

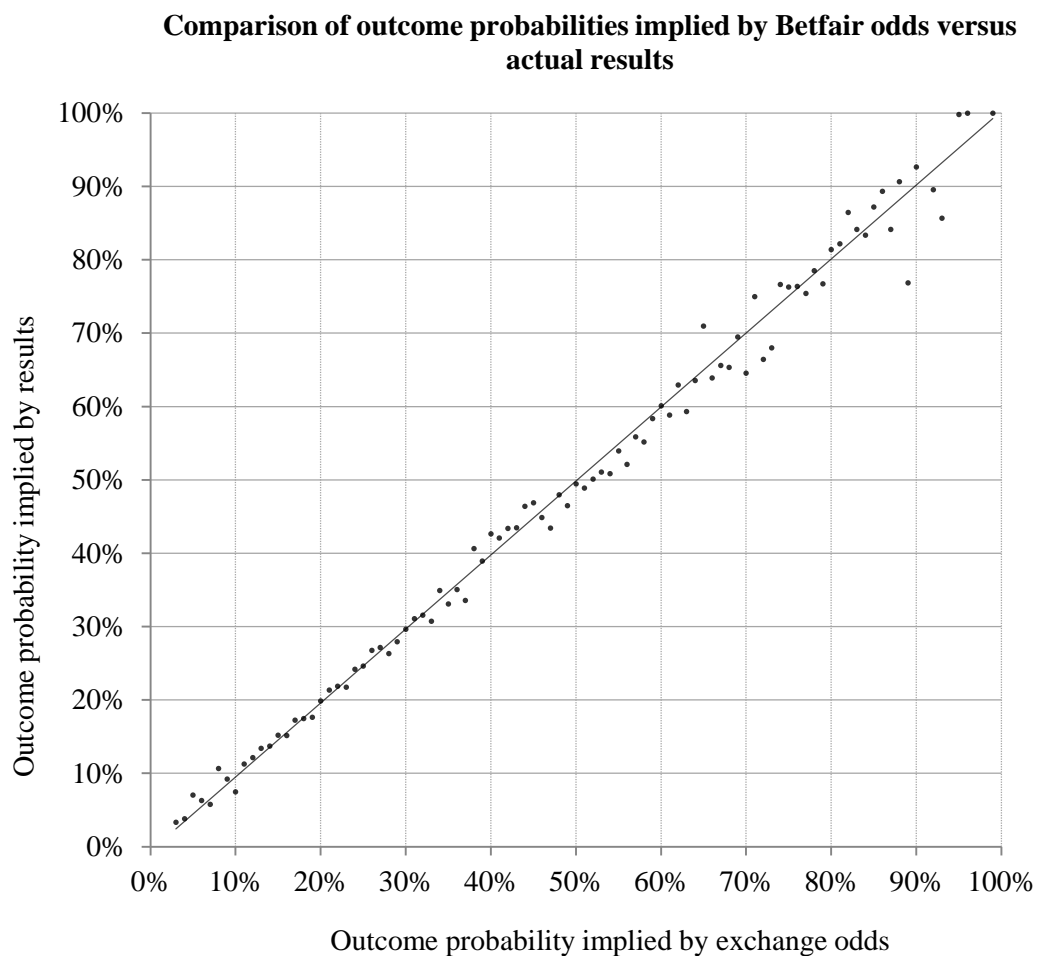
## *The Contrarian Betting System*

In a football betting market, the odds of a team winning represent the probability of such an outcome occurring. For example, fair odds of 2.50 imply the team has a 40% chance of victory ( $1/2.50$ ). Of course, a bookmaker's odds are never fair. Built into them is a margin to help him make a profit, no matter what the result. For now, however, let's assume there is no bookmaker's margin. How do we know that odds of 2.50 really mean a team has a 40% chance of winning? Of course, before the event, we can't. The odds, essentially, just represent an estimate of what we think the probability is. However, we can test how accurate our estimates were, by checking the results for lots of teams that were priced 2.50. If 40% of them ended up winning that's pretty good evidence that our estimates, in aggregate, were reliable.

