USING THE WISDOM OF THE CROWD TO FIND VALUE IN A FOOTBALL MATCH BETTING MARKET

Prelude to the Updated Version

This publication represents an updated version of my 'Wisdom of the Crowd' first published on Football-Data in August 2015 following its inclusion within my latest book *Squares & Sharps, Suckers & Sharks: The Science, Psychology & Philosophy of Gambling*. In addition to reproducing some of the original material and rationale behind this betting system, I have also presented two new interesting models for removing a bookmaker's margin in an attempt to estimate the 'true' betting odds for a football match. Furthermore, I have also investigated whether some bookmakers' odds are better suited to carrying out this task than others. Finally, I have updated the analyses with the most recent football results and odds data (through to May 2017) and retested the methodology for profitability.

Introduction

At my local fair recently there was a competition to guess the number of flower petals stuffed into a box. Rather than bother to take a close look, I simply asked the stall holder for the list of previous guess. So far there had been 40 guesses from which I calculated the average: 245. I submitted that figure as my guess. It turned out the correct answer was 295, meaning I was out by 17%, not too bad. Additionally, following my attempt there were a further 13 guesses, a total of 54, with an average of 272, just 8% lower than the correct answer. Seemingly, 54 completely unconnected people acting independently of one another collectively produced an answer that was a pretty close estimate of the true one. When you consider that the range of guesses was from 53 to 866, perhaps you'll agree that something magical is going on here. Not only was the final average pretty close to the actual number, it was also more accurate than the vast majority of individual guesses. That magic is called the 'Wisdom of the Crowd.'

The Weight of an Ox

The wisdom of the crowd phenomenon was first observed in the early 20th century by the eminent anthropologist, Sir Francis Galton. At a 1906 country fair in Plymouth, 787 people participated in a contest to estimate the weight of butchered ox. Galton calculated that the median guess to be 1207 pounds, a figure accurate to within 1% of the true weight of 1198 pounds and again more accurate than the majority of individual estimates. The wisdom of the crowd is a very real and repeatedly observable phenomenon in life, not least in the world of betting markets that are dominated by player psychology. Remove the influence of the bookmaker's margin or betting exchange's commission and we find that betting markets actually do a phenomenal job of replicating the 'true' probabilities of outcomes, despite not knowing *a priori* what the results of sporting events will be. [The word 'true' used in this context is placed in inverted commas to denote that we don't (nor can't) know the true probability of a result *a priori*, because sports, unlike casino games, are not the product of simple probabilistic functions. Rather, they can only be proved empirically *a posteriori* by analysing samples of results retrospectively.] This is perhaps best observed at betting exchanges. The chart below, based on 52,411 Betfair odds from worldwide football league matches during the period 29th October 2004 to 31st October 2005, compares the probabilities implied *a priori* by volume weighted average betting prices with the probabilities calculated *a posteriori* by the actual results. There is an almost perfect correlation (r = 0.995).

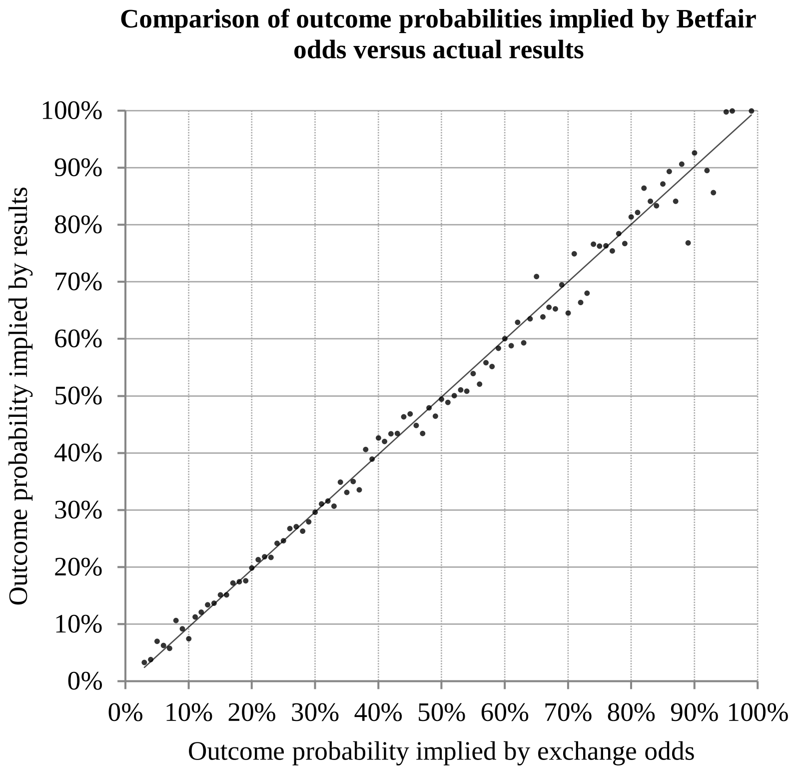

The Signal & the Noise

What is it about a betting crowd that makes their collective opinion so accurate, where individually so many can be inaccurate? Provided individual errors are not systematic and in the same direction they will tend to cancel each other out. Each individual guess has two components: signal (information) and noise (error). Remove the noise and what's left behind is the signal, that is to say the collective wisdom. Four conditions are necessary for a crowd to be wise in this manner: diversity, independence, decentralisation and aggregation. Arguably they are all present in betting markets.

Diverse Opinions

Having a diverse set of opinions is a prerequisite for collective wisdom. Where everyone is thinking or doing exactly the same thing, the probability of systematic error or bias increases. Diversity is the basis for any competitive market; let different ideas or products compete against one another and that's usually a recipe for the best ones succeeding. Google inherently understand the significance of diversity; it invests a lot of time, money and effort into lots of little start-up ideas like Google Earth, Google Glass and a driverless car. Not all of them will succeed but by having lots of eggs in your basket, you increase the chances that at least some of them hatch. In prediction markets like betting, diversity is virtually a given because of the environment of uncertainty and the typically large number of people acting in them with difference opinions, risk preferences and approaches to forecasting. When attempting to forecast the outcome of a game, for example, there are potentially limitless ways to skin that cat. Some prediction methods try to determine the intrinsic probability of outcome (for example value betting). Others adopt a more psychological approach, believing a market to be more a reflection of opinions (and more importantly opinions about opinions) with all their biases, making use of methods such a technical analysis to study trends and directions in betting prices. Then there are different types of prediction models: linear or nonlinear, static or dynamic, deterministic, probabilistic, or dynamic. Most will be wrong but the pooling of diverse ideas encourages collective accuracy.

Independent Thinking

Diversity arises out of independence of thought. Arguably, this is the most important ingredient for crowd wisdom. If everyone thinks the same way and does the same thing we frequently end up with poor outcomes. Everyone was betting on the 2015 UK General Election to return a hung parliament, because that's what all the polls and pundits were saying was going to happen. Evidently they were all doing the same thing and failing to take account of a couple of very important influences: the shy Tory effect and the lazy Labour voter. (Of course, it's easy to say this with hindsight.) A similar thing happened in the 2017 General Election. This time a Conservative Majority was predicted but it failed to materialise. Pollsters and pundits collectively underestimated the higher turnout amongst younger voters, who predominantly chose to vote with Labour. Perhaps if more polls and more pundits had demonstrated a greater independence of thought using what economists call 'private information', such ideas (and others) could have been used to collectively improve their predictions. Of course this isn't always the case. Teaching people to serve a tennis ball or to do differential calculus requires a narrow field of learning through repetition. But such activities are sufficiently predictable with clear relationships between cause and effect. In prediction markets under uncertainty, by contrast, learning through pattern recognition is limited because the patterns are largely random. What signal exists is deafened by noise, with good and bad luck dominating outcomes. In such environments having people acting independently helps to eliminate that noise, because it offers the best chance for keeping people's errors from becoming correlated. When mistakes are random they will cancel out.

An Invisible Hand

The final two pieces of the jigsaw that make the wisdom of the crowd such a powerful mechanism are decentralisation and aggregation. A system is said to be decentralised if it's not acting under the influence of a top-down central authority. By definition, independence and diversity of thought and decision making will be encouraged where central regulation is not restricting outputs. Decentralisation ensures that a crowd of self-interested, independent people working without top-down interference will collectively find a better solution than anything else you could come up with. The process happens as if my magic. It's the mechanism behind bird flocking, fish shoaling and insect swarming, the emergence of complex and seemingly coordinated behaviour out of a few simple rules followed by the self-interested individuals. In the case of birds there are just four: stay close to the middle; keep sufficient distance between neighbours; avoid collisions; and flee predatory attack. For human interactions, the 18th century economic Adam Smith labelled this magic the *'invisible hand,'* describing the unintended social benefits resulting from individual actions.

Public Property

Decentralisation, however, will only be of benefit if there exists a way of coordinating or aggregating together all the information. In a betting market that aggregation process is explicit: the conversion of private information and expression of opinions into a piece of public property – the odds. The betting price for a football team publically aggregates all the private information that exists. It represents the current balance of opinions about the likelihood of a team winning as expressed by the amounts of money wagered for and against it.

