GRESHAM

COLLEGE

FACTORIALS, FORTUNES AND FALLACIES

A Lecture by

PROFESSOR IAN STEWART MA PhD FIMA CMath
Gresham Professor of Geometry

25 October 1995

Gresham Geometry Lecture
25 October 1995

Factorials, Fortunes, and Fallacies

or

How to win the Lottery (or at least not lose by much)

The odds against winning the jackpot on the National Lottery are precisely 13,983,815 to 1. If you placed one bet every week you would expect to win, on average, once every quarter of a million years.

How do we know that?

This lecture examines the mathematics behind the lottery — probability theory — and exposes several popular fallacies about it. For example: do you improve your chances of winning by sticking to the same numbers every week, or by swapping? Are your chances improved by avoiding last week's numbers? Does a random-looking sequence like 4,19,25,27,42,46 have a better chance than a patterned one like 1,2,3,4,5,6?

It also examines strategies and systems for maximizing winnings. What sequence you bet on, and when you bet, *do* make a difference.

In 1995 the British government introduced a National Lottery, similar to those run by several other countries including Germany. The proceeds are divided between Camelot (the company that runs it), the government, various worthy causes such as charities, sport, and the arts, and the winning punters. The lottery predictably proved controversial — not because of any irregularities in its running, but for political and ethical reasons. (It was, for example, dubbed a 'desperation tax' by the Methodist church.) Setting politics and ethics aside, let's try out the methods of probability theory on this topical example to get some insight into its mathematical structure.

The rules are simple. For £1, a player buys a lottery card and chooses six distinct numbers between 1 and 49 inclusive — for example 5, 6, 17, 31, 32, 47. A seventh 'bonus number' must also be chosen, but for simplicity I'll ignore that. Every week six random numbers (plus a bonus number which I've just promised not to mention) are generated by a machine that mixes up plastic balls with numbers on and blows them one by one into a tube. In order to win prizes, players must match the numbers marked on their cards to those selected by the machine. The order in which the numbers are chosen or drawn is irrelevant.

The top prize, the jackpot, is typically worth about £8 million, and goes to those players who get all six numbers correct. Successively smaller prizes are given for getting five numbers out of the six main ones, plus the bonus; five numbers out of the six main ones; four numbers out of the six main ones; and three numbers out of the six main ones. (To keep the discussion simple I will ignore everything except the jackpot.)

In the first draw for the lottery nobody won the jackpot, which was therefore 'rolled over' to the next week — meaning that the prize money for the second week's jackpot was added to that for the first. In week two a factory worker from the North of England duly won £17,880,003, with the choice of numbers 26, 35, 38, 43, 47, 49, and Lottery Fever struck. For example, an insurance company offered a policy to protect employers against a key member of staff winning the jackpot and leaving. How worthwhile would such a policy be to a company? How do the odds of winning the jackpot compare with those for being killed in a road accident — against which few companies insure their employees, however important?

There are many people who offer to sell 'systems' for improving your chances of a win. Are any of these systems effective? Do you stand a better chance of winning if you stick to the same numbers every week, or if you swap? Do 'random' sequences like 4, 11, 15, 23, 31, 42 have a better chance of winning than patterned ones like 1, 2, 3, 4, 5, 6?

What are your chances of winning, anyway? How do we know?

Probability

To answer these and other questions about the lottery we need to introduce some ideas from probability theory. In fact my main aim in this section of the lecture is to build some intuition about randomness and probability that applies to many real world situations, not just the lottery.

Randomness is a tricky concept. For the moment let's say that a sequence of events is *random* if the next event does not depend in any manner upon the previous events. If I roll an unbiased die, and it comes up 6 five times in a row (unlikely, but possible), then this has no effect on the next throw. It will be either 1, 2, 3, 4, 5, or 6, and there is no reason to expect any number to be more likely than the others.

Random systems have no memory.

The actual definition of probability that mathematicians use is that it is a numerical quantity associated with certain *events* which satisfies certain properties. There is a theorem (called the Law of Large Numbers) which allows us to *interpret* probability as a long-term frequency. Suppose that you conduct some 'trial' that has a number of different events A, B, C, etc. Then the probability p(A) of outcome A can be interpreted as the proportion of trials in which event A occurs. More precisely, it is the limiting value of that proportion in a very long series of trials. I repeat that this is an interpretation, not a definition: the failure to appreciate this point underlies several misunderstandings about the nature of probability.

