

[Please cite as:

Forsyth, R.S. (1987). Chapter 12: The Art of Being Lucky. In: Drapkin, T. & Forsyth, R.S. *The Punter's Revenge*. London: Chapman & Hall Ltd. 230-242.]

## 12

### The art of being lucky

In this chapter we conclude by drawing together the main themes of the book, and present some general guidelines on the scientific approach to gambling.

#### 12.1 The gambling industry

It is very important to 'know your enemy', as it were. Gamblers do not operate in a vacuum, closed off from society at large, even if they sometimes seem to inhabit a peculiar private world. They need someone to accept their wagers. Most of the time that means one of the large bookmaking companies. These firms are not in business for fun: they are out for a return on capital employed, in the classic capitalist tradition.

It has been said of British restaurants that we in this island get the catering we deserve. There is more than a garlic-clove of truth in this remark. In France, people are educated to expect high-quality cooking and would kick up a tremendous Gallic fuss at some of the pap that passes for food in this country. The British, on the other hand, know no better, and meekly lap up what is placed before them.

Much the same can be said of betting. On the whole, we get the bookmakers we deserve. If the public cannot tell an overround book (at 115%) from a complete rip-off (at 150% and upwards) what incentive do the bookmakers have for staying on the straight and narrow? The unfortunate truth is that most punters wouldn't know a fair bet if it came up and bit them on the ear. The Government can slap on fresh taxes, the betting shops can hike their margins, and still the great punting public rolls up to the counter with cash in hand. Apart from one or two outraged racing journalists, nobody seems to notice when the bookies go 'over the top'.

This suits the bookmakers very well. They do not really want smart alics who study the form seriously and know a value-for-money bet when they see one to monopolize the winnings, because if word leaked out that gambling required skill and judgement, it might discourage the vast majority who merely want a 'flutter',

In the bookmakers' ideal world, all gambling would aspire to the condition of bingo. The skill required by this game can be mastered by a normal ten-year-old child, and after that it is purely a matter of luck. People come along for the atmosphere and the organizers rake off a steady, predictable percentage of the takings. The beauty of bingo, as far as those who run it are concerned, is that anyone can win. You can look at a big winner and say, truthfully, to yourself "if it can happen to him (or her) it can happen to me". It certainly can; but the probability is that it won't.

The Government is quite happy to connive at this state of affairs. As long as the tax revenue is maintained, or increased, they have few qualms about soaking the punter. There are, of course, various legal checks on how bookmakers can operate, but the example of the football pools shows how flimsy such punter protection can be. The levy on football pools has risen and risen over the years till it now stands at 42.5% -- effectively making the pools companies into extensions of the tax collection service. But lo and behold! The public keeps on coming back for more punishment. So long

as the top prize remains enormous in relation to the average stake, people seem prepared to go on losing week after week. This is depressing not only for the professional gambler, but for the part-time punter who wants a fair contest of wits.

So the scientific gambler is faced on all sides with the encroachment of 'junk betting' just as the gourmet faces the incursion of 'junk food' eating houses which are displacing restaurants all over the country.

Remember that the major betting organizations are not, for the most part, run by men with the turf in their blood any longer. They are, like much of UK industry, parts of larger conglomerates. They see themselves as being in the 'leisure industry' which encompasses hotel chains, tourist agencies, cinemas, fast-food outlets and much else besides. What the sober-suited men at the top care about is a healthy balance sheet at the end of the financial year. The only kind of gambling that truly concerns them is the sort of speculation that goes on in the City.

But capitalism has its virtues as well as its vices, and if the majority of punters were well-enough educated to perceive a reasonable bet and bold enough to demand it, that demand would be met -- if not by the major chains then by competitors who would spring up to satisfy that market. We hope that this book will play a small part in that education process, helping to halt the slide from informed gambling to mindless punting that has taken place over the past twenty-odd years.