Market Equilibrium

The magic lies in the emergence of wisdom without individuals having a complete understanding of what the market is doing and without anyone knowing what the 'true' answer, if there is one, will be. People with only partial knowledge and limited calculating abilities actually arrive collectively at the right answer. A wonderful demonstration of this magic was accomplished by Nobel Prize-winning economist Vernon Lomax Smith. In 1956 he set out to determine whether people with limited information would confirm to the hypothesis of market clearing, where prices of traded assets adjust up or down such that quantity supplied at the market-clearing price equals the quantity demanded at the market-clearing price. Such a price is also called the equilibrium price. Giving his 22 students cards with a dollar price tag, he made half of them buyers and half of them sellers. The sellers were instructed not to sell at less than this price, whilst the buyers were instructed not to buy at more than their card value. A difference achieved between card value and actual contract price could be regarded as profit for the player. Strict anonymity was applied such that no one knew the value of anyone else's card. The students were then asked to start trading, calling out bids and offers which may, or may not, be accepted. Sellers and buyers were free to accept a bid or offer. If they were refused, further price compromise or bartering would be required until they were accepted. The successful trades were recorded publically on the classroom blackboard. Economic theory was matched by reality. Traded prices quickly converged on one price, the equilibrium price or what we might also call the expectation price, despite players being completely unaware of their competitors' demands and despite none of them preferring this outcome (self-interested traders after all want more profit). Collectively the convergence on the market-clearing price yielded the best possible outcome, even if some of the players had been blessed with additional knowledge telling them how they should trade. The brilliance of this experiment was that it demonstrated that for markets under uncertainty, imperfect people could collectively produce near-perfect outcomes. What allowed it to happen was a decentralised independence of action and the aggregation of privately anonymous information via the publication of a price.

Price Discovery

Essentially, price convergence is exactly what happens at a betting exchange like Betfair. This invisible hand is a kind of Bayesian process in which prices are continually and dynamically updated to reflect changes in information and opinion, which in turn are reflected in changes in supply and demand. At a betting exchange odds move simply in response to supply and demand. The market maker sits completely outside the contest, skimming his commission percentage from the action. This process is otherwise known as 'price discovery', a mechanism for determining the price of an asset in the marketplace through the interactions of buyers and sellers, or in this case backers and layers. Remarkable as it may seem, the betting public collectively 'knows' the 'true' probability of outcome of a sporting event through their betting actions. Odds shorten on the fancied competitors and lengthen on the least fancied, settling at values that reflect all the private information that has been consumed by the players. This price discovery is dynamic with the equilibrium price never completely stationary because there will always be new information arriving randomly on to the market.

For a bookmaker, things are a little different but only because they are part of the action; the fundamental process remains the same. Odds shorten because too much money has been bet on one outcome, giving the bookmaker a large liability in the event that it happens. Bookmakers are always looking to reduce their liability; in this case they can achieve this by shortening the odds to discourage further interest from customers. At the same time they lengthen the odds on the opposition to attract money. Through this Bayesian price clearing process they attempt to balance their book. If they get it right they won't care which team or player wins, and in effect they become more like an exchange. This won't happen all the time. Sometimes traditional bookmakers choose to take some sort of position on an event, either with a view to offering attractive prices better than the rest of the market or because they believe the rest of the market and their own customers to be wrong. Despite this, however, as a generalisation the punter should understand that, like for an exchange, he is not really betting against the bookmaker at all but against his some of his customers who collectively, it might be argued, form a wise crowd.

Knowing that betting odds are often wise can be put to use to make a profit. In the rest of this publication I'll show how this can be achieved specifically for the football home-draw-away match betting market.

The Bookmaker's Margin

The chief difference between a bookmaker and an exchange is the margin. The latter simply applies a commission to winning trades, but to all intents and purposes the published prices represent almost 'true' probabilities, at least collectively across a book of prices for all outcomes. In other words they collectively sum to 100% (or very close to it). The former, by contrast, shortens those prices, or rather increases those outcome probabilities. By offering odds that are less than fair, the bookmaker aims to build in his profit margin.

The traditional way of measuring the bookmaker's advantage is by means of the overround, known elsewhere as the 'vig' (short for Vigorish) or 'juice'. We can calculate the overround simply by summing the inverse of all the betting odds in a book. The simplest of books is one with just two outcomes: either player 1 wins or player 2 wins. Suppose Any Murray plays Novak Djokovic in the Wimbledon final, with both players priced at odds of 1.90. The inverse of these decimal odds is the implied outcome probability, with the bookmaker's advantage built in. In this case the implied probability for both players will be 1/1.90 = 0.526 or 52.6%. Hence, the overround, V, for this betting book will be given by:

$$V = \frac{1}{1.90} + \frac{1}{1.90} = 1.053$$

Typically, we express the overround as a percentage. To do this we simply multiply the decimal answer by 100%. In this case, then, the overround is 105.3%. Logically, the sum of all outcome probabilities in a contest should equal 100%, and indeed when the odds are fair this will be true. A sum of probabilities of 105.3% doesn't really make any sense mathematically, but is a reflection of how unfair odds have become, and how much advantage the bookmaker has given himself. In this example our bookmaker has built a 5.3% advantage into his prices. Overrounds for books with a larger number of outcomes are calculated in exactly the same way. For example, the overround for a standard football match bet with home win, draw and away win outcomes will be given by:

$$V = \frac{1}{H} + \frac{1}{D} + \frac{1}{A}$$

where H, D and A are the odds for the home win, draw and away win respectively. Subtracting 1 (or 100%) from the answer gives us a measure of the bookmaker's advantage, what we might call his margin, M. For example, the bookmaker Coral had odds for the 2017 Champions League Final between Juventus and Real Madrid of 2.88, 3.00 and 2.62. Plugging those numbers into the equation above gives V = 1.062 (or 106.2%), hence M = 0.062 (or 6.2%).

Removing the Margin

If we remove the bookmaker's margin we should theoretically get a rough idea of what the 'true' outcome probabilities of a match are, assuming the betting odds to be wise. The question is how should the margin be removed? That in turn requires an answer to another question: how did the bookmaker apply his margin in the first place? The simplest answer is that he applies his margin equally across all outcomes. For example, in Coral's book for the Champions League final, this would mean that the margin of 6.2% was applied equally to home, draw and away odds. Consequently, it is then a simple matter of multiplying the published odds by the margin to arrive at their implied 'true' odds, in this example 3.06, 3.19 and 2.78 respectively. For fairly evenly matched outcome probabilities as in this example, this equal margin distribution might very well be a reasonable assumption. But does it hold where one team is much more likely to win than the other?

It is now well-established that bookmakers apply differential margin weights to their odds, with greater weights the longer the odds. In other words if you purely bet randomly, you will lose less as a percentage betting on favourites than you will on longshots. This holds for a variety of sports, including football, tennis, basketball, snooker, darts and horse racing. It is called the favourite–longshot bias and arises because punters have a tendency to overbet long prices relative to short ones. It remains unclear whether this is because punters choose to behave in this manner or whether bookmakers are encouraging them to do so, but explanations for why it happens are beyond the scope of this document. For those wishing to learn more about this bias, and others in the world of sports betting, I refer you to my latest book: [*Squares & Sharps, Suckers & Sharks: The Science, Psychology & Philosophy of Gambling*]. Here I will now present three alternative methods for applying a differentially weighted margin, and the method for removing it.

1) Margin Weights Proportional to the Odds

The first method for applying (and removing) a differentially weighted margin is one that I published in the original version of this document (and in my book). It assumes that the margin applied to each outcome is proportional to the size of the betting odds (or inversely proportional to the outcome probability). Hence, for a book with n runners and overall profit margin M, the specific margin applied to the fair (or 'true') odds for the i^{th} runner (O_i) will be given by:

$$M_i = \frac{MO_i}{n}$$

For a home-draw-away football match betting market with 3 possible outcomes:

$$M_H = \frac{MH_f}{3}$$

$$M_D = \frac{MD_f}{3}$$

$$M_A = \frac{MA_f}{3}$$

where H_f, D_f and A_f are the fair home, draw and away odds respectively.

For example, a home-draw-away book with fair odds of 1.5, 5 and 7.5 and a margin (M) of 0.08 or 8%, the differential margins for home, draw and away would be 0.04, 0.133 and 0.200 respectively. To calculate the actual prices one then simply divides the fair price by the margin weight plus 1. For the home odds, for example, this is 1.5 ÷ 1.04 = 1.44. Similarly, the draw and away prices are 5 ÷ 1.133 = 4.42 and 7.5 ÷ 1.200 = 6.25. If a margin weight of 8% had been applied equally to home, draw and away prices, we would have 1.39, 4.63 and 6.94 respectively. You can see from this exercise that a differential weighting of odds in this manner shortens longshots more significantly than favourites.