Example

Trial: toss a coin.

Events H (heads) and T (tails). [E (edge) is excluded in this particular example, although for a real coin it does occasionally happen. *Very* occasionally.]

Probabilities: $p(H) = \frac{1}{2}$, $p(T) = \frac{1}{2}$.

This mathematical system is called a *fair coin.* The coins in your pocket are very well modelled by a fair coin if you perform a real-world experiment known as 'tossing'.

Here is the result of an actual experiment carried out 20 times using a £1 coin: I promise I didn't just invent it.

T T T T H T H H H H H H T T T H T T T H (1)

There are 11 T's and 9 H's, frequencies of 11/20 = 0.55 and 9/20 = 0.45. These are close to 0.5, but not exactly equal.

You may object that my sequence is unusual — it doesn't look random enough. Actually, it does — but our psychological makeup misleads us on what randomness looks like. You'd probably be much happier with something like

H T H H T T H T T H T H H T H T H H T T (2)

with frequencies 10/20 = 0.5 and 10/20 = 0.5. As well as getting the numbers spot on, the second sequence looks more random.

But it isn't.

What makes sequence (1) look non-random is that there are long repetitions of the same event, such as T T T T and H H H H H H. Sequence (2) lacks such repetitive sequences, so we think it looks random. But random sequences *should* contain repetitions. For example, if you look at successive blocks of four events, like this:

T T T T H T H H H H H H T T T H T T T H
T T T T
T T T H
T T H T
T H T H
etc.

then T T T T should occur about one time in 16 (I'll explain why in a moment). In fact here it occurs once in 17 times — pretty much spot on! Agreed, H H H H H H should occur only once in 64 times, on average — but I didn't throw my coin enough times to see whether it came up again later. *Something* has to come up, and H H H H H H is just as likely as H T H T H T or H H T H T T.

Random sequences often show occasional signs of patterns and clumps. Do not be surprised by these: they are to be expected. They are *not* signs that the process is not random. Unless the coin goes H H H H H H H H H H H H..., in which case the shrewd person would guess that perhaps it is double-headed...

Simulated Lottery Draws

I wrote a computer program to simulate the lottery machine using random numbers.. The first 10 draws were like this:

4	6	32	36	40	43
6	23	25	27	35	47
16	27	28	46	47	48
8	17	23	33	35	37
3	21	26	27	37	38
7	12	23	28	33	37
3	15	18	24	28	49
16	19	33	40	44	47
11	21	34	36	38	46
7	9	13	24	26	30

You can see all sorts of apparent 'patterns' and 'clumps' here. But in fact (because of the way they are generated) we know these sequences are entirely typical of randomly drawn ones.

On the 23rd trial, by the way, I got

30 31 34 35 36 43

On the 26th I got

3 4 5 12 14 34

and on the 85th

1 3 8 8 17 19

so don't expect random choices to be evenly spread out.

Suppose you toss four coins in a row. What can happen?

Fig.1 (next page) summarizes all the possible results. The first toss is either H or T (each with probability $\frac{1}{2}$). Whichever of these happens, the second toss is also either H or T (each with probability $\frac{1}{2}$). Whichever of these happens, the third toss is also either H or T (each with probability $\frac{1}{2}$). And whichever of these happens, the fourth toss is also either H or T (each with probability $\frac{1}{2}$). So we get a 'tree' with sixteen possible routes through it. According to the rules for manipulating probabilities, each route has probability

$$\frac{1}{2} \times \frac{1}{2} \times \frac{1}{2} \times \frac{1}{2} = \frac{1}{16}.$$

This is plausible, because there are 16 routes, and each should be equally likely.