Of course, when we go to a bookmaker, we don't define the odds, the odds are their waiting for us to back. What, then, defines what price is published? Naturally, the opening price posted by a bookmaker represents his opinion for the probability for that particular result. Thereafter, what happens to that price will be determined by his customers who choose to back it, or alternatively choose to oppose it by backing other possible results. If disproportionately more people choose to back one result over other possible results, that price should shorten. If fewer, then it should lengthen. Essentially, the odds simply represent the public expression of all privately held opinions about the chances of that result occurring, much as the price of a share on a financial market represents the balance of opinions between buyers and sellers. Bettors who back a price presumably believe that the result has more chance of happening than implied by the odds, much as buyers of a share must presumably believe it is underpriced. Similarly, bettors who oppose a price (by backing alternatives) presumably believe that the result has less chance of happening than implied by the odds, as sellers of shares must believe them to be overpriced. In this sense, a football betting market is really just the same as any other type of buyers and sellers market.

The interesting question to ask in such market environments is, who is right: the backers (and buyers) or the opposers (and sellers)? As I stated earlier, it's not possible to know before the event, but checking retrospectively after many outcomes are known, we can arrive at a good estimate for the answer. Much of the time, it turns out, both backers and opposers are right. In other words, the odds do a very good job of defining the fair price or true probability of the result even before the result is known. Backers and opposers (as buyers and sellers) essentially

engage in what economists call a process of price clearing, arriving, as if by magic, at a price that closely represents true probabilities (or in the case of financial markets, the true value of assets), an equilibrium price, if you like. A good illustration of these equilibrium prices can be seen at the betting exchange Betfair. The chart below compares the outcome probabilities that were implied by volume-weighted average betting prices for 52,411 football matches played between October 2004 and October 2005 with their actual results. The correlation is close to perfect. Such a market, where prices accurately predict the probabilities of real outcomes, is said by economists to be efficient.



A consequence of market efficiency and prices that closely match true probabilities in a betting market is that a consistent profit is rather difficult to come by, particularly after the costs of playing in the market (via the bookmaker's margin or overround) have been taken into account. If, on average, both backers and opposers are right about the odds, how can anyone consistently find an advantage? Of course, bettors won't be able to accurately predict a true price all the time, but if their errors are random (sometimes overestimating the true

probability and sometimes under estimating it), their betting performances will be largely random too, that is to say, driven by luck, both good and bad.

In recent decades, however, the hypothesis of market efficiency has, to some extent, fallen out of favour. Many now argue that human beings do not behave in a manner that leads to such efficient pricing, on account of the expression of systematic (non-random) errors or biases that result in less than fully rational markets. One reason for this is our inability to properly weight large and small probabilities. Specifically, it is observed that people systematically overestimate the probability of unlikely events in terms of the money that they are willing to wager on their occurrence. Such a bias, for example, explains why so many people are willing to play the lottery, despite such a small possibility of winning and with such poor mathematical expectation. Another example can be seen in football betting; it is called the favourite–longshot bias. Underdogs, or longshots, in football match betting markets (as well as in other sports for that matter) frequently have prices that are shorter (after accounting for the bookmaker’s margin) than they should be, as implied by actual outcomes. In contrast, favourites are found to be priced longer than they really ought to be. The implication is that bettors are over-betting longshots (because they are overestimating the chances of lower probability outcomes) whilst under-betting favourites. Sadly, this bias appears to have weakened in recent years as more and more bettors have exploited it, and because of the bookmaker’s profit margin, its strength is not sufficient to consistently push the prices for favourites into positive mathematical expectation, where betting them over the longer term would yield a profit. For example, betting average football match odds shorter than about 1.50 will typically lose the punter about 3 to 4% on turnover, compared to over 20% betting teams priced longer than 5.00. By using a number of bookmakers and backing the best possible prices, betting favourites should enable the bettor to roughly break even.

Fortunately, the misjudgement of probabilities is not the only way people can express a bias in a betting market. One of the most common ways people exhibit a systematic bias is via the hot hand fallacy, sometimes also called the reverse gambler’s fallacy. The error initially arises because the extent to which randomness or luck in a repetitive pattern or streak is present is underestimated. As we’ve observed before, with betting prices doing a pretty good job of describing true probabilities, profitability from betting those prices will, in the absence of systematic errors expressed by other players, be largely random. In its place, other causal explanations for such streaks are assumed to be more relevant, causes which should prolong the longevity of the streak. In expressing such a fallacy, the influence of regression to the mean, the tendency for a variable to be closer to the average on a subsequent measurement

following a previous extreme one, will be ignored. 'What goes up has a tendency to come down' is substituted with 'what goes up will probably stay up for longer.'