With a little bit of algebraic rearranging, we can reverse the process to calculate what fair odds the bookmaker will have estimated in the first place, given his book margin and applying this model of differential margin weighting. Hence, for any published odds, O, the fair (or 'true') odds from which they came will be given by:

$$O_f = \frac{nO_{bookmaker}}{n - MO_{bookmaker}}$$

For a 3-outcome market like football match betting this will be:

$$O_f = \frac{3O}{3 - MO}$$

So in our example above, a published price of 6.25 will have fair (or 'true') odds of (3 x 6.25) ÷ (3 – (0.08 x 6.25)) = 7.50. I have made available a simple Excel odds **calculator** which will perform this task. Input the bookmaker's odds and calculator will remove the margins (columns H, I and J – 'Margin proportional to odds').

Clearly, as the odds get longer, the differential margin weight that must be applied to them gets larger. Ultimately, there will be a limit to how high those published odds will go, determined by the size of the bookmaker's book margin, M. The chart below illustrates the evolution of actual odds that might be offered by bookmakers with varying profit margin, M, for a 3-outcome book like home-draw-away. As the fair (or 'true') odds increase, so do the actual odds offered by a bookmaker but at ever decreasing rates. The theoretical limit to the highest odds a bookmaker will offer according to this model will actually be given by 3/M. For example, with a margin of 10% or 0.1, the highest odds you would see would be 30; for a margin of 0.05 it would be 60; for a margin of 0.02, 150; and so on. Evidently less generous bookmakers get close to their limits much more quickly. More generically, for a book with n runners, this model would imply maximum odds given by n/M.

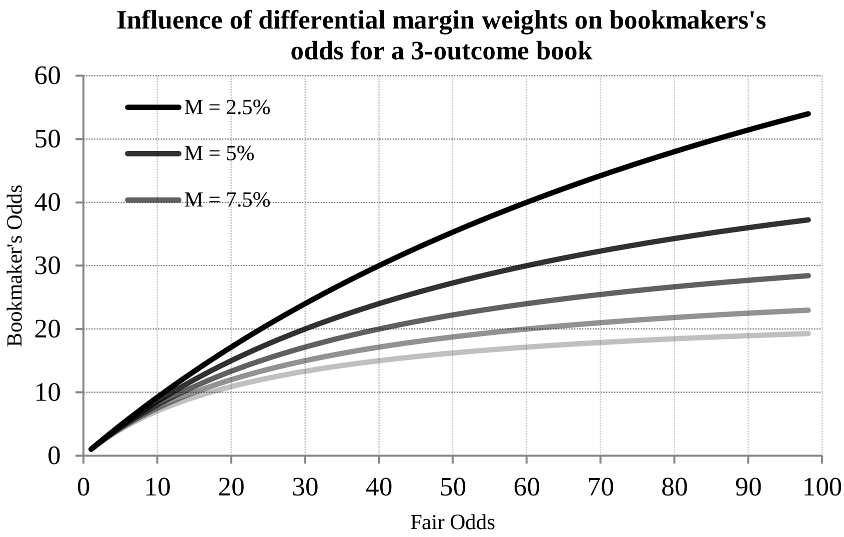

2) Using the Odds Ratio

In a blog for Sports Trading Network Keith Cheung describes how something called the Odds Ratio can be used to define the distribution of the bookmaker's margin. As Cheung explains, the odds ratio (OR) measures the relative probability of two events and is defined by:

$$OR(a,b) = \frac{O_a}{O_b} = \frac{p_a(1-p_b)}{p_b(1-p_a)}$$

where p_a and p_b are the probabilities of *a* and *b* occurring.

An OR(a,b) of 2 implies that outcome *a* is twice as likely as outcome *b*. To illustrate the usefulness of the odds ratio Cheung considers the following example.

> *"[W]hile we might instinctively want to say that 60% is twice as likely as 30% (because 60/30 = 2), there is no answer to the question of what is twice as likely as 60% using this definition. But using the odds ratio, we can say that 75% is twice as likely as 60%, because OR(75%, 60%) = 2."*

We can use the odds ratio to apply disproportionate margin weights to odds in a simple 2-outcome book via the assumption that the difference between the 'true' odds and the bookmaker's odds is equal for all odds in a book when defined by means of the odds ratio. So, again following Cheung, suppose that a 'true' outcome probability p maps to x after application of the bookmaker's margin. Similarly a 'true' probability q maps to y.

Since:

$$OR(x,p) = OR(y,q)$$

Therefore:

$$\frac{x(1-p)}{p(1-x)} = \frac{y(1-q)}{q(1-y)}$$

If we let c equal the odds ratio, rearranging we get:

$$x = \frac{cp}{1-p+cp}$$

and

$$y = \frac{cq}{1-q+cq}$$

Since p + q = 1,

$$y = \frac{c-cp}{p+c-cp}$$

Furthermore, since $x + y = M + 1$ (or the overround, V) it is possible, although not easy, to solve for c, the odds ratio, by means of the quadratic formula and finally obtain the probabilities x and y. Alternatively, it can be performed iteratively in Excel by means of either Goal Seek or Solver.

We can also reverse the process to calculate the 'true' odds if we already know the odds offered by the bookmaker. Now:

$$p = \frac{x}{c+x-cx}$$

and

$$q = \frac{y}{c+y-cy}$$

Whilst the process is the same for a 3-outcome book, the mathematics does become even more complicated. Because we've now introduced a third odds ratio OR(z,r), this time we will need a cubic formula to solve for it. A casual look at the formula would be enough to put anyone off trying. Furthermore, some of the solutions to the formula can involve complex numbers. Fortunately I managed to find a ready-made Excel worksheet that does the job and have adapted it for my own odds **calculator**. Again, simply input the bookmaker's odds and the calculator will compute the 'true' values (columns K, L and M – 'Odds Ratio').

Returning to our original fair book with odds 1.5, 5 and 7.5, applying an 8% margin via the odds ratio method would result in bookmaker's odds of 1.43, 4.42 and 6.55 (with the odds ratio equal to 1.17). Compared to my first methodology (1.44, 4.42 and 6.25), we can see that the differential margin weight applied to the favourite is slightly stronger resulting in slightly shorter odds relatively speaking, whilst the weight applied to the underdog is slightly weaker, resulting in relatively longer odds. In other words, the strength of the favourite–longshot bias applied to a fair book is slightly weaker for this methodology.

3) Using a Logarithmic Function

A third method for applying an overround uses a logarithmic function. Suppose, for a 3-outcome book, we have the 'true' outcome probabilities p, q and r, which sum to 1. Suppose we also have the bookmaker's probabilities x, y and z, which sum to V = M + 1. According to this methodology, $p \rightarrow x$, $q \rightarrow y$ and $r \rightarrow z$ in such a way that:

$$p^n + q^n + r^n = x + y + z = M + 1$$

In other words there is a common exponent, n, which will transform the 'true' probabilities to the bookmaker's probabilities given his margin M. The bigger the 'true' odds the more they will be shortened, in accordance with the favourite–longshot bias. n must obviously be < 1 since $p < x$, $q < y$ and $r < z$.

Hence:

$$Log_x p = Log_y q = Log_z r = n$$

Conversely, if we wish to remove an overround to transform the bookmaker's odds back to 'true' odds:

$$x^{\frac{1}{n}} + y^{\frac{1}{n}} + z^{\frac{1}{n}} = p + q + r = 1$$

Hence:

$$Log_p x = Log_q y = Log_r z = \frac{1}{n}$$

As for the 'odds ratio' methodology, these equations are best solved with Excel's Solver facility via means of iteration. Again my 'true' odds **calculator** will perform the task (columns N, O and P – 'Logarithmic function'). For our fair book with odds 1.5, 5 and 7.5, an 8% margin applied using the logarithmic method shortens odds to 1.45, 4.34 and 6.29 respectively.

Whilst there are differences between the three methodologies, it is clear they each give rise to a greater shortening of longshots in comparison to an equal application of the margin across all outcomes. But just how valid are these methodologies? Naturally the bookmakers won't tell us precisely how they apply their margin. One way to test their validity or usefulness is to analyse betting returns from these calculated 'true' odds. The more accurate those 'true' odds are, the closer to break even or zero our net profit should be (betting them to level stakes) assuming good and bad luck in results cancels out. Let's take a look at some data.

Testing the Margin Models

To test the validity of these four margin models, I have used a sample of Pinnacle's closing odds for the home-draw-away football match betting market purchased from the website Indatabet.com which includes games played worldwide from 11th August 2007 until 4th June 2016. A total of 136,876 games make up the sample, and consequently 410,628 odds values. This sample size should be sufficient to ensure that neither good nor bad luck in the results of games relative to those odds will predominate. That is to say, if we could genuinely know *a priori* what the true chances of a football game were, whether home, draw or away, our returns would be exactly zero. The table below summarises the theoretical returns from blind level stakes betting on all outcomes.

Model	**Profit over turnover (yield)**
Pinnacle's odds	-3.36%
Equal margin distribution	-0.56%
Margin proportional to odds	+0.1%
Odds ratio	-0.20%
Logarithmic function	+0.07%

All models for removing Pinnacle's margin do a reasonable job, but evidently some are better than others. Clearly an equal distribution of the margin falls short relative to the other three. Unsurprisingly, the likely explanation for this will be its failure to take into account the favourite–longshot bias. Relative to the other three models, a removal of the margin on longshots gives rise to 'true' prices which are arguably too short. Over a large sample of level stakes wagers, that shortfall will manifest itself as losses. Obviously favourites will be priced a little too long as a consequence giving rises to greater profits on those, but their influence will be disproportionately weaker and over a large sample the net result will be a sizeable negative deviation from zero as is witnessed.

