gamma family of error spending functions is as described by Hwang, Shih & De Cani (*Statistics in Medicine*, 1990).

Efficient Error Spending GSTs

A test with 5 planned analyses, type I error probability $\alpha = 0.025$, power 0.9 if $\theta = \delta = 1$, and type I and II error spending functions

$$f(\mathcal{I}) = \min\{(\mathcal{I}/\mathcal{I}_{\max})^2, 1\}\,\alpha, \quad g(\mathcal{I}) = \min\{(\mathcal{I}/\mathcal{I}_{\max})^2, 1\}\,\beta.$$

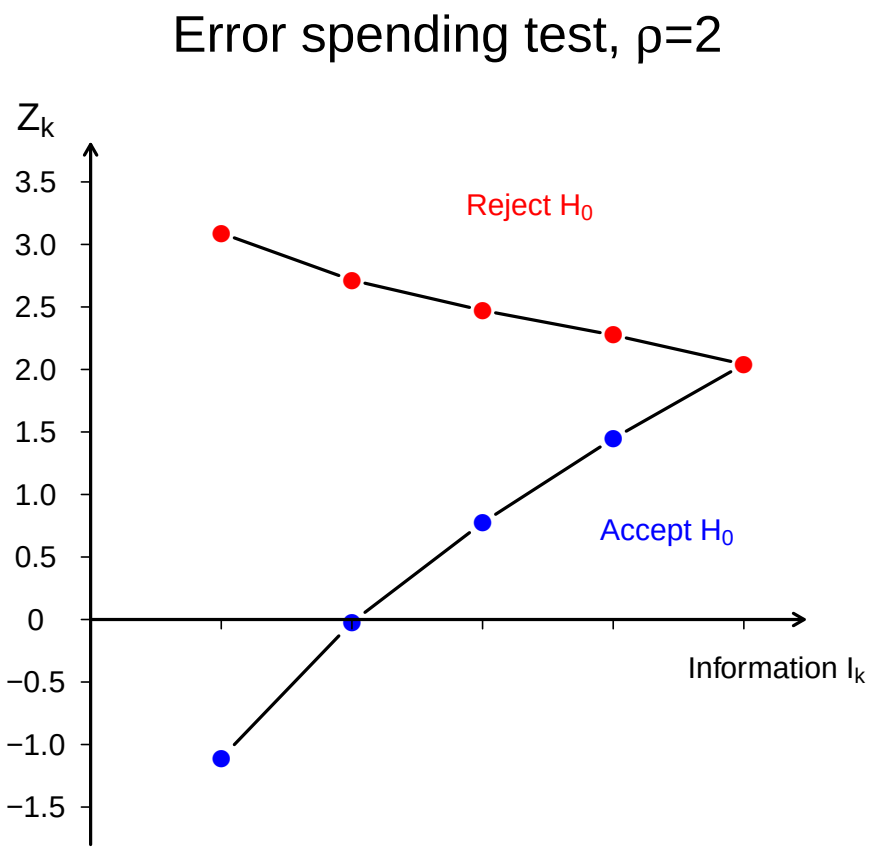

$$R = 1.132.$$

Efficient Error Spending GSTs

This stopping boundary makes a nice book cover.