The phenomenon of regression, or reversion, to the mean was first uncovered by Sir Francis Galton, the Victorian polymath, as he experimented with the heredity of sweet peas. In cross breeding trials, Galton noted a tendency for the size of the offspring to show a smaller distribution than that of the parents. Crucially, whilst the offspring of larger parents tended to be smaller, the offspring of smaller parents tended to be larger. It is important to realise that there is no requirement for any teleological cause for this regression in a strictly deterministic sense, merely a random process that sees extremes become less extreme. As if to demonstrate this point, paradoxically, regression to the mean is not time dependent; if subsequent measurements are more extreme, the tendency will be for their earlier ones to be closer to the average. Regression to the mean, then, is entirely reversible. Crucially, this principle informs us not that things **must** return to the average, just that they have a **tendency** to do so. Previous good fortune, then, has a tendency to be followed by outcomes that are less fortunate. The same is true for bad luck as well.

Let's consider a football example: a team on a 6-match winning streak. If the (fair) odds for winning those games were respectively 1.50, 2.75, 1.72, 3.80, 1.66 and 2.50, and those prices were fairly accurate representations of the true outcome probabilities, then arguably the probability of such a hot streak occurring by chance would be less than 1%; in other words, pretty lucky. When such an extreme streak occurs, however, the probability that something caused it to happen (for example players playing well) will, by many, be overestimated, as a consequence of the expression of the hot hand fallacy and ignorance of regression to the mean. To put it another way, because we have assumed that the probability such an outcome arising as a result of luck is so small, given that it happened, luck cannot have had much to do with it. If such a thought process increases the likelihood of backing the team in their 7th match, bettors expressing the hot hand fallacy may show a tendency to over-bet teams with better recent form relative to their opposition. As a consequence, we might expect the odds for such 'hot' teams to be shorter than a more objective assessment of outcome probabilities would suggest they ought to be. Similarly, odds for 'cold' teams might be expected to be priced longer. That, at least, is the theory. Is there any way this can be tested in practice? Perhaps more importantly, if any systematic expression of the hot hand fallacy exists, will such a bias provide consistently profitable opportunities for contrarians looking to exploit it, in a way that the favourite-longshot bias appeared largely unable to.

To test such a hypothesis we need some way of measuring how ‘hot’ or ‘cold’ teams are. One way is to use the betting odds themselves. If we assume that, on average, the ‘fair’ odds represent the ‘true’ probability of a team winning then over the long term our expected return betting such odds should be approximately zero. Consequently, teams on ‘hot’ streaks will show positive returns over the short term, whilst those on ‘cold’ streaks will show negative returns. As we know, a bookmaker’s odds are not fair, but we can use them to estimate what the fair odds would be with his margin removed. To remove it we have to know how he applies his margin to a typical match betting market. I’ve already mentioned that longer prices will tend to be shortened disproportionately more than shorter prices, on account of the favourite–longshot bias. Whilst the bookmakers will never tell us exactly how they do it, I have attempted to estimate the process by means of a simple algorithm which assumes that the weights applied to each possible outcome are in inverse proportion to the probabilities of those outcomes, that is to say, bigger weights for longer odds. Hence, for a book with  $n$  runners and overall profit margin  $M$ , the differential margin applied to the fair odds for the  $i^{\text{th}}$  runner ( $O_i$ ) will be given by:

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

For a home-draw-away football betting market,  $n = 3$ . Hence:

$$M_i = \frac{MO_i}{3}$$

where  $M$  is given by

$$\frac{1}{H} + \frac{1}{D} + \frac{1}{A} - 1$$

and  $H$ ,  $D$  and  $A$  are the bookmaker’s (unfair) home, draw and away odds respectively.