The other three models all do a much better job of getting close to an expected 0% yield. Is there any way to tell them apart? The time series graph below hopefully provides an answer.

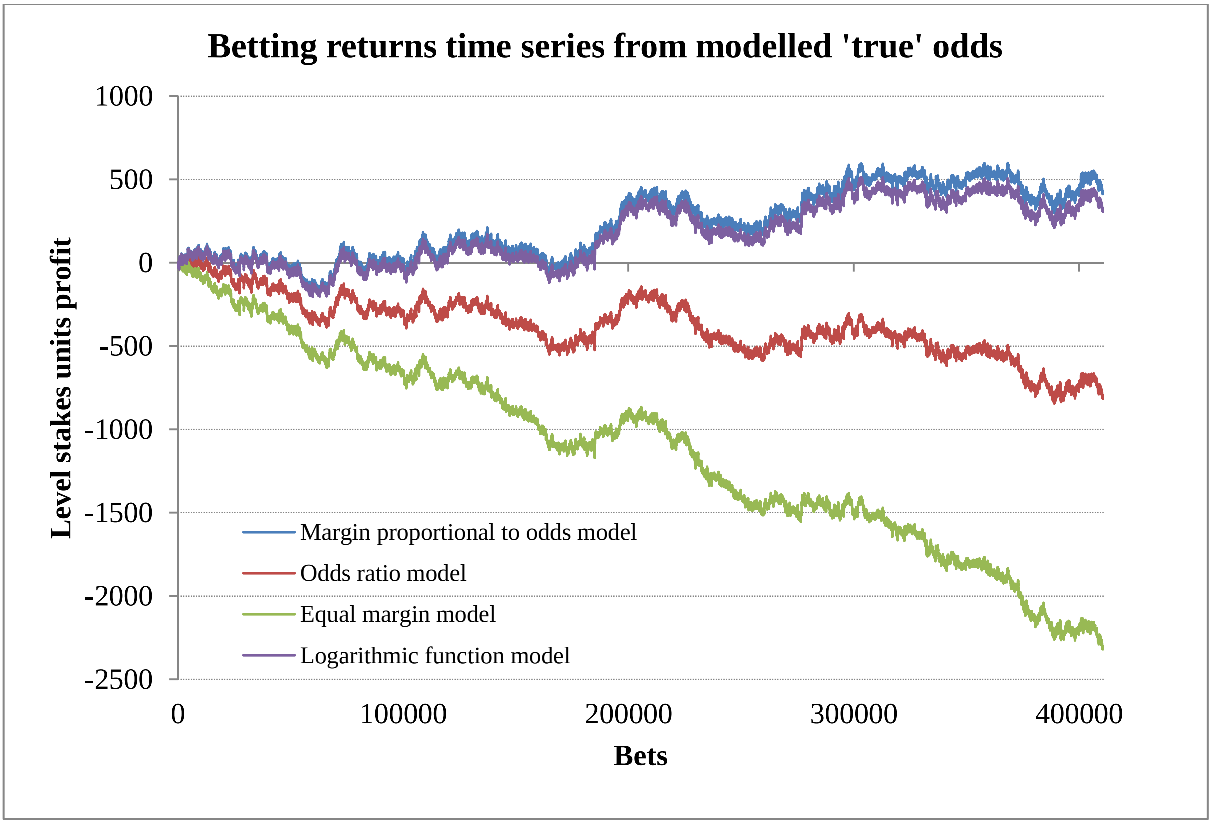

Again it is clear that the 'equal margin distribution' model fails to accurately define the 'true' odds. There is a strong negative trend over time in losses (green line) for the reason identified above. A similar albeit weaker negative trend is also evident for the odds ratio model (red line). Arguably, some of this might be down to random variance – perhaps even a sample size of over 400,000 betting odds is insufficient to eliminate it completely – but the most likely explanation is that this model also underestimates the strength of the favourite–longshot bias.

The 'margin proportional to odds' model (blue line) – my model used in the first version of this document and in my book *Squares & Sharps, Suckers & Sharks* – and the 'logarithmic function' model (purple line) yield almost identical time series. Yet in contrast to the 'odds ratio' model, they both show a slight positive trend towards profitability, implying a slight overestimation of the strength of the favourite–longshot bias. In other words, the 'true' odds they define for longshots are slightly too high. Nevertheless, given the tiny overall deviation from zero it is probably fair to say they both do a pretty good job of eliminating a margin in an attempt to guess what the bookmaker perceived the 'true' odds to be.

Using 'True' Odds to Find Value

An obvious question now arises: if we can identify what the 'true' probability of the result of a football match is, can we use this information to identify bookmakers' odds setting errors with a view to making a profit? Those of you familiar with odds comparison websites like Oddschecker, Oddsportal and Betbrain will know that bookmakers' prices are not all the same. If we have used one of them – in this example Pinnacle – to identify what these 'true' probabilities are, then other brands arguably might well be offering positive expectation, that is to say a chance to make a profit over the long term. For example, if after removing the margin we have 'true' odds of 2.00, a bookmaker with odds of 2.1 will be offering an expectation (or value) of 5% (or 2.10/2.00).

Using a sample of data collected through my own service Football-Data.co.uk, I have been able to analyse the returns of best market prices (as published by the odds comparison Betbrain.com) when those prices exceeded the calculated 'true' odds as estimated by my four models. The sample includes games played in 22 leading European football league divisions from 27th July 2012 to 17th May 2017, a total of 37,303 matches and 111,909 match betting odds. Whilst these odds were neither closing nor opening prices, they were collected at the same time as those from Pinnacle, allowing us to determine the availability of expected value. The four charts which follow show the time series of theoretical betting returns (with level staking) from best market prices where they exceeded the 'true' odds estimated by each of the four models.

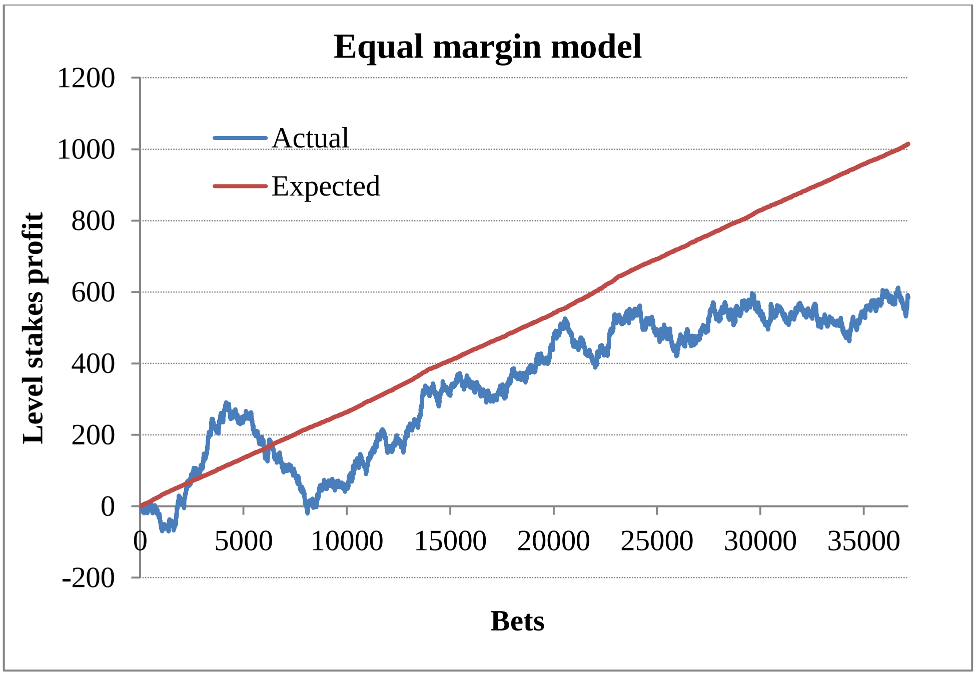

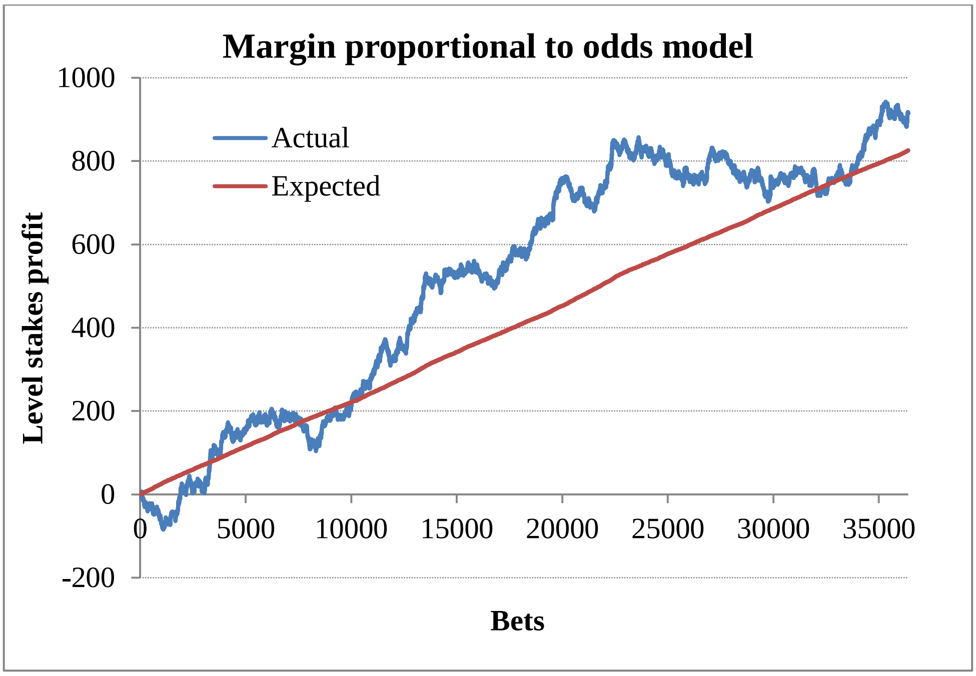
Margin proportional to odds model
Actual
Expected
Level stakes profit
1000
800
600
400
200
0
-200
0
5000
10000
15000
20000
25000
30000
35000
Bets