## 12.2 Money management

Assuming that skilled gambling survives the rising ride of pure-luck lotteries, the astute gambler will need to know the fundamentals of money management. This is a topic we have touched on already, in Chapter 6, when dealing with losing runs; but it deserves fuller treatment.

Successful money management demands realism, so the first step is to forget all your aspirations towards fabulous jackpots. No gambling system on earth can guarantee a steady stream of successful 66-1 outsiders. If that is what you are after, especially if you want them linked up in a fourfold accumulator, you are living in dreamland.

The proper attitude is the attitude shown by the bookmakers themselves in the way they conduct their business. They are seeking a return on their investment, and so should you.

If you are a real professional gambler, you are only reading this book for amusement, but if you are an amateur who nevertheless takes gambling seriously as an interesting and profitable sideline, you will have to make your plans just as if you were embarking on a part-time business venture.

To give an illustration: suppose you decide to develop a fixed-odds football forecasting system on your computer -- possibly building on the foundations laid in Chapters 8 and 9.

First you will have to cost your time. You will arrive at a time budget rather like the following.

Entering data	1 hour per week
Checking data	0.5 hours per week
(If you don't check it at once, you probably never will.)	
Running the program	0.5 hours per week
Filling in coupon	0.5 hours per week
Visiting betting shop etc.	0.5 hours per week
Further research and development	2 hours per week

You have just set aside five hours of your time every week, during a 39-week season. Two hours of that is for further development and exploration, but this is absolutely vital because you are aiming at a moving target. The world does not stand still. The football league has altered its structure and points-system in the past and will no doubt do so again in the future. Teams adopt and discard attacking and defensive fashions of play. The bookmakers monitor and adjust their own practices and rules. New taxes are imposed from time to time. In short, you must expect the unforeseen. Even a successful system will not go on delivering the goods for ever.

If you cost your time at £8 per hour, you are already clocking up £40 per week -- a substantial investment.

Next you must add in your other outgoings. You will need to keep abreast of the results, and that means a newspaper bill which will not leave much change from £4 a week. There is also the little matter of floppy disks and printer paper to take into account, even though we are writing off the cost of the computer against other activities. It is easy to neglect these things, but if you sit down and do your sums you will find that they do add up.

So you now know that you have to win at least £44 per week just to break even. This gives you a target to aim at.

To make £50 per week, which is a very modest profit indeed, you need a system that yields a 50% post-tax profit on turnover and an average weekly stake of £100. That is a pretty fantastic level of performance.

Alternatively you could make do with a more realistic aim of 20% post-tax profit on turnover provided you were prepared to stake £250 per week. You could even accept a more businesslike 10% profit margin with stakes averaging £500 a week.

Now you see the dilemma clearly. If you have £500 a week to spare, why bother with all this extra work on such a risky project? Why not just spend the money on living well? Or invest it in a secure commercial enterprise?

From a gambling point of view, £500 a week puts you in rather dubious company. You are not quite up with the high-rolling sheikhs, but you have left the casual punter a long way behind. On any reasonable reckoning you would have to be prepared for a losing run of 10 weeks at the minimum, so this entails a starting bank of £5000 or more.

You can attack this problem from several directions. The first and most common response is to devalue your own time. You say "well, it's my hobby, I'll only count my time at £2 an hour", or words to that effect. At a stroke you have reduced your break-even level to around £15 per week. But the exercise has been worthwhile, because you have been forced to look at your gambling activities as a serious business proposition -- and have recoiled in horror!

That's fine. It is up to you how you spend your free (or nearly free) time. But you must admit that you are not making money out of gambling in the same way as -- let us say -- IBM is making money out of selling computers. You are merely satisfying your pride by picking winners a bit more often than the next man.

The second way round the problem is to improve your forecasting program or system. We all try to do this all the time; but of course it requires a further investment of time, with no guarantee of a

pay-off. Research of this kind can easily swallow up a lot more than the 2 hours a week we originally bargained for.

A third solution is to give up the idea altogether. Once again, this is a perfectly rational response to the situation. If you think you can make more money by taking a part-time job, or get more satisfaction by staying longer in bed each morning, you are entitled to do so.