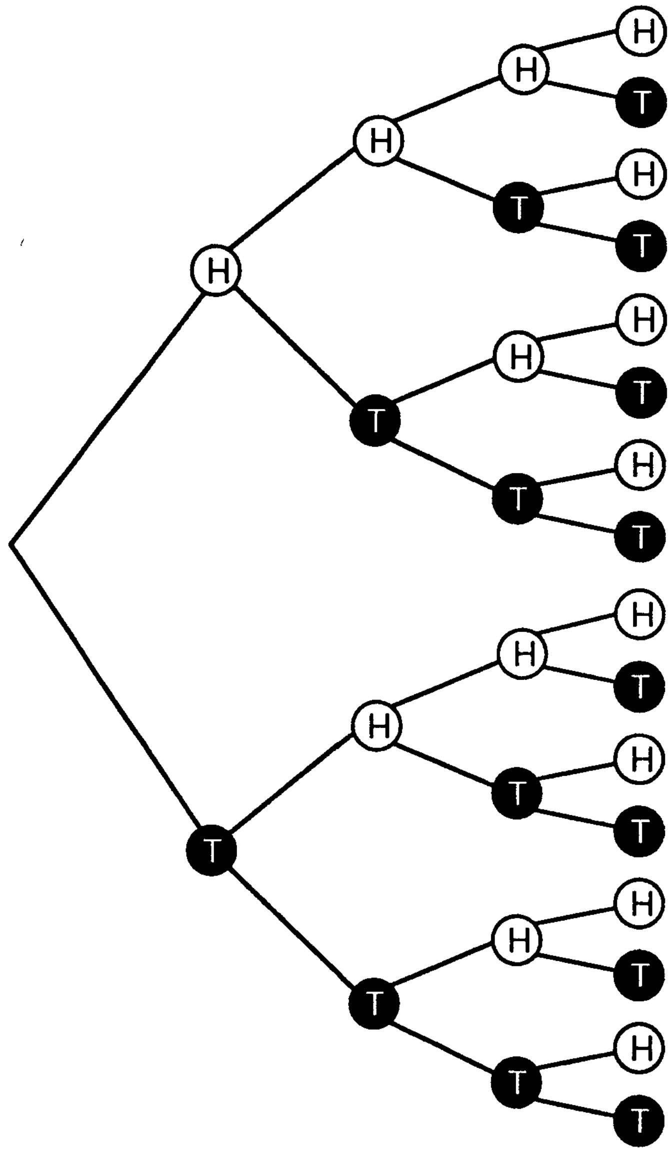

Fig.1 Tree of successive events shows all possibiities for four consecutive coin tosses.

Notice that T T T T has probability $\frac{1}{16}$, and H T H H (say) *also* has probability $\frac{1}{16}$. Although H T H H looks 'more random' than T T T T, they have the same probability.

The point to understand here is that it is the *process* of tossing a coin that is random. This does not oblige the *results* to look random too. Usually they do — but that's because most sequences of H's and T's don't have much pattern. In fact Gregory Chaitin has *defined* randomness as lack of pattern — more precisely, he says that a sequence is random if if cannot be generated by a computer program shorter than the sequence itself — and he has proved that in this sense almost all *long enough* sequences are random.

The 'Law of Averages'

The basis of many misconceptions about probability is a belief in something usually referred to as 'the law of averages', which asserts that any unevenness in random events gets ironed out in the long run. For example, if a tossed coin keeps coming up heads, then at some stage there will be a predominance of tails to balance things out.

We'll see later that even though there is a sense in which this is true, the usual way people interpret it is wrong. For example, tails do not become 'more likely' if you toss lots of heads. The probability of a tail remains at $\frac{1}{2}$.

You might like to try an experiment to test the alleged 'law'. Take ten coins. Toss each coin repeatedly until you get a run of (say) four heads, H H H H. Does such a run somehow improves the odds on getting tails on the next throw? Toss each of the ten 'pre-prepared' coins once, and see how many tails you get. (For a more sophisticated experiment, do this lots of times and see how the results vary: then compare with the theoretical predictions for ten tosses of an unbiased coin, to which we now turn.)

Factorials

If you toss ten independent fair coins, what is the probability of getting exactly five heads?

Let's start with an easier one. If you toss a coin four times, what is the probability of getting exactly two heads? Look at Fig. 1. There are 16 different sequences of H's and T's, and you can count how many of them contain exactly two heads. There are 6 of them:

HHTT
HTHT
HTTH
THHT
THTH
TTHH

so the probability of exactly two heads is $\frac{6}{16} = 0.375$. Devoteees of the law of averages should note that this is *less* than the probability of *not* getting exactly two heads, which is $1 - 0.375 = 0.625$.