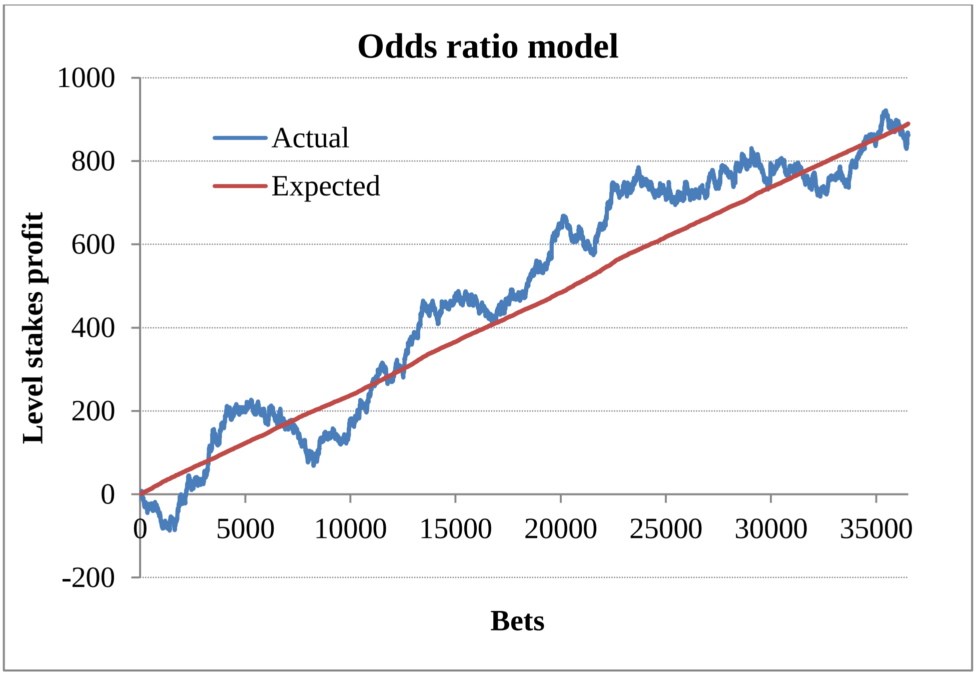
Odds ratio model
Actual
Expected
Level stakes profit
1000
800
600
400
200
0
-200
0
5000
10000
15000
20000
25000
30000
35000
Bets

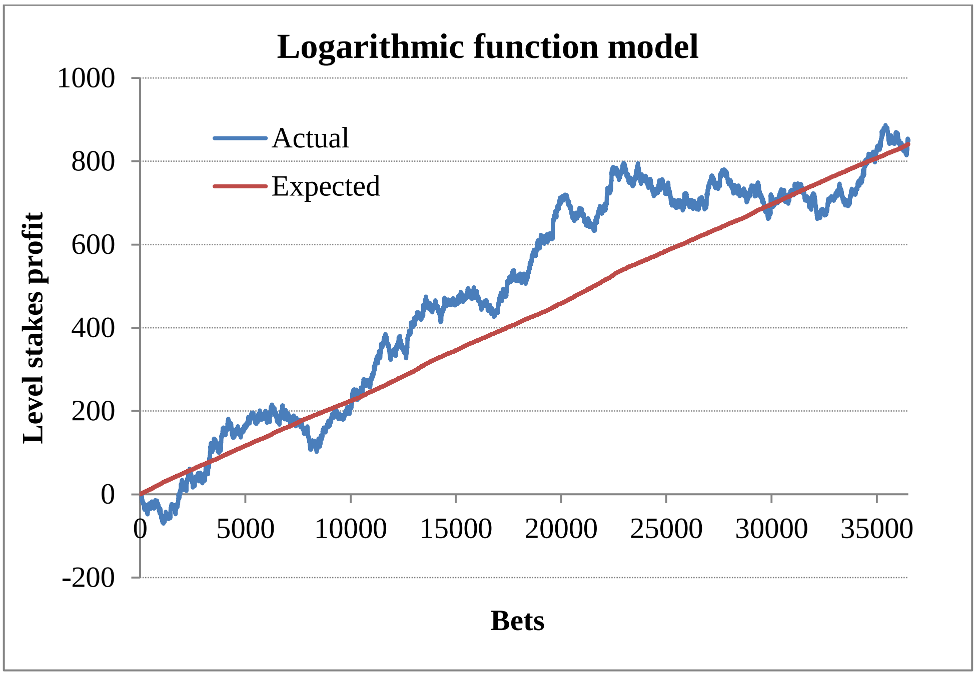

All four margin models allow us to find value, although clearly the 'equal margin' model underperforms relative to the other three for reasons I have already discussed. Essentially, because it fails to apply a large enough margin to longer prices, too many non-value bets will slip through the net. In contrast the 'margin proportional to odds' and 'logarithmic function' models, for sustained periods, slightly over perform relative to expectation for more or less the opposite reason.

The next table compares the performances of the four models for increasing value thresholds. The charts above plot the returns from all positive-value bets. The figures shown for the row 'Value > 2%', for example, show the returns for all bets where expected value is greater than 2%. Also shown are the number of bets, the expected value and average odds for each sample.

	Equal margin				**Margin proportional to odds**			
Value	Bets	Actual	Expected	Avg. Odds	Bets	Actual	Expected	Avg. Odds
>0%	37157	1.57%	2.73%	4.15	36395	2.51%	2.27%	3.31
>1%	24691	2.47%	3.87%	4.63	24057	3.59%	3.19%	3.49
>2%	16119	3.34%	5.16%	5.31	14967	5.71%	4.23%	3.82
>3%	10690	3.45%	6.52%	6.13	9065	8.30%	5.39%	4.32
>4%	7298	4.86%	7.94%	7.04	5519	12.05%	6.63%	4.95
>5%	5107	6.77%	9.43%	8.02	3400	11.78%	7.98%	5.74

	Odds ratio				**Logarithmic function**			
Value	Bets	Actual	Expected	Avg. Odds	Bets	Actual	Expected	Avg. Odds
>0%	36529	2.36%	2.44%	3.7	36496	2.32%	2.31%	3.36
>1%	24073	3.19%	3.45%	4.02	24239	3.26%	3.23%	3.57
>2%	15213	4.47%	4.61%	4.54	15169	5.36%	4.28%	3.94
>3%	9542	5.92%	5.88%	5.27	9193	7.22%	5.46%	4.49
>4%	6186	7.36%	7.20%	6.17	5656	11.15%	6.72%	5.2
>5%	4038	8.84%	8.67%	7.21	3508	13.45%	8.10%	6.11

Naturally, as we increase the value threshold this will reduce the number of betting opportunities available but will increase the expected yield. That actual returns also increase roughly in line with expected returns is an indication that using Pinnacle's implied 'true' odds to find betting value is a useful methodology. Average betting odds also increase with increasing value threshold since more higher-value opportunities exist at longer betting prices. This is simply because there is a greater variance in the ratio between best market odds and 'true' Pinnacle odds the longer they are.

A number of interesting comparisons can be made from the data in the table. Firstly, as the charts demonstrated before, failing to take into account the favourite–longshot bias when considering the bookmaker's margin leads to an underperforming model and underperforming returns relative to those expected. We can also see that average odds for the 'equal margin' model are higher than for other models. As previously noted, because the model's 'true' odds for longer prices are effectively too short (because the margin removed was too small), there are many more opportunities for best market prices to appear as though they offer value. For example, Sampdoria was a best price of 13 to beat Lazio in the Serie A match on 7th May 2017, with Pinnacle pricing them 12.5 at the same time. Assuming their 2.37% margin for this match was applied equally across home draw away would imply 'true' odds of 12.8. A price of 13 would thus represent a theoretical 1.6% value. Yet we have now confirmed this model is inappropriate as a method of applying and removing a margin. In fact it was far less likely this was a value bet; indeed the other models which assume a disproportionate weighting of margin towards longshots confirm this, with implied 'true' odds of 13.9, 13.3 and 13.8 for 'margin proportional to odds', 'odds ratio' and 'logarithmic function' models respectively. In fact in my sample of 111,909 odds there were 3,957 occasions when the 'no favourite–longshot bias' model suggested a value bet whilst the other three models did not.