For example, if the fair odds are 1.5, 5 and 7.5 and the bookmaker’s margin,  $M$ , is 8% (or 0.08), the differential margins for home, draw and away would be 0.040, 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  $1.5 \div 1.04 = 1.44$ . Similarly, the draw and away prices are  $5 \div 1.133 = 4.42$  and  $7.5 \div 1.200 = 6.25$ . You can see from this exercise that a differential weighting of odds in this manner shortens longshots more

significantly than favourites, in accordance with the favourite–longshot bias. Whilst the basis for this simple odds model is purely conjecture, it does appear to closely reflect the betting prices for many of the major brands.

Of course, here, we wish to use prices quoted by the bookmaker to estimate what prices he implied were fair. With a little bit of algebraic rearranging, we can reverse the process, using the following equation.

$$O_i = \frac{nO_{\text{bookmaker}}}{n - MO_{\text{bookmaker}}}$$

For a football match betting market, therefore:

$$O_i = \frac{3O_{\text{bookmaker}}}{3 - MO_{\text{bookmaker}}}$$

Of course, if the bookmaker’s odds we are using to estimate ‘fair’ odds themselves result from the systematic bias we are trying to find, this presents a kind of circular reference problem. For our purposes here, however, it is probably safe to assume that the relative rates of winning and losing by teams will provide a much bigger influence towards short term returns than small inaccuracies in odds arising because of expression of the hot hand fallacy, or any other systematic error for that matter. For example, the difference between winning and losing at odds of 2.00 is 200% of the stake. By contrast, the difference between winning at 2.00 and winning at 1.95 is just 5% of the stake. If our hypothesis is correct, backing relatively ‘colder’ teams should prove to be more profitable (or at least less unprofitable) than backing relatively ‘hotter’ teams, by virtue of the fact that disproportionately fewer people are backing them. The following analysis would appear to offer considerable support to the presence of such a systematic bias in a football match betting market.

For the 5 domestic seasons 2010/11 through 2014/15, average home-draw-away match betting odds (collected from the odds comparison Betbrain.com) were used to estimate ‘fair’ odds for 36,126 league matches in 22 European divisions. For each match, the winning team was awarded a risk adjusted score of  $[1 - 1/\text{odds}]^1$ , whilst the losing team (or both teams

---

<sup>1</sup> This is simply the profit won from a bet with stake 1/odds. For this model I have preferred to use risk-adjusted returns over level stakes returns to minimise variance in scores. A lucky win at odds of 10/1, for example, will have a much bigger (and potentially unwarranted) influence on short term cumulative returns calculated with level stakes. The same would be true for an unlucky loss at odds of 1/10.

where they drew) was awarded a score of  $[-1/\text{odds}]$ . For each team, these scores are consecutively added during the course of a season, and reset to zero at the start of the next one. The process is perhaps best illustrated by means of an example, in this case for Blackpool's first 10 league games in the 2010/11 season, as shown in the table below. We can see that, during this period, Blackpool had over performed relative to what the betting market had expected the team to achieve. After 10 games, betting risk-adjusted stakes at these theoretical fair odds would have netted the player over 2 units.