What about five heads in ten tosses? There are 1024 different sequences of ten H's and T's, so listing all cases is not such a good idea. We need a more general method, and I'll illustrate the beginnings of an area nowadays often called 'discrete mathematics' or 'combinatorics' that can be used to answer such questions.

First, a simpler question. If I have, say, seven objects A B C D E F G, how many different ways can I arrange them in order?

There are 7 choices for the first object. Having selected that one, there are only 6 choices for the second object. Then 5 choices for the third, 4 choices for the fourth, 3 choices for the fifth, 2 choices for the sixth, and 1 choice for the seventh. So the total number of possibilities is obtained by *multiplying* these numbers of choices:

$$7\times6\times5\times4\times3\times2\times2\times1.$$

This works out as 5040, and it is called *factorial* 7 [people often say 7 *factorial* too] and denoted 7!

Similarly, the number of ways to arrange n objects in order is

$$n! = n(n-1)(n-2)...3.2.1.$$

Now back to the question: how many sequences of ten H's and T's contain exactly 5 H's? One way to think about this is to imagine choosing the positions of the 5 H's from the 10 'slots' in the sequence.

The first H can go into any of 10 slots.
The second H can go into any of 9 slots.
The third H can go into any of 8 slots.
The fourth Hcan go into any of 7 slots.
The fifth H can go into any of 6 slots.

Multiplying up, we get

$10 \times 9 \times 8 \times 7 \times 6 = 30240.$

Hang on, that can't be right: there are only 1024 sequences. What's happened is that we have counted each arrangement many times. For example H H H H H T T T T T comes up if we allocate heads to slots 1 2 3 4 5 in that order, or 2 1 3 4 5, or 5 3 1 4 2, or... In fact the number of times that a given sequence, such as H H H H H T T T T T, comes up is the number of ways to place five objects (the slots) in order. This is 5! = 120. So we have to divide 30240 by 120, which gives

$30240/120 = 252.$

This means that the probability of getting exactly five heads with ten coin tosses is 252/1024 = 0.246.

This is even *smaller* than the probability of getting two heads in four tosses. But aren't the numbers supposed to balance out in the long run? Doesn't look as if they are likely to balance out *exactly*, does it?

If you look back to our calculation you find that the number of ways to get 5 heads in 10 trials (252) came about as

$$\frac{10.9.8.7.6}{5.4.3.2.1}.$$

By similar arguments, the number of ways to get r heads in n trials is

$$\frac{n(n-1)...(n-r+1)}{r(r-1)...3.2.1}$$

which is called a *binomial coefficient* and written as $\binom{n}{r}$. This can be rewritten using factorials as the famous formula

$$\binom{n}{r} = \frac{n!}{r!(n-r)!}.$$

Thus armed, you can calculate the number of sequences of ten coin tosses that contain exactly r = 0, 1, 2, 3, ..., 10 H's. Divide those numbers by 1024 and you get the probabilities of obtaining r heads in 10 tosses. The results are:

r	$\binom{10}{r}$	$\binom{10}{r}/1024$
0	1	0.000976
1	10	0.009765
2	45	0.043945
3	120	0.117187
4	210	0.205078
5	252	0.246093
6	210	0.205078
7	120	0.117187
8	45	0.043945
9	10	0.009765
10	1	0.000976

OK, that's given you some technique. Before applying it to the lottery, though, I'd like to pick up on the 'law of averages' point again.

Random Walks

The calculations and experiments we've just done make it clear that there is no 'law of averages', in the sense that the future probabilities of events are *not* changed in any way by what happened in the past.

This means that lottery systems that are based on analysing past draws are totally useless. The more clever the pattern-detecting software is, the more cleverly it's analysing the wrong thing. The patterns it thinks it finds are spurious and irrelevant, and predictions based on those apparent patterns are nonsense.

However, there is an interesting sense in which things *do* tend to balance out in the long run. You can plot the excess of the number of H's over the number of T's by drawing a graph of the *difference* at each toss. You can think of this as a curve that moves one step upwards for each H and one down for each T. So the opening sequence (1), which was

T T T T H T H H H H H H T T T H T T T H,

produces the graph of **Fig.2**.