We can also note that, as the time series charts illustrated, the 'margin proportional to odds' and 'logarithmic function' models show an over performance relative to expectation. Some of this may simply be the result of good fortune. The rest of it may be due to the models' overestimation of the strength of the favourite–longshot bias as we have previously observed. For example, when St Etienne played PSG on 14th May 2017, the 'logarithmic function' model considered the true price to be 9.57. Thus, the best market price of 10 offered 4.5% value. However, if this model applies a margin weight that it a little too strong to such a longshot, the 'true' value will be higher. Suppose, instead, the 'true' price was 9.34 as reflected by the 'odds ratio model'. The 'true' expected value would then be 7.1%. If this was in fact the correct model then we would see actual returns (over a large sample of similar bets) of 7.1% rather than the 4.5% predicted by our 'logarithmic function' model. Of course, the opposite will be true at shorter prices, but due to the greater variance for longshots their influence will dominate these statistics.

Taking all the information presented here into consideration it is not clear which model for applying (and removing) a bookmaker's margin offers the best means of replicating the real world of the favourite–longshot bias. Clearly the 'margin proportional to odds', 'odds ratio' and 'logarithmic function' models do a much better job than the simple 'equal margin' model which fails to consider the bias at all. Mathematically, the first of these is considerably easier to reproduce in Excel than the other two, although as previously mentioned I have made available a **calculator** for all models. Based on actual returns the 'odds ratio' model appears to fit expectation best of all, although it's impossible to determine if that is simply the result of chance. Perhaps the 'logarithmic function' model offers the best intuitive explanation, given that it accords with much standard economic theory of risk management and utility. It was Daniel Bernoulli in the 18th century with his rational choice theory who proposed that the subjective value of something to someone takes a logarithmic shaped utility function, a theory that has underpinned economics to this day. [*"[T]he utility resulting from any small increase in wealth will be inversely proportionate to the quantity of goods previously possessed."* For more on this, see my book.] Essentially, bettors are less sensitive to (in other words less bothered by) price changes the longer the odds, and this price sensitivity may very well adopt a logarithmic function. For example, assuming such logarithmic sensitivity a bettor would regard the difference between 2.00 and 2.14 to be equivalent to the difference between 10.00 and 12.59 ($2.00^{1.1} = 2.14$, $10.00^{1.1} = 12.59$).

No model appears to perfectly match reality as confirmed by the general positive or negative trend in returns to 'true' odds – if they were perfect there should be no trend at all discounting for the effects of chance.

Nevertheless, all three models appear to offer a reasonable means of removing the bookmaker's margin to determine the 'true' odds of a football match. They also offer a reliable method of finding value in the football match betting market far in excess of simply blindly betting best market prices. Those best market prices in my sample had an average overround of 100.6%. In other words, if we had blindly bet all possible outcomes with appropriate staking we would have lost about 0.6% on turnover. This method of value hunting is also superior to traditional arbitrage betting in the sense that it reveals far more betting opportunities. This is because we don't always need an underround book to have a value price. Typically about a third of books in my sample had some sort of value price available to bet on. Around 60% of these value opportunities were found in books that were still overround at best prices. Yet there is a final question we need to address before closing: can any bookmaker be used – here we have used Pinnacle – to estimate the 'true' price of a football result? In other words, are all bookmakers offering 'wise' prices on which this 'wisdom of the crowd' methodology must hinge?

Some Bookmakers are Wiser than Others

I began this publication by arguing that a bookmaker's price, like an exchange's, represents the public expression of a trading process between opposing points of view which ultimately 'discovers' the 'true' probability of an outcome. This process is known as the 'Wisdom of the Crowd'. But do all bookmakers really operate like this? Are they really so *laissez-faire* as to just sit back and let their customers battle amongst themselves, thereby arriving at collectively wise prices?

Suppose bookmaker 'A' offered 1.50 for Arsenal to beat Real Madrid at the Bernabéu. Suppose bookmaker 'B' offered 2.00. Would we really use bookmaker 'A' as a means of estimating the value offered by bookmaker 'B' to be 33%? Of course we wouldn't. Well, at least anyone who knew the slightest thing about the relative abilities of Arsenal and Real Madrid and the home advantage in football wouldn't. A less obvious example: bookmaker 'A' offers 10 for Leicester to beat Manchester City, whilst bookmaker 'B' offers 12. All other bookmakers offer 8. Does 12 represent value? Does 10? It's not inherently clear. To determine whether a bookmaker, on average, is offering something close to a 'true' price (discounting his margin) we need to analyse a much larger sample of betting odds. We already did this for the bookmaker Pinnacle in the earlier part of this document. It is pretty evident that Pinnacle's odds offer a reasonable measure of the 'true' price since the ratio of best market price to their implied 'true' price is a reasonable predictor of betting value (based on actual returns) for any pricing model which considers the favourite–longshot bias. Let's now repeat the exercise for another four brands: bet365, Betvictor, Ladbrokes and William Hill.

Using my original 'margin proportional to odds' model to estimate the true odds, the scatter plot below correlates the expected return (as defined by the ratio of odds from one of the four leading UK bookmakers and Pinnacle's implied 'true' odds) with the actual returns (to level stakes) betting those UK bookmaker odds (again using my sample of odds data collected through Football-Data.co.uk). Actual returns were calculated at 0.01 (1%) intervals in the expected return. A 5-point running average was then used to smooth the variance in the data. So for example the actual return for the point at 0.95 on the expected return axis is in fact the average of the actual returns for 0.93 through to 0.97 on the expected return axis. With such a large dataset many of these pairs are based on tens of thousands of odds, and can thus be considered statistically reliable. Naturally at the extremities of the expected returns (less than 0.8 and more than 1.1) the number of included odds diminishes substantially and variance in actual returns increases dramatically. I have restricted the range to that shown for visual clarity and to emphasise the conclusions that we can draw from this analysis.

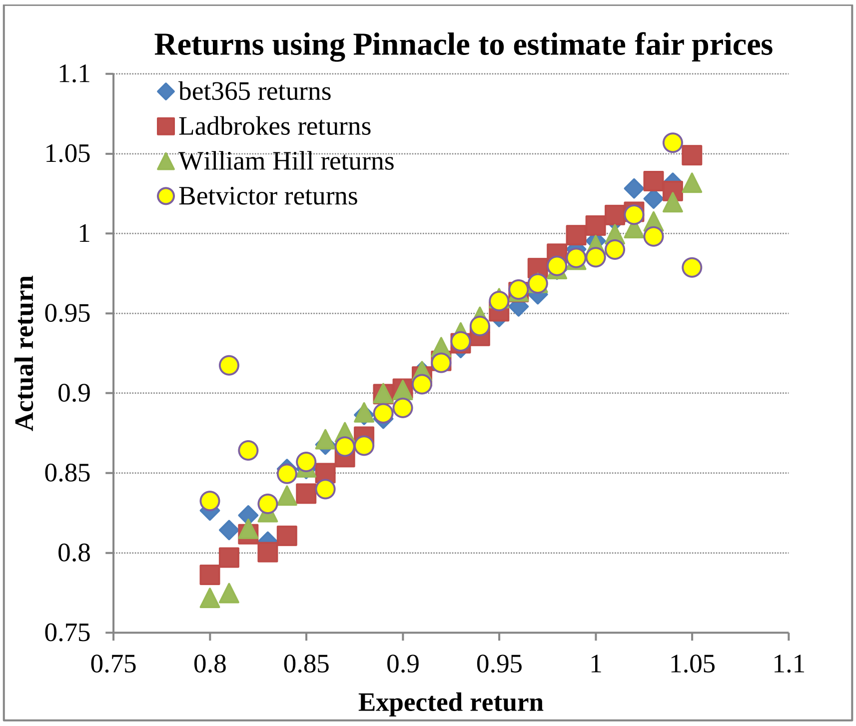

The correlation is strong and almost 1:1 for each of the four brands. Thus, where expected return is 0.9, we see about 0.9 for our actual return; where expected is 0.95, actual is 0.95; and so on. Although I've not shown them, the trend line gradients for all four bookmakers are essentially 1, and the statistical strength of the correlations (as measured by R-squared) is high. In other words, when a UK bookmaker offers 2.10 for a result and Pinnacle's implied true price is 2.00, we hold a 5% expectation. Should we bet such prices many times, our returns are likely to give us a 5% profit over turnover.

The next plot reverses the process. This time, the odds from the UK brands are used to estimate 'true' odds with actual returns calculated from Pinnacle's odds.