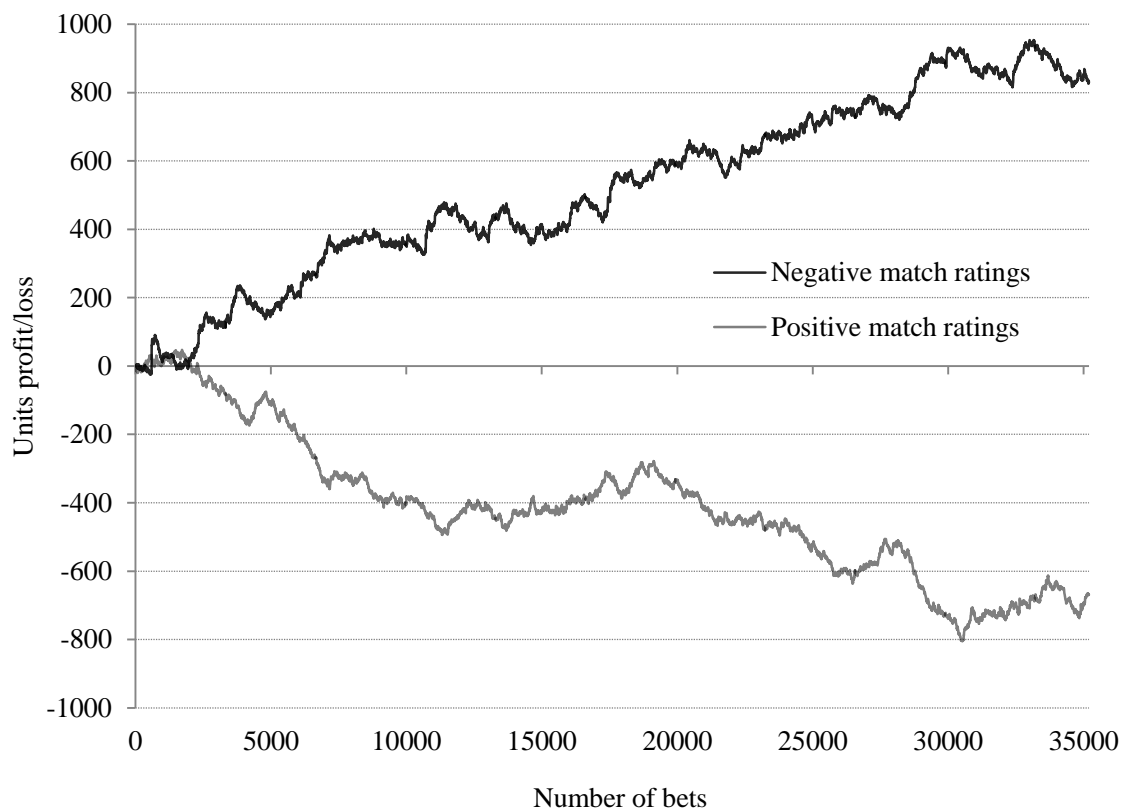
Team	Opposition	Date	Fair Odds	Result	Profit	Running Score
Blackpool	Wigan	14/08/10	4.96	Won	0.798	0.798
Blackpool	Arsenal	21/08/10	27.40	Lost	-0.037	0.762
Blackpool	Fulham	28/08/10	3.29	Lost	-0.304	0.458
Blackpool	Newcastle	11/09/10	6.92	Won	0.855	1.313
Blackpool	Chelsea	19/09/10	48.20	Lost	-0.021	1.292
Blackpool	Blackburn	25/09/10	3.34	Lost	-0.299	0.993
Blackpool	Liverpool	03/10/10	12.80	Won	0.922	1.915
Blackpool	Man City	17/10/10	6.39	Lost	-0.157	1.758
Blackpool	Birmingham	23/10/10	4.72	Lost	-0.212	1.546
Blackpool	WBA	01/11/10	2.95	Won	0.662	2.208

The next step is to utilise these cumulative scores to design a predictive rating. After Blackpool's first game, for example, their cumulative score was 0.798, on account of winning their match against Wigan. This score is therefore taken as their team rating for their next game. In other words, it is a measure of how 'hot' or 'cold' they are going into their next game. Naturally, prior to their first game of the season, their rating will be 0. Repeating this process for every team, we can then finally produce a match rating, defined by subtracting the rating of one team away from the rating of their opposition. Swapping the teams around will simply provide a rating of opposite sign with equal magnitude. So, for example, Blackpool entered their game with Manchester City on 17 October 2010 with a team rating of 1.915. Manchester City, similarly, entered the game with a rating of 0.521. Hence, we can calculate the match rating by  $1.915 - 0.521 = 1.394$  (or  $-1.394$  if calculated the other way around). In other words, this is equivalent to saying that prior to their match, Blackpool had been performing relatively better than expected compared to Manchester City. Of course, Manchester City, with a positive team rating themselves, had also been over performing, but

just not to the extent that Blackpool had been. In contrast, when Blackpool met Birmingham in their next game, Birmingham had a team rating of -1.509, implying they had been doing worse than the betting market had predicted. The match rating for that game was 3.267 in favour of Blackpool.