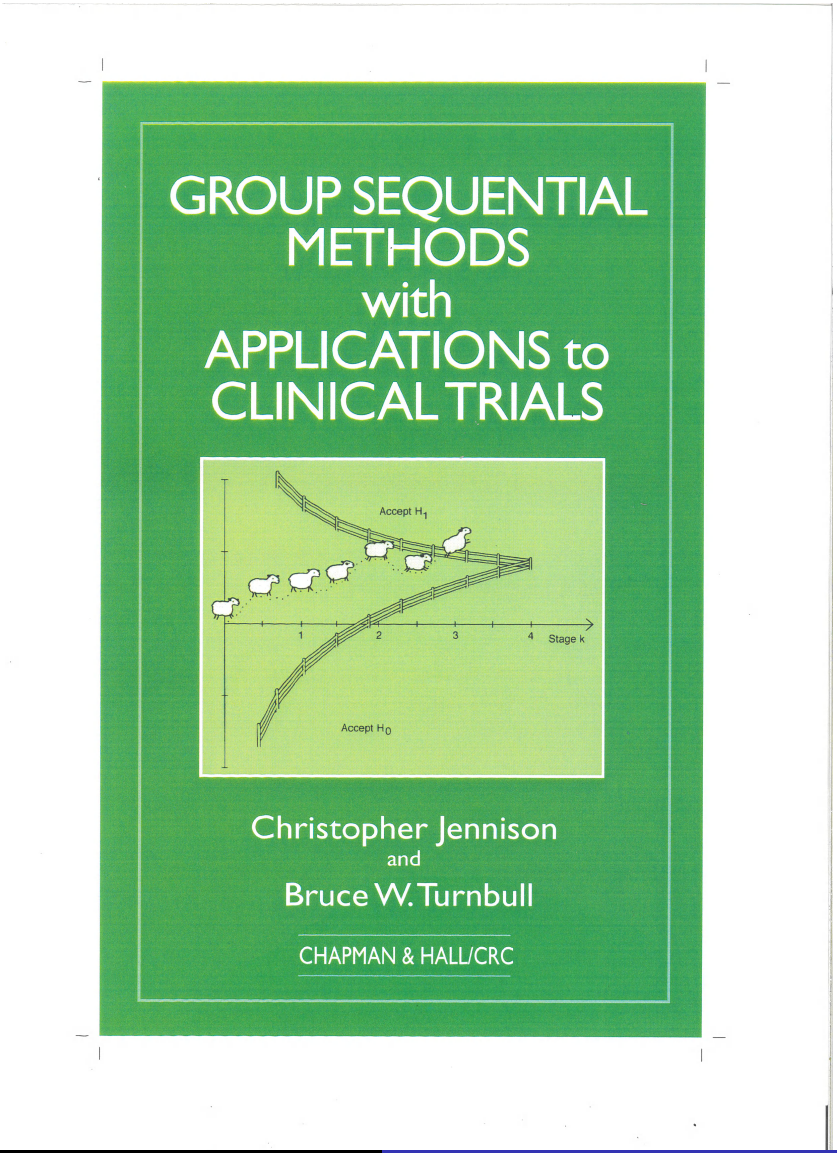

Finding optimal GSTs by Dynamic Programming

Similar methods can be used to

Optimise timing of analyses,

Allow data dependent group sizes (Jennison & Turnbull, *Biometrika*, 2006).

However, there is little to be gained from these embellishments.

Other applications of this method of optimisation include:

Group sequential tests of superiority and non-inferiority (Öhrn & Jennison, *Statistics in Medicine*, 2010),

Group sequential tests that can deal with "pipeline data" (Hampson & Jennison, *JRSS, B*, 2013),

Optimising gain functions from financial models (Robbie Peck, *University of Bath, PhD thesis*, 2020).

Optimising dose allocation in a First in Human trial

Reference: Lizzi Pitt, *University of Bath, PhD thesis*, 2021.

Phase I, First in Human, trials are conducted to investigate the safety of a new molecule and find the maximum tolerated dose.

Patients are treated a few at a time, e.g., in cohorts of 3.

The dose escalation scheme only moves on to a higher dose when lower doses have been shown to be safe.

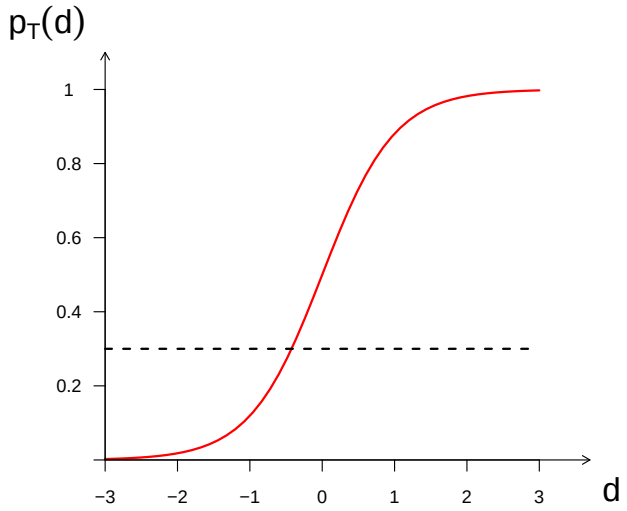

The aim is to determine the Maximum Tolerated Dose (MTD), at which the probability of a Dose Limiting Event (DLE) is, say, 0.3.