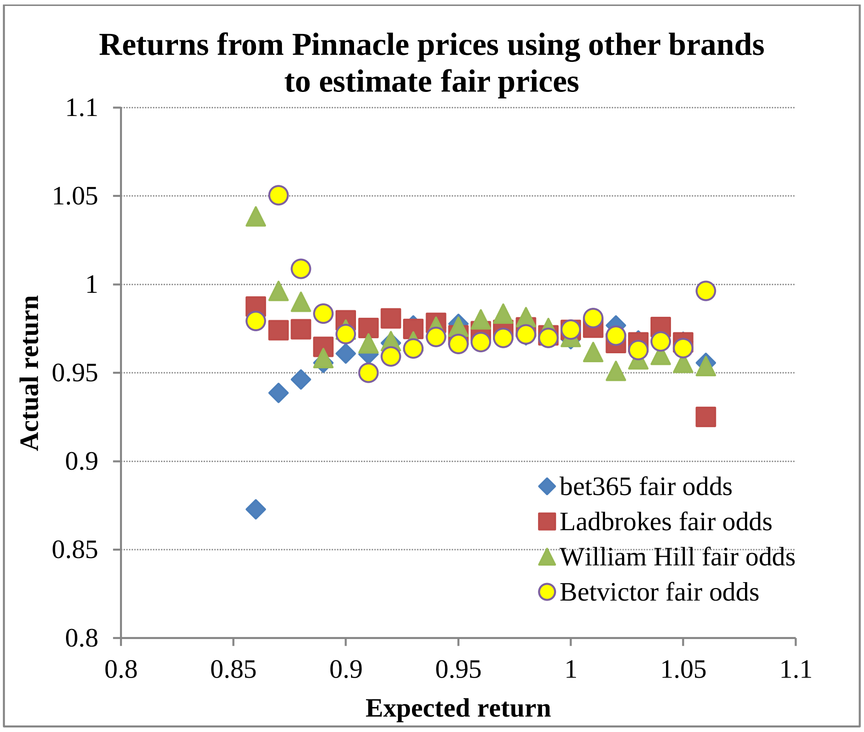

There is essentially no meaningful correlation between expected return (as predicted by the fair prices of the four UK bookmakers) and actual return betting Pinnacle's odds. Regardless of whether Pinnacle's odds are shorter or longer than the estimated 'true' prices, it seems your return will basically be the same: a loss roughly equivalent to Pinnacle's margin.

The conclusion here has to be this: relative to Pinnacle's home-draw-away match betting market, the four UK brands analysed here are not efficient and do not provide an accurate measure of 'true' outcome probabilities. So for example, when Ladbrokes offers a price of 2.00 implying a 'true' price of 2.13 whilst Pinnacle decide to offer 2.20, this will not mean that the punter holds an expectation of 1.033 (or 2.20/2.13). Rather, it will mean that the implied true odds are something like 2.25 and you can expect to make a loss over the long term betting these propositions with Pinnacle. Relative to Pinnacle, Ladbrokes and the other UK brands, on average, fail to offer any meaningful information about the 'true' odds of a football result.

There are two possible explanations to account for these findings. The first is that my model for calculating fair prices lacks reliability, and the greater the margin applied by the bookmaker, the greater the inaccuracy. Pinnacle, of course, have the smallest margins. The average in this sample of matches was 2.7%. This compares to 4.1%, 5.5%, 7.4% and 7.8% for Betvictor, bet365, William Hill and Ladbrokes respectively. Perhaps the 'margin proportional to odds' model I used here for estimating the 'true' prices implied the by the four UK bookmakers just doesn't work as well when margins are larger. However, given that theoretical level stakes returns from blind betting all 'true' odds in this 5-year sample are 99.88%, 100.13%, 99.80%, 99.59% and 100.00% for Pinnacle, bet365, Ladbrokes, William Hill and Betvictor respectively, this explanation doesn't seem credible. Furthermore, I have tested the correlation between expected and actual returns without removing the margin at all and broadly the same patterns exist. Finally, arbitrage bettors who use Pinnacle will know that it's typically the other bookmaker who sits on the value side from the bettor's perspective. It's obviously no surprise, then, that Pinnacle are the only bookmaker to openly welcome arbitrage bettors, knowing that, unlike the other brands, they are not, on average, offering positive expectation by doing so.

Alternatively, we might speculate that some bookmakers are intentionally shifting prices away from market efficiency in favour of pursuing the interests of their business models. Pinnacle's pricing model as much as possible utilises crowd wisdom and accepts sharp players to tighten its lines. Other bookmakers from Europe and the UK prefer to encourage a steady flow of squares via promotional offers, a wider variety of low-liquidity markets and the regular availability of best market prices (if not the lowest margins). With respect to the last of those, a casual perusal of any odds comparison will reveal numerous matches where bookmakers are significantly out of line with Pinnacle's market, and in the extreme offering loss-leading value to the player. So for example, Pinnacle offered a price of 1.22 for Barcelona to beat Espanyol on the 2nd January 2016 that was longer than Betvictor's implied 'true' price of 1.20. Yet, conceivably, Betvictor artificially shortened Barcelona's price relative to what they knew was fair with a view to being best-in-market for Espanyol, at 19. This price, and not Pinnacle's price for Barcelona, would then represent the value, based on Pinnacle's implied 'true' odds of 16.8.

Whilst it is pure conjecture that brands manipulate their markets like this, the regular restriction and banning of customers who regularly exploit best market loss-leading value would be a pattern consistent with such a conclusion. If bookmakers know they are intentionally offering loss-leading value to customers it's not going to prove very cost effective for them if they allow those customers to continually exploit such value while the other parts of the book are attracting little or no custom at all. Of course, Pinnacle, with a much more passive price management system that reflects customer turnover will not suffer the same problem. The consequence is that Pinnacle, relative to other brands, will offer more efficient and accurate markets as a reflection of 'true' outcome probabilities.

This is not to say that other brands' markets are completely inefficient, nor that Pinnacle is wholly efficient. Rather, this analysis is simply illustrating that Pinnacle is relatively more efficient than the rest. We can also compare the UK brands with each other, to establish some sort of efficiency hierarchy. The following four plots of actual returns against expectation reveal that Betvictor and bet365 are more efficient than Ladbrokes and William Hill, and therefore more useful at estimating 'true' prices (relatively speaking).

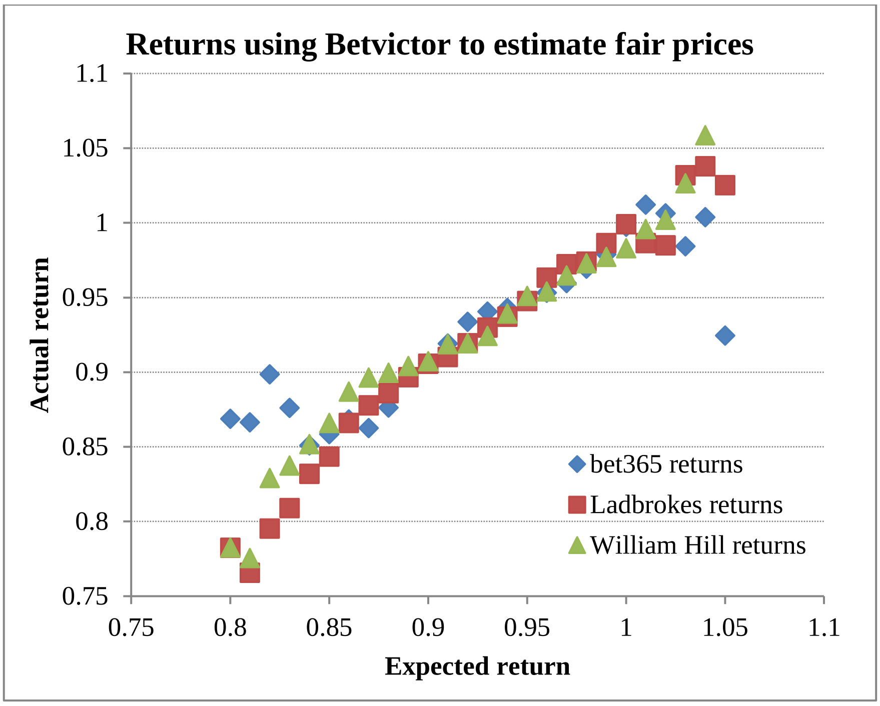

Returns using Betvictor to estimate fair prices
Actual return
Expected return
1.1
1.05
1
0.95
0.9
0.85
0.8
0.75
0.75
0.8
0.85
0.9
0.95
1
1.05
1.1
bet365 returns
Ladbrokes returns
William Hill returns