A fourth response is to cut costs, though there is not much scope for this. The newspapers are relatively cheap, and the main cost is your time. We assume you are doing everything at a good pace anyway, but you will get a little quicker with practice.

So you end up trying to do a little of everything:

- (1) You accept somewhat a lower margin on turnover;
- (2) You work a bit faster;
- (3) You keep trying to improve your system,
- (4) You raise the stakes.

However, the fourth point -- increased stakes -- is the most dangerous. You only raise your stakes above pocket-money level after at least six weeks of low-level staking have convinced you that you are on to a winning proposition. And even then, you never put more than a tenth of your current balance (preferably only a twentieth) on a single bet.

It is even more important to limit your losses than to make a profit, because once your betting bank has gone you can never recoup your earlier investments. (And if you keep replenishing your betting bank from other sources of income, Heaven help you! That really is the road to ruin.) So although the figures given in this section will not apply exactly to your individual case, the general principles will.

It is strange how bad luck seems to dog the footsteps of those who ignore the discipline of sound money management.

### **12.3 Respect for the data**

We have considered gambling as a business. Now we will consider it as a science. All the difficulties described in the previous section are eased by having a better forecasting system, but you cannot expect magic.

What you can do is follow well-established scientific principles. You are engaged in an on-going research programme, and you should behave accordingly. The first thing to do is to look at how successful scientists behave. The second is to attempt to mimic them.

Let us take medical research as a case in point.

For many years the origins of Parkinson's disease (which affects the muscular co-ordination of sufferers and ultimately makes it impossible for them to walk or speak) have been a mystery. It was labelled an 'idiopathic disease', which is a Greek way of saying that the scientists were baffled. But good scientists do not like being baffled, and some of them were determined to tackle the problem.

An obvious step was to eliminate various potential causes. Was it caused by a bacterial infection? No one could find any evidence of contagion. Was it caused by a virus? Nobody could find any traces of

a common virus in the corpses of dead victims, nor any obvious signs of the immune system responding to viral attack in the bodies of live ones. Was it then, an inherited condition? This, until recently, was the most favoured hypothesis; but it was disposed of by an extensive study of twins, identical twins have the same genetic make-up, so if one inherited a predisposition to Parkinsonism, the other would as well. However, the proportion of identical twins of sufferers who also had the disease was not significantly greater than the proportion of fraternal (non-identical) twins of sufferers who also had the disease -- and very little different from the rate in the population at large.

Negative evidence is always disappointing, but it has its own value. In this case it set one team of researchers off on a hunt for environmental pollutants as a possible cause for the illness. Eventually, after much patient sifting of evidence and many dead-ends, they found a common factor: all the localities they had identified where Parkinsonism was unusually frequent were near paper-pulping plants. They had not (and still have not) found the cause, still less a treatment, but at last, after literally centuries of fruitless investigation, they were on the right track in the search for a way of dealing with this crippling disability.

Now gambling on horses or football is not such a worthy aim as a cure for Parkinson's disease, but if you are going to do it at all, you might as well do it properly. That entails treating the enterprise in a scientific spirit. It also means focusing your attention. No medical researcher tries to find the cause of every disease under the sun: they concentrate on one or a few related maladies.

Similarly, you should home in on a particular form of betting and specialize in that. Let us say, for the sake of example, that you decide to specialize in predicting away wins in football matches where the odds on offer are 3-1 against or greater. These are games that the home side should win, or at least draw, and so most of them will not be away wins. In fact, somewhat less than 20% will result in away wins. That is too few, but if you could get the proportion of your selections up to 25% you would be on to a bet that was profitable after tax (since average odds are bound to exceed 3-1); and at 30% you would be laughing. In other words, you only need a slight edge.