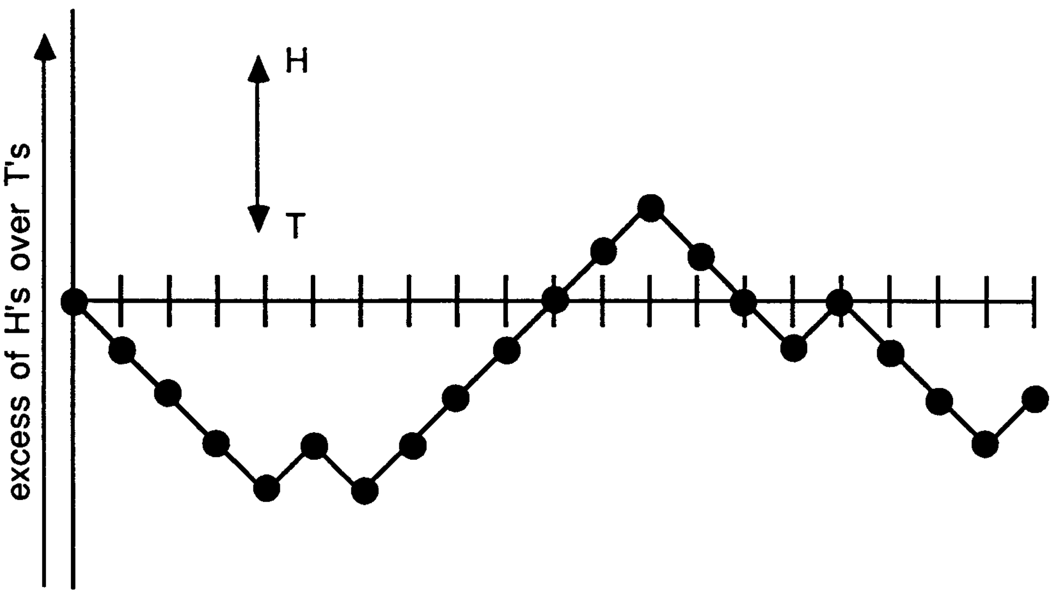

Fig.2 Random walk representing excess of heads over tails.

That establishes the principle, but the picture may still make you think that the numbers ballance out pretty often. So here's a graph of a typical random walk corresponding to 10,000 tosses. Heads spend a lot of time in the lead.

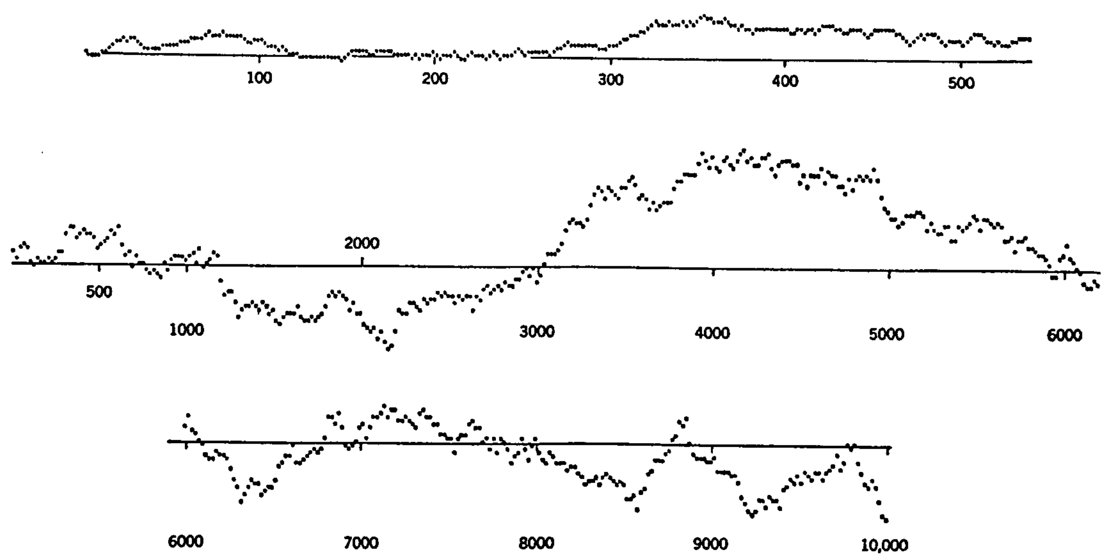

Fig.3 Random walk for 10,000 tosses.