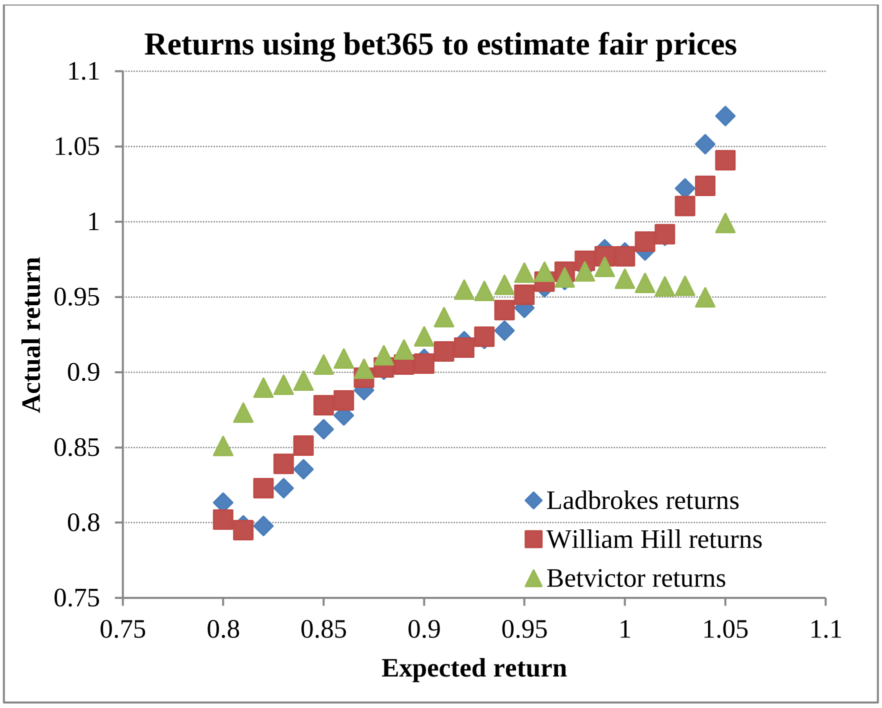

Returns using bet365 to estimate fair prices
Actual return
Expected return
1.1
1.05
1
0.95
0.9
0.85
0.8
0.75
0.75
0.8
0.85
0.9
0.95
1
1.05
1.1
Ladbrokes returns
William Hill returns
Betvictor returns

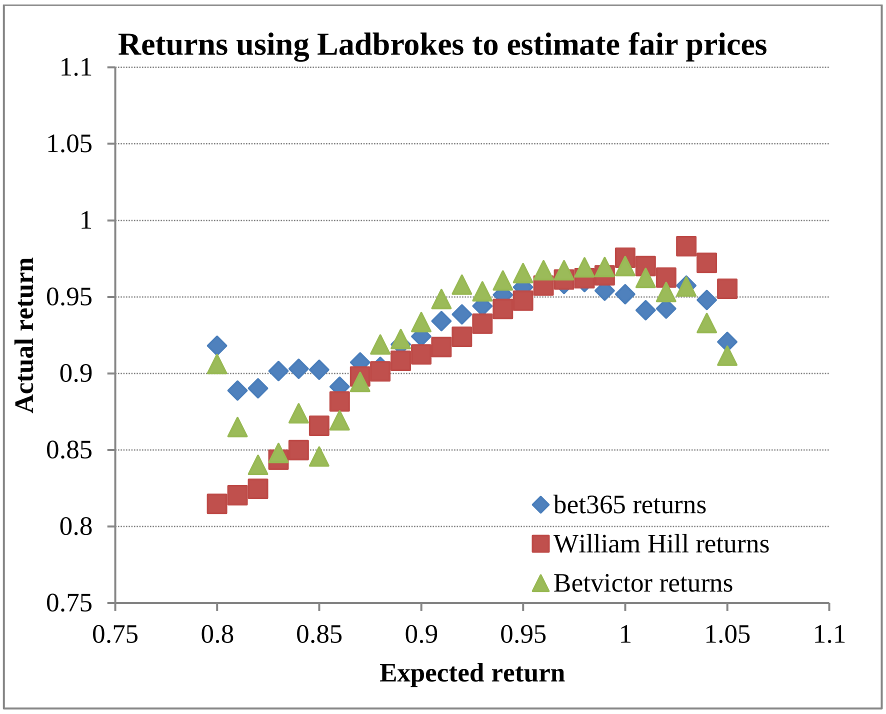

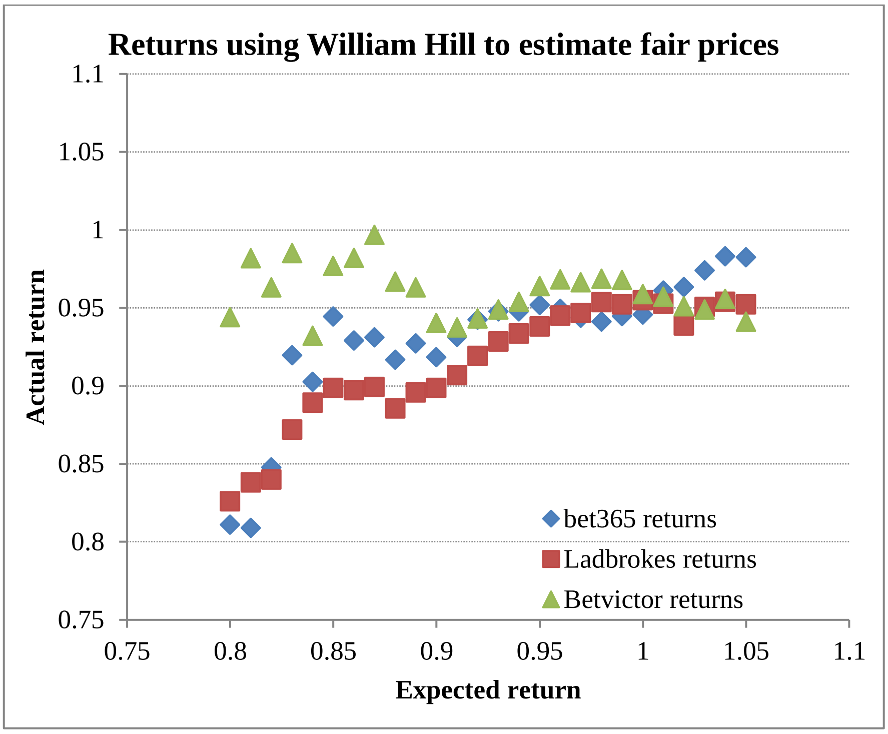

Wrapping Up

Expanding on my original discussion of the 'Wisdom of the Crowd' this updated publication has attempted to achieve two goals: 1) further investigate how bookmakers might apply their margin to a football match betting market (home-draw-away) and how, subsequently, we can remove that margin to find 'true' odds; 2) use those 'true' odds and the assumption that they are wise (efficient or accurate) to find betting value and make a profit. Regarding the first, I've proposed four models which describe the application (and removal) of a bookmaker's margin. Three of these take into account the well known phenomenon that is the favourite–longshot bias where odds on longshots are shortened more than they are on favourites. All three of them do a considerably better job of replicating what bookmakers are most likely doing, and consequently of reproducing their implied 'true' odds than a simple model which ignores the bias and assumes the margin is spread equally across home, draw and away odds.

When I wrote the original version of this document in August 2015 I used only one model – the 'margin proportional to odds' model. Whilst this model was formulated purely from conjecture it arguably does as good a job as two new models introduced in this piece of work – the 'odds ratio' model and the 'logarithmic function' model, which may very well have a superior rationale underpinning them. For this reason and for the sake of simplicity I will continue to use my existing model for estimating the 'true' odds of a football result when issuing my bet suggestions for my 'Wisdom of the Crowd' betting system on Football-Data.co.uk. Followers are free to experiment with the other two models at their leisure by means of the **'true' odds calculator** I have made available.

This work has also made clear that the stand out bookmaker for assessing 'true' odds is Pinnacle. When you see their odds you can be confident, relative to other bookmaker brands, that what you see is a real reflection of the 'true' probability of a football match result. More than any other brand their model relies on encouraging turnover rather than large customer numbers, the acceptance of sharp players for the wisest possible market, and the management of betting odds that in many ways resembles what takes place at a betting exchange. Consequently their football betting market will be the wisest and most accurate. Odds published by other bookmakers, by contrast, do not necessarily reflect those 'true' probabilities. Rather, they may at times reflect other motivations of the bookmaker, most specifically the advertising of best prices (and sometime loss-leading prices too) to attract new customers and to create favourable publicity. Consequently using them to estimate 'true' odds and find value may not be appropriate.

Naturally, there are a couple of caveats with this approach to finding value via this 'Wisdom of the Crowd' approach. Firstly, given the relatively modest yields of perhaps a few per cent and at fairly long average odds one should reasonably expect to suffer fairly long periods of treading water, or worse still, losing, lasting hundreds and perhaps thousands of bets. Of course, the same is true, to a greater or lesser extent, for all and any betting one engages in. Secondly, it is to be expected that the sort of bookmaker who will offer betting prices in excess of Pinnacle's implied 'true' odds will also be the sort of bookmaker who won't like a customer consistently exploiting such generosity. As already mentioned, this is usually offered to attract new customers or to advertise the impression that the brand offers good value. If customers repeatedly take advantage of those superior prices they can often expect to have their betting activity curtailed. Advising how a punter can avoid detection in this respect is beyond the scope of this document, although two important things to consider are ways of hiding your IP address and adopting betting patterns that make you look more like a square. There are a number of websites that offer some guidance in this respect [Daily25, Sports Arbitrage Guide, Betting Expert, Betting Systems that Work, Kick Off Profits, Punter2Pro, Gamble Geek]. Nevertheless, the information published in this document has at least identified that a 'wisdom-of-the-crowd' approach can identify where bookmakers have made mistakes, and that technically at least it should be possible to exploit them.

Good luck in trying.

Joseph Buchdahl