The first step is to shop around a little. On this particular score, a preliminary investigation has revealed that William Hill appears to offer slightly better odds against away wins in general than Mecca bookmakers (whereas Mecca tend to offer better odds with draws). If you were pursuing this line seriously you would want to compare the big four (Joe Coral, William Hill, Ladbroke and Mecca) for a few weeks to see which offers the best prospects,

Then you collect a dossier containing a large number of such matches, several hundred at least. Initially this need not be on the computer at all. In fact there are some advantages in using a large notebook or ledger in the early stages, since it forces you to get close to your data. The computer cannot do your thinking for you, though it can be very helpful in testing your hypotheses at a later stage.

Once you have a sufficient number of examples, you start looking for ways of discriminating the real banker homes from the more dubious ones which may end in aways and draws. All you need to do is push the rate of home wins down from around 55% to under 50% in your chosen sub-group and the job is virtually done.

So you look for various tell-tale signs. You suspect that if the home team has just lost its last match, it is not quite so good as the bookies think. Accordingly you classify the data into two groups, those where the home team lost its last game and those where it did not. But the proportion of aways turns out to be only marginally higher in the former group. Don't give up yet: this is only the beginning. Try the away team's latest result. Then try them in combination: perhaps if the home

team lost and the away side won, that may be a good sign. You find (let us suppose) that it is, but it is too rare to give enough bets each week. (These matches are by definition between a good home team and a poor away team.)

Never mind, you may well need several tests in the filtering process, and there are plenty of other variables to look at -- league standing, FRAN rating and goal difference among others. This is exactly where the computer comes in; you supply the hypotheses and the machine runs over the data to check them. Most of them will be wrong, but that's how it is in real life. You only need one or two hits among scores of misses and you have the basis for a winning strategy.

What this brief outline illustrates is the classic scientific method of conjecture plus refutation. You think up all kinds of stories to explain the facts and then try to refute them. The ones that stand up are more likely to be true.

The conjectural side poses very few problems. Most people can think up a dozen fanciful theories before breakfast. The refutation is where we all tend to fall down. It is fatally easy to keep elaborating a theory to fit the facts (see Chapter 10).

With a small sample you may find that if the home team lost its last game, it is a good predictor of an away win. As your sample increases, the rule becomes less effective, but you are already emotionally committed to your theory, so you look for a way to patch it up. Eventually you decide that if the home side lost and the away side won or drew that is good enough. Then more data arrives that does not quite fit the pattern. Rather than starting again, you modify your rule with a couple of escape clauses.

You end up with a rule like: if the home side lost by more than one goal and the away side drew and there is an R in the month or the away side won and the home side lost or drew or the away team's name begins with a vowel or the home team scored an odd number of goals in its last two games, then forecast an away win. Your rule will probably need repair every week.

Such over-elaboration is not unknown in the history of science, despite the long-established principle called Occam's Razor which states that entities should not be postulated unnecessarily and which is a bias towards parsimony. But if you indulge in these baroque explanations, you will not make any money.

Knowing when to abandon a pet theory, then, is as important as being able to construct theories in the first place. This is where respect for the data comes in.

You try to keep your data accurate; you try to pull in a large and representative sample; and you listen to what it is trying to tell you, rather than imposing your own ideas on it. Respect for the data means, among other things, that you do not look at a 2-2 result which your system picked as an away win and say: what a near miss. Nor do you count a horse beaten by a short head into second place as 'nearly a winner' just because your program selected it. If you have a fair-sized sample these marginal cases will cancel themselves out, so you must accept even the borderline decisions of the ultimate referee, reality, with a good grace -- however much they go against your preconceptions.

You do not need to be the world's greatest statistician, but you do have to have a healthy respect for the way things pan out on the day. Abandon that, and you are back in the dark ages of witchcraft and superstition.

You also need the perseverance to keep worrying away at the data like a dog with a bone in its mouth until the treasure falls into your lap.