If you reverse the above random walk (corresponding to tossing the coins in backeward-moving time) you find that now

HEADS ARE IN THE LEAD	TAILS ARE IN THE LEAD
7804 tosses	8 tosses
2 tosses	54 tosses
30 tosses	2 tosses
48 tosses	6 tosses
2046 tosses	

This sort of wildly unbalanced behaviour is entirely normal. In fact random walk theory shows that the probability that, in 10,000 tosses, one side leads for 9930 tosses and the other for only 70, is about one in ten.

However, random walk theory also tells us that the probability that the balance never returns to zero (that is, that H stays in the lead *forever*) is 0. (The average time taken to return, by the way, is infinite.) This is the sense in which the 'law of averages' is true. It carries no implications about improving your chances of winning, though, if you're betting on whether H or T turns up. The point is, you don't know how long the long run is going to be — except that it is likely to be very long indeed.

OK. Suppose you toss a coin 100 times and at that stage you have 55 H's and 45 T's — an imbalance of 10 in favour of H's. Random walk theory says that if you wait long enough, the balance will (with probability 1) correct itself.

Isn't that the 'law of averages'?

No. Not as that 'law' is normally interpreted. If you choose a length in advance — say a million tosses — then random walk theory says that those million tosses are unaffected by the imbalance. In fact, if you made huge numbers of experiments with one million extra tosses, then on average you would get 500,055 H's and 500,045 T's in the combined sequence of 1,000,100 throws. (On average, imbalances *persist*. Notice however that the *frequency* of H changes from 55/100 = 0.55 to 500055/1000100 = 0.500005. The 'law of averages' asserts itself not by *removing* imbalances, but by swamping them.)

What random walk theory tells us is that if you wait long enough, then eventually the numbers will balance out. If you stop at that instant, you may imagine that your intuition about a 'law of averages' is justified. But you're cheating: you stopped when you got the answer you wanted. Random walk theory also tells you that if you carry on for long enough, you will reach a situation where the number of H's is a million more than the number of T's. If you stopped *there*, you'd have a very different intuition!

Enough theory. Now let's take a look at the lottery.

Winning Probabilities

What *is* your chance of winning the jackpot? Assume that the machine that chooses the numbers is statistically random. All the evidence supports this assumption, and indeed there are good mathematical reasons to expect it to be true, based on the idea that the machine 'does not know' which numbers are written on the balls. That is, in principle it treats all balls equally. So the probability of any particular number being chosen is 1/49. A human player may think that the numbers 3 and 7 are 'lucky' and 13 is 'unlucky', but statistical analysis of the numbers produced over many years in other countries shows that the machine isn't superstitious.

Assume you've chosen 5, 6, 17, 31, 32, 47. What is the probability that the machine will produce those numbers?

The probability that the first number drawn by the machine matches one of your six is clearly 6/49, because there are six possible numbers to choose out of a total of 49. If that number *is* one of yours, then the probability that the second number drawn by the machine matches one of your other five choices is clearly 5/48, because now there are only five possible numbers left to to choose out of a total of 48. Similarly the probability of the machine choosing one of the remaining four numbers on the next try is 4/47, and so on with successive probabilities 3/46, 2/45, and 1/44.

According to probability theory, then, the probability of all six numbers matching those you chose can be found by multiplying all of these fractions together:

$$\frac{6}{49} \times \frac{5}{48} \times \frac{4}{47} \times \frac{3}{46} \times \frac{2}{45} \times \frac{1}{44} = \frac{1}{13983816}$$

So the chance of winning the jackpot is precisely 1 in 13,983,816 each week. In bookie's jargon, the *odds against* are 13,983,815 to 1.

This calculation has a factorialish look to it, so let's do it using the methods established earlier. You are choosing 6 balls out of 49, and the number of ways to do this is

$$\binom{49}{6} = \frac{49.48.47.46.45.44}{6.5.4.3.2.1} = 13{,}983{,}816$$

This is the same fraction as before but upside down.

Notice that this calculation produces the same answer for *any* set of six numbers between 1 and 49. All choices of six numbers, no matter what patterns do or do not occur, have exactly the same chance of being chosen by the machine.

Expectation

Your *expectation* is the amount you win, on average, per attempt. What is it?