Had someone been able to bet every home and away result for each of the 36,126 matches in this sample at the theoretical 'fair' odds and to level stakes (a total of 72,252 bets), their profit over turnover would have been 0.22%. The fact that it wasn't exactly zero will be a consequence of either model inaccuracy in the way the 'fair' odds have been calculated, slight (lucky) over performance of longer odds relative to shorter ones during this 5-season period, or a combination of the two. Nevertheless, the figure is reasonably close to what we could expect betting at 'fair' odds to yield. The time series of accumulated profits/losses, furthermore, shows a fairly typical random walk about the break-even line. Contrast that to the time series for betting, on the one hand, all negative match ratings (where we favour a relatively 'colder' team over a relatively 'hotter' team), and on the other hand, all positive ratings (where we favour a 'hotter' team over a 'colder' team). Again, bets are struck at theoretical 'fair' odds and to level stakes. The results are graphed below.

**Theoretical bankrolls from level staking**





The conclusions appear stark. Backing teams (at theoretical 'fair' odds and to level stakes) that had been performing relatively 'colder' to their opposition in this sample would have delivered a +2.36% yield. What is more, the profit-taking appears consistent across the 5 seasons. In contrast, backing relatively 'hotter' teams would have shown losses of -1.90% on turnover. The difference between these two yields for this sample is statistically significant, with just a 0.4% probability that it could have arisen by chance only. The implication must be that relatively 'colder' teams are overpriced as a consequence of fewer people backing them and, as such, provides evidence of a systematic bias arising out of the hot hand fallacy and a disregard for regression to the mean. Being a contrarian in football match betting might very well be profitable.

We should not, however, let excitement get the better of us just yet. Firstly, statistical significance compared to 0% expectation (or break even) is weaker, with a 1.5% probability that the profitability from backing relatively 'colder' teams has arisen by chance only. Secondly, it must be admitted that profitability is only small. A further consequence of that will ensure that short term variance in the swings of profits and losses will be considerable, as you can see from the chart. Finally, we need to remember that the odds used to produce this performance were theoretical. How would things look in the real world? In fact, blindly backing best available market prices to level stakes for this sample would have seen a -0.51% yield. Backing relatively 'colder' and relatively 'hotter' teams would have seen yields of +1.41% and -2.39% respectively. The lower figures are a consequence of the fact that, on average, even best prices advertised at Betbrain.com are not quite 'fair'. In this sample, the mean overround was 100.6%.

We could, of course, choose to be more selective in our betting criteria. What if our expression of betting interest was limited to progressively 'colder' teams, that is to say limited to match ratings that were not simply negative but below -1, -2, -3 and so on? The effects of such selective betting can be seen in the next table.

<b>Match ratings &lt;</b>	<b>Bets</b>	<b>Yield from 'fair' odds</b>	<b>Yield from best odds</b>
0	35,200	2.36%	1.41%
-1	23,215	4.44%	3.46%
-2	14,230	4.87%	3.77%
-3	8,355	5.27%	4.17%
-4	4,742	2.85%	1.72%
-5	2,564	-0.17%	-1.25%
-6	1,315	-7.01%	-8.19%
-7	672	-7.90%	-8.94%
-8	316	-3.58%	-4.86%
-9	159	-7.48%	-9.08%
-10	61	18.56%	16.21%

As one might reasonably expect, restricting our betting to ever decreasing match ratings improves returns, but apparently only to a point. Below match ratings of -4 the improvement ceases, and indeed results in unprofitable returns below match ratings of -5. Of course, it should be observed that as our selectivity increases, the number of betting opportunities vastly diminishes. The apparent failure of our hypothesis to explain returns at these low match ratings could simply be a consequence of greater variance and bad luck due to the much smaller sample sizes. Alternatively, however, it could be a result of some other factors not taken into account, where the relative difference between the 'heat' of two teams is large. Nevertheless, despite these caveats, there are strong reasons to believe that this apparent bias in the football match betting market represents something more than a spurious statistical association. After all, in contrast to most data mining exercises that go looking for something profitable until it is found, this exercise began with a hypothesis of the hot hand fallacy, and sought to test it with real data.

I chose to charge a fee for this document, rather than give the material away freely, so as to protect the underlying bias that may exist here. As more bettors test the system, the inefficiency will begin to disappear. That is the nature of a market, where advantages, once discovered, will be exploited until they no longer exist. You are free either to try this methodology as described, or perhaps use it alongside other betting strategies that you may follow. I will be freely available to assist anyone looking for advice as to how to use Excel to develop the match ratings as I have described them here.