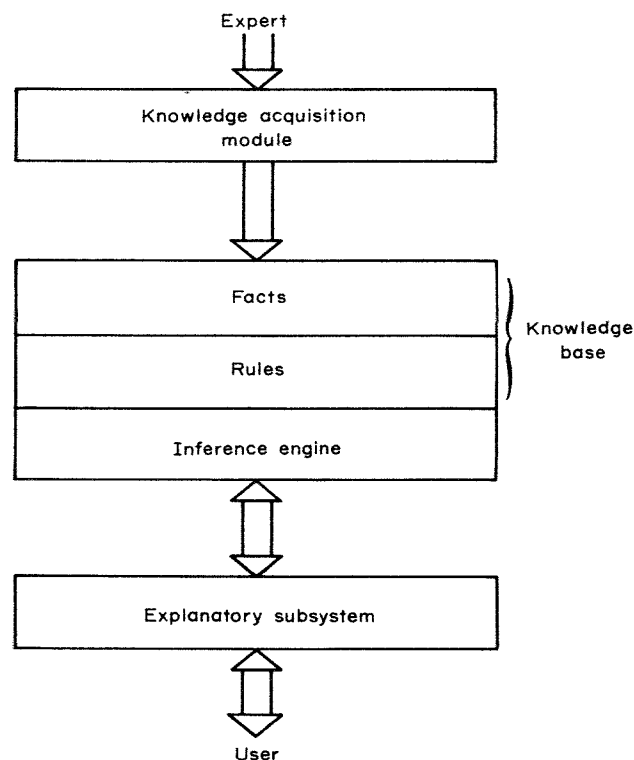
#### 12.4 Artificial Intelligence and animal cunning

Speaking of dogs brings us on to the subject of Artificial Intelligence (AI). One of the most interesting developments to come out of AI in the past decade is the Expert System. As we saw in earlier chapters, an expert system shares many of the characteristics of a human expert -- including the ability to handle uncertain inference in ill-defined situations, by means of Fuzzy Logic or probabilistic reasoning.

Figure 12.1 sketches the component parts of a typical expert system.

When an expert system is being developed for use in prospecting for precious ores or diagnosing illness, the first thing the designers do is interview a human specialist who can do the job with a high degree of skill. The expert's knowledge is then codified and refined for use by computer. Why not do the same with, say, horseracing? There are an awfully large number of inexpert 'mug' punters; but there are also a few genuine professional gamblers who know the form and profit from it.

So far nobody has seriously tried this approach (or if they have it has gone unreported!). But there have been a few promising excursions into the area which suggests that the application of AI techniques to gambling could pay off rather handsomely, certainly for those who get in early. We consider one of them here.



**Figure 12.1** Expert Systems. The typical expert system is really just a way of getting expertise from a human specialist to a user. The computer acts as a delivery vehicle so that the expert's skill can be disseminated more widely than would otherwise be possible.

#### 12.4.1 A BEAGLE on the racetrack

A pilot study, with an evolutionary flavour, was carried out by the present authors. Here we used the BEAGLE learning system (Bionic Evolutionary Algorithm Generating Logical Expressions) on a sample of horserace data from the summer of 1982.

This data covered 159 horses from 51 all-age handicaps with ten runners or less. Only the top three horses in the betting forecast of each race were considered. The database was split into two parts, 99 records were used as a training set (for forming the rules) and the remaining 60 were used as test data, to see if the rules generalized to fresh examples. (This is good statistical practice.) Each horse was measured on 18 variables, including the following.

SPEED	weight-adjusted speed figure for the horse
RATING	whether highest, next highest speed figure etc.
FC	position in betting forecast
LAST	place in last outing
LASTONE	place in last but one outing
LASTTWO	place in last outing but two
DAYS	days since last race
PEN	weight penalty, if any, added for a recent win
SPOTFORM	Daily Mirror Spotform in rank order
BTN	distance beaten in best recent performance
GOING	1 if the horse had won on the prevailing going in the past two seasons, otherwise 0
WT	weight carried
DIST	distance of the race in furlongs
SF	stable fancies (according to the Daily Mirror) coded as: 1 -- strongly fancied (SF) 2 -- fancied (F) 3 -- each way chance (EW) 6 -- no fancy (blank)

In addition, variable WIN was used to record whether the horse won.