The typical jackpot is around £8 million, and you bet £1. On the simplifying assumption that at most one person wins the jackpot — which we'll shortly see is not the case — your *expected* winnings -- what you win on average if you make a very large number of bets -- are £$(\frac{8000000}{13983816} - 1)$, that is, a *loss* of 43p.

Of course there are the other prizes, but the expected 'winnings' of any player still amounts to a substantial loss — which is *why* the Lottery can pay the company that runs it, the government, charities, sport, and the arts, as well as the winners. In gambling terms, the lottery is a sucker bet.

One implication of this calculation is that the more you play, the greater your expected loss becomes. (That's what *makes* it a sucker bet.) Many people will try to sell you a 'system' that involves betting on lots of combinations of 6 numbers — say a fancy way to start with 8 numbers and guarantee that if at least five of those are selected then you are guaranteed to win the jackpot. (This is easy to do. For example, list all $\binom{8}{5}$ = 56 ways to choose five numbers from your eight. Complete each of these by adding in each of the 44 numbers not already chosen in those five. *You must win!*)

The snags are twofold. First, you have to make 56×44 = 2464 bets. (Your expected *loss* is now 2464 × 43 p = £1059.52.)[1]

But what purveyors of such systems try very hard to hide from you is that they only work provided you get those five numbers right from among your chosen eight. And that's still very unlikely. Indeed, its probability neatly balances out, so that your chances of fulfilling this condition are the same as your chances of winning the jackpot if you place 2464 different bets. The point is that *any* 2464 different bets have the same chance of

1 It's not *quite* that bad in reality — there are other prizes than the jackpot. The numbers are slightly more in your favour, but the principle is the same.

winning as the 2464 you select using the system. So the system doesn't offer much added value. (There is one tiny advantage: you only need look at your initial eight numbers to see if you've won. With 2464 bets to look through, you might miss a win — or wrongly think you've got one — by mistake.)

The Rational Gambler

This doesn't mean that it's not worth your while to play. The company, the government, charities, sport, and the arts, are effectively betting huge numbers of times and thus can expect to *receive* their expectations. The typical player, however, bets only a small sum every week, and for them the expectation gives an incorrect picture — just as the statement 'the average family has 2.4 children' fails to apply to any individual family. The pattern for most players is that they pour small sums down the drain — typically about £100 per year — hoping for numbers that, for them, never turn up. Occasionally they get three numbers right and win £10, but that's just a psychological sop to keep them interested, and it's more than cancelled out by their losses. What they buy is £100 worth of excitement, and a possibility — albeit remote — of instant riches.

They could achieve just about the same effect by dropping 50p down the drain every week and hoping that they have a lost aunt in Australia who is about to die and leave them millions. The probabilities are very similar.

The pattern for the very few who hit the jackpot is utterly different. What they buy is a total change of lifestyle.

Should a 'rational' person play? There are many considerations, among them 'do I really *want* a total change of lifestyle?'. But even if we focus solely on mathematical issues, probability theory offers guidelines, not proofs. There are many cases where it is rational to bet when your expectation is negative: the clearest example is life insurance. The mathematics of life insurance is very similar to that for the lottery. The insurance company makes huge numbers of bets with all of its clients, but each client's bet is small. The expectations are similarly weighted in the insurance company's favour. The main difference is that normally the players hope *not* to hit the jackpot. It is entirely rational to play the life insurance game, negative expectation notwithstanding: players pay an affordable premium in order to purchase financial security in the event of an unlikely — but possible — disaster.

Should a company insure against key personnel winning the jackpot and heading off for a lifetime of indulgence in warmer climes? The chance of being killed in a car crash in the UK is about 1 in 500,000 per week. This is about 28 times as likely as winning the lottery. So it is definitely irrational to insure against an employee winning the lottery if you aren't insuring against far more likely, and more prosaic, disasters. Advertising for the lottery emphasises that 'it could be YOU', but your chances of being mugged, or of contracting AIDS heterosexually, are a lot higher. However, in any given year you are about two hundred thousand times more likely to win the jackpot than to be hit by a meteorite.

How to win the lottery

Now we've finally reached the bit you came along for. A mathematician is going to tell you how to win a fortune on the lottery. Some amazing system based on a mixture of probability, chaos theory, and your great aunt's cat's initials.

