

The Independence Referendum: Predicting the Outcome¹

David N.F. Bell

Division of Economics
Stirling Management School, University of Stirling,
IZA, Bonn and
ESRC Centre for Population Change



¹ Thanks are due to my colleagues, Angus Armstrong and Liam Delaney for helpful comments on earlier drafts of this paper. As usual, the responsibility for its content rest with the author.

Introduction

What is the best way to forecast the outcome of the Scottish referendum on independence? This is a vitally important question for individuals, organisations and businesses trying to plan for the post-referendum period. This paper introduces an alternative way of forecasting the referendum outcome, which under certain assumptions will produce a forecast with good statistical properties. It is based on the use of what are now commonly referred to as “prediction markets”. In turn, the superiority of the use of prediction markets in forecasting over other types of forecast is based on a somewhat abstruse theorem in financial economics that is known as the “efficient markets hypothesis” (EMH).

This paper begins by describing prediction markets and how they relate to the EMH. It then uses these arguments to develop a current forecast of the likely outcome of the referendum before going on to discuss how these forecasts relate to opinion polls which measure voting intentions in the referendum.

Prediction Markets

The EMH is an important building block in the understanding of financial and economic markets. The basic idea is that the market price of an asset or good reflects all the information relevant to setting that price. If the EMH is true, there is no point in trying to beat the market because any information that might have a bearing on the price is instantly reflected in the price. One of the best examples of this effect was documented by Roll (1984). He showed that the production of concentrated orange juice was concentrated in a relatively small area around Miami, Florida. The yield, and therefore the price, of the orange juice were highly dependent on weather conditions in this area. He showed that the futures price of orange juice was closely linked with local weather forecasts and indeed outperformed these forecasts in day-ahead temperature predictions.

This particular market is one in which most relevant information (i.e. weather forecasts) to the setting of the futures price is publicly available. The issue of what information is available to whom and whether it can be acted upon is an important aspect of different flavours of the EMH. Under the “strong” form, on average, prices reflect the true value of the underlying good or asset and react instantaneously to new public and private information. In the “weak” form of the EMH, asset prices follow a random walk and cannot be used to predict future prices. Though the EMH is contested, particularly in its strong form, it is a useful rule of thumb as to how markets behave. The orange futures market is a good example.

But, I hear you ask, what does this have to do with the Scottish referendum? Some economists, notably Wolfers and Zitzewitz (2004) and Leigh and Wolfers (2006), have noticed that the EMH may be a useful in helping predict many different kinds of events. This is because there are many markets, like orange futures, where prices are determined by the likelihood that an event will happen sometime in the future. If such markets are efficient (i.e. the EMH holds), then the futures price is the best possible guess of where the actual price will be in the future (after adjusting for the cost of capital over the relevant period).

Commodity futures markets form one part of the set of all “prediction markets”, also known as “event futures”. Another part of this set is the domain of the gambling industry. It provides a myriad

of prediction markets where individuals can bet on the outcome of all kinds of events. Sports and politics are popular prediction markets.

Bets can involve various forms of contract from "winner take all", "index" and "spread". This article focuses on the "winner-takes-all" contract, which is perhaps the type that is most familiar to amateur gamblers. In this form of betting, it is worth taking a bet if you think that the chance of winning is better than the odds available. For example, if a bookmaker offers odds on a horse of 5/1 and you estimate that the chance of the horse winning is 33 per cent (equivalent to odds of 2/1), then the bet is worth taking. If you are right and in the (unlikely) event that you were in this situation 3 times then you could expect 1 winner (paying 5) and 2 losers (losing 1 each time) giving a profit of 3.

There is a complication: bookmakers have to make a living. They have to ensure that a gambler can never be in a position where he/she can be certain not to make a loss. To do this, they ensure that if someone places bets on both sides of a gamble that would generate the same winnings, these winnings will be less than the value of the stakes. The difference is the bookmaker's margin. However, the bookmaker is not free to set this margin, so long as there are many bookmakers. In these circumstances, competitive pressures will drive down the size of the margin, since bookmakers need to attract business by offering competitive odds, in order to make any profit at all. In the limit, when the margin is driven down to zero, the market will be efficient and will accurately reflect the odds of the event occurring based on all of the available information at the time.

Even in an efficient market punters may place irrational bets. This creates a short-run opportunity for rational gamblers to increase their returns. As Wolfers and Zitzewitz put it:

"In a truly efficient prediction market, the market price will be the best predictor of the event and no combination of available polls or other information can be used to improve on the market-generated forecasts. This statement does not require that all individuals in a market be rational, only that the marginal trade in the market is motivated by rational traders"

The range of events on which a bet can be placed has increased hugely in recent years. It now includes most political contests. As implied by Wolfers and Zitzewitz, other predictions of political events are available – from opinion polls to econometric models. But if it is appropriate to apply the EMH in such markets, then no other forecast can improve on the market-prediction.

Not surprisingly, this argument is contested. Erikson and Wleizen (2008) argue that opinion polls reflect opinion on the day they were collected, and therefore should not be naively interpreted as forecasts. Further they suggest that, if opinion poll data are appropriately adjusted, they will outperform market predictions. However, Berg et al (2001) compare opinion polls and market based predictions from 19 national elections. Their findings suggest that market predictions provide a serious alternative to opinion polls:

The market outperformed polls in 9 of 15 cases according to each measure (election eve closing prices and last week average prices). Across all elections, the average poll error was 1.91% while the average market error was 1.49% and 1.58% by the two measures.

In a few cases (the 1988 and 1992 U.S. Presidential elections) the market dramatically outperformed polls. Berg et al (P 4)

Rhode and Strumpf (2004) examine betting patterns in US Presidential elections. There were very active betting markets in the outcomes of Presidential elections. They argue that the market was fairly efficient despite the limited information of the gamblers and the attempts by political parties to influence the odds. Their results show that, of the 15 elections between 1884 and 1940, the mid-October betting favourite won 11 times (US Presidential elections occur between Nov 2nd and Nov 8th). On one occasion the favourite lost and in the remaining three cases, it was difficult to determine which candidate was favourite because the races were very close. At the time, opinion polls were not available and information on trends in political campaigns was less pervasive. In the circumstances, the predictive markets did rather well.