The BEAGLE program uses an evolutionary learning scheme like that outlined in section 11.2.4. When given the horseracing data it came up with rules such as (SPEED > 60) and (WT > (SPEED \* 2.18)) which between them help to distinguish likely winners from losers. Its best rule set contained four rules (shown below). These four rules are used jointly to give a 'signature' or 'fingerprint'. That is to say, if rules 1 and 3 were true and rules 2 and 4 false the signature would be the binary number 1010 or 10 decimal. This index 10 points to position 10 in a table where information about that particular configuration of rule values is accumulated.

Applied to unseen data by the LEAF module (Logical Evaluator And Forecaster) the BEAGLE rules were correct 73% of the time. Of course, you can be right 72% of the time simply by saying No, since most horses lose, so this is not at all exciting. But an examination of the printout demonstrates that the rules were acting as an effective filtering system. BEAGLE only predicted one winner from the 60: it duly obliged. This is a good sign because selectivity is the essence of scientific betting. You cannot expect to predict every race; the important thing, as we have stressed, is to wait till you find a 'racehorse among donkeys' and then strike. Moreover the program's top 11 forecasts contained five winners. In other words the horses in the 0.3175 probability group did rather better than expected.

**Listing 12.1** BEAGLE output and rules.



LEAF report, sorted on estimated values of:  
(WIN = 1.00)

Rank	Pos.	Actual	Estimate		NAME	WIN	(groupsize)
1	10	1.00	0.7407	+	BALATINA	1	2 ??
2	25	1.00	0.4444	+	PAULAGER	1	4 ??
3	40	0.00	0.4444	+	F-SOUND	0	4 ??
4	54	0.00	0.4444	+	BRIGHT EDEN	0	4 ??
5	16	0.00	0.3175	+	REKAL	0	6 ??
6	22	1.00	0.3175		BIGPAL	1	6 ??
7	26	0.00	0.3175	+	RING-BID	0	6 ??
8	31	0.00	0.3175	+	ABERFIELD	0	6 ??
9	49	0.00	0.3175	+	DENMORE	0	6 ??
10	50	1.00	0.3175		L-WIMPY2	1	6 ??
11	52	1.00	0.3175		STEEL-C	1	6 ??
12	1	0.00	0.2675	+	DURANDAL	0	26
13	2	1.00	0.2675		PONTINLAD	1	26
14	6	1.00	0.2675		LITTLE-M	1	26
15	9	1.00	0.2675		GLIDE-P	1	26
16	11	0.00	0.2675	+	OVER-T-R	0	26
17	17	0.00	0.2675	+	T-MORGAN	0	26
18	19	0.00	0.2675	+	BRAC	0	26
19	21	0.00	0.2675	+	DENMORE	0	26
20	24	0.00	0.2675	+	WELL-M-A	0	26
21	32	1.00	0.2675		VALENTIN	1	26
22	33	0.00	0.2675	+	LULAV	0	26
23	48	1.00	0.2675		OUR-B	1	26
24	53	0.00	0.2675	+	P-LAD	0	26
25	13	1.00	0.2222		ILL SEE YOU	1	0 ??
26	20	0.00	0.2222	+	A-LUCKY	0	0 ??
27	28	1.00	0.2222		HIGH-O-T	1	0 ??
28	37	1.00	0.2222		D-M-BOY	1	0 ??
29	59	1.00	0.2222		H-O-TIME	1	0 ??
30	15	0.00	0.1528	+	TOPORI	0	7 ??
31	34	0.00	0.1528	+	IMPLICATION	0	7 ??
32	42	0.00	0.1528	+	M-CURE	0	7 ??
33	45	0.00	0.1528	+	DLO	0	7 ??
34	57	1.00	0.1528		L-WIMPY1	1	7 ??
35	58	0.00	0.1528	+	B-BOY	0	7 ??
36	35	1.00	0.1111		L-ROCK	1	1 ??
37	3	0.00	0.0732	+	EPSOM IMP	0	43
38	4	0.00	0.0732	+	HAWALI	0	43
39	8	0.00	0.0732	+	BADAYOUN	0	43
40	12	0.00	0.0732	+	SANDRAS-S	0	43
41	14	0.00	0.0732	+	SILVER	0	43
42	18	0.00	0.0732	+	S-PALM	0	43
43	23	0.00	0.0732	+	APERITIVO	0	43
44	27	0.00	0.0732	+	HAVOC	0	43
45	29	0.00	0.0732	+	S-GLORY	0	43
46	30	0.00	0.0732	+	E-COURT	0	43
47	36	0.00	0.0732	+	SYMPA	0	43
48	38	0.00	0.0732	+	REDDEN	0	43
49	39	0.00	0.0732	+	J-WISHING	0	43
50	41	1.00	0.0732		MANILOW	1	43
51	43	0.00	0.0732	+	ALDENHAM	0	43
52	44	0.00	0.0732	+	A-DICKENS	0	43
53	46	0.00	0.0732	+	HAMPSHIRE	0	43
54	47	0.00	0.0732	+	D-SHARP	0	43
55	51	0.00	0.0732	+	AMOROUS	0	43
56	55	0.00	0.0732	+	H-COOL	0	43
57	56	0.00	0.0732	+	MINMAX	0	43
58	60	0.00	0.0732	+	FRANCESCO	0	43
59	5	0.00	0.0555	+	MOSES-S	0	3 ??
60	7	0.00	0.0555	+	SVALBARD	0	3 ??