Come on, get real. If I knew how to win, do you think I'd give the secret away for the price of a paperback? *The same goes for all the other people who try to sell you their systems.* Because the expectation is negative, most systems have the same fundamental flaw: they are just ways to lose your money faster. Buying ten thousand tickets raises your chance of winning the jackpot to one in 1398, but your average loss goes up to around £4300. You'd do better putting your £10,000 on a four-horse accumulator with each horse at 11 to 1.

It's true that if you can lay hands on really big sums, you can guarantee winning. For a mere £13,983,816 you can bet on every combination. You will usually win about £8 million — a neat way to throw away £6 million. If several other people get the same winning numbers as you, then you throw away a lot more. The one case where this calculation fails is when there is a roll-over, in which case the jackpot becomes about twice as big as usual. Then you might come out ahead — but only if nobody else gets the numbers right. This danger made itself apparent on 14 January 1995, when a roll-over jackpot of £16,293,830 was shared between 133 people, who chose the numbers 7, 17, 23, 32, 38, 42.

Statistically, only about four people should have chosen this sequence if everybody picked numbers at random, so there must be some curious effects of human psychology at work. For example, people often think that because the lottery machine draws sequences at random, they must always *look* random. A sequence like 1,2,3,4,5,6 doesn't look random, so it will never come up. Not true. On average it will come up once every 13,983,816 times — exactly the same odds as those on the sequence that *did* come up on 15 January 1995. Odds are about the potential, not about the actual; and a random *process* does not always produce a result that *looks* irregular and patternless. To compound the misunderstanding, most of us have the wrong intuition of what a random sequence looks like. If you ask somebody to write down a random sequence, they usually space the numbers apart too much. Randomly generated sequences are often 'clumpy', but we human beings don't expect them to be. Compare the nice, evenly but not *too* evenly spaced sequence 7, 17, 23, 32, 38, 42 that 133 people picked, with the clumpy sequence 26, 35, 38, 43, 47, 49 that won £18 million for the only person who chose it.

And here, finally, is a system that might just work, and your faith in mathematicians is justified. If you want to play the Lottery, the smart move is to make sure that if you do win, then you win a bundle. There are at least two ways to achieve this aim, and the best strategy is to combine them.

The first is *only bet when there is a roll-over.* Paradoxically, if everybody did that, the roll-over wouldn't be worth anything, because the company running the Lottery wouldn't receive any stake money. So you just have to hope most punters aren't as smart as you are — which, if you're reading this advice and following it, is of course true.

The second tactic was made dramatically clear when those 133 disappointed people won a jackpot of £16 million and got only a miserable £122,510 each. It is *never bet on numbers that anybody else is betting on.* It's not so easy to do this, because you have no idea what numbers the other 25 million punters are choosing, but you should avoid too many numbers under 31 because of all those birthdays that everybody falls back on when asked to choose numbers. Avoid 3, 7, 17, and other numbers people think are lucky. You could sprinkle in a few 13's: many people think they're unlucky, but those little plastic balls aren't superstitious. I personally recommend something *really* stupid, like 43, 44, 45, 46, 47, 48.

'But surely a sequence like that never comes up?' Not so. It comes up, on average, once every 13,983,816 times — once every half a million years, to put the chances in perspective. The sequence 43, 44, 45, 46, 47, 48 is exactly as probable as the sequence 7, 17, 23, 32, 38, 42 that came up on January 14th. It is exactly the same as the odds on the six numbers your sweaty little paws have just marked on your Lottery ticket as you handed over your hard-earned £1 coin.

GRESHAM COLLEGE

Policy & Objectives

An independently funded educational institution, Gresham College exists

- to continue the free public lectures which have been given for 400 years, and to reinterpret the 'new learning' of Sir Thomas Gresham's day in contemporary terms;
- to engage in study, teaching and research, particularly in those disciplines represented by the Gresham Professors;
- to foster academic consideration of contemporary problems;
- to challenge those who live or work in the City of London to engage in intellectual debate on those subjects in which the City has a proper concern; and to provide a window on the City for learned societies, both national and international.

Gresham College, Barnard's Inn Hall, Holborn, London EC1N 2HH
Tel: 020 7831 0575 Fax: 020 7831 5208
e-mail: enquiries@gresham.ac.uk