Dose response function (O'Quigley, et al. 1990, *Bmcs*)

We assume a dose response function of the form

$$p_T(d,a) = \left(\frac{\tanh(d)+1}{2}\right)^a = \left(\frac{1}{1+e^{-2d}}\right)^a,$$

where $p_T(d,a)$ is the probability of a DLE at dose level d, on some transformed scale (e.g., logarithmic).

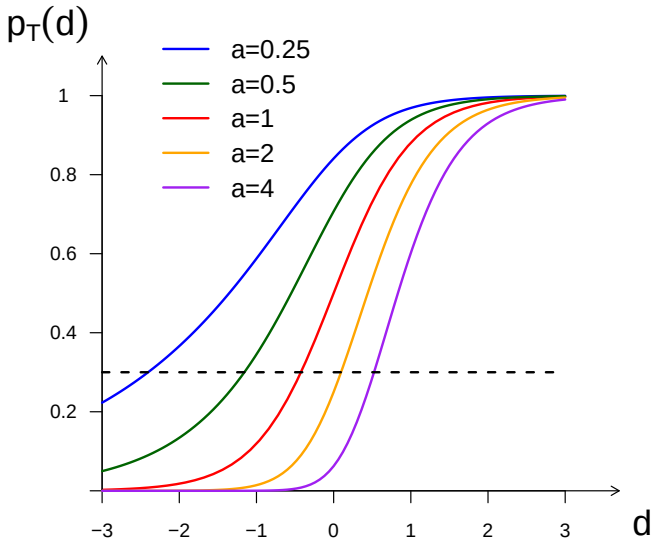

The parameter a determines the likelihood of a DLE.

As a increases, $p_T(d,a)$ decreases.

Problem formulation and loss function

Suppose a maximum of J cohorts of n_c subjects are to be treated with doses selected from the set $\mathcal{D} = \{d_1, \ldots, d_m\}$.

At the end of the study, with data x_J, a dose $d^* \in \mathcal{D}$ will be selected as the estimate of the maximum tolerated dose.

We define the loss function

$$L(d^*, a, x_J) \;=\; |p_T(d^*, a) - 0.3| + \delta\, n_T,$$

where n_T is the number of DLEs that occur in the trial.

We place a prior $\pi(a)$ on the dose response model parameter a.

Our aim is to minimise the expected loss

$$\int (E\{|p_T(d^*, a) - 0.3|\} + \delta\, n_T)\, \pi(a)\, da.$$

over possible final decision rules and dose allocation rules that satisfy specified constraints on dose escalation.

Optimising by Dynamic Programming

The state variable

After j cohorts of n_c have been observed, the data so far comprise

$$\{U_1, Y_1, \ldots, U_j, Y_j\},$$

where $U_j \in \{1, \ldots, m\}$ is the index of the dose chosen for cohort j and Y_j the number of DLEs observed in cohort j.

It is (usually) sufficient to summarise by the state variable

$$X_j = (N_{1,j}, \ldots, N_{m,j}, V_{1,j}, \ldots, V_{m,j})$$

where $N_{i,j}$ is the total number of cohorts allocated dose d_i and $V_{i,j}$ the total number of DLEs observed at that dose.

With 10 cohorts of 3 and 6 doses, there are 16 million possible states at the final analysis and 25 million over the whole trial.

Solution by Dynamic Programming

Lizzi showed this can be done!

Conclusions

We now know how to play Who Wants to be a Millionaire?

We have seen the method of Dynamic Programming

Time for the "Pub" part of this Pub Lecture