Crude success rate: 73.3333%

Mean target value = 0.2833

Success rate in non-queried groups = 82.8571% [ 35 cases].  
Av. target value for all-YES group = 1.0000 [ 1 cases].  
Av. target value for all-NO group = 0.0455 [ 22 cases].

Data-file was : gg02.dat  
Rule-file was : gg02.rrr

Rule-set being used (with logical target expression):

```
( WIN = 1.00) $  
  
( SF <= ( PEN > RATING))  
$      5      2      17      75  
( SPEED >= 64.00)  
$      17      27      6      50  
(( LAST <= DIST) = LAST)  
$      10      12      12      65  
( SPOTFORM = ( 0.00 >= BTN))  
$      6      7      16      70  
  
( ) $  
      0.2222      99  
0000      3.0000      43  
0001      0.0000      3  
0010      1.0000      7  
0011      2.0000      2  
0100      7.0000      26  
0101      2.0000      4  
0110      2.0000      6  
0111      0.0000      1  
1000      0.0000      0  
1001      0.0000      1  
1010      0.0000      0  
1011      0.0000      0  
1100      0.0000      0  
1101      0.0000      0  
1110      3.0000      4  
1111      2.0000      2
```

Looking at this from the other angle the number of winners out of the bottom 31 was only three. The category in which all the rules were against a win, the 'no-hopers', contained only one mistake out of 22, a horse called Manilow -- and who wants to win money backing a horse called Manilow? So the rules are good at filtering out the rubbish. This in fact is their most valuable function. (N.B. The question marks in the forecast listing are put there by LEAF to warn about predictions made on the basis of small samples.)

All in all, therefore, there is evidence that a learning technique developed for AI research can pay its way on the racecourse. This is only a preliminary study, but the results are highly encouraging. (Note that the rules given in Chapter 7 are founded on a far larger sample and are hence more reliable.)

There is a wealth of techniques in the AI literature just waiting to be exploited. There are learning algorithms, pattern recognition methods, schemes for handling uncertain inference and much more besides. We have only scratched the surface here. Who knows what riches wait in store for the gambler well-informed enough to put them into practice?

Science is all about prediction; so is gambling. Yet the typical gambler works in a pre-scientific mode. The application of scientific principles (especially from AI) on a wide scale would transform the betting industry. But we can be sure of one thing: whoever gains as a result, the bookies won't be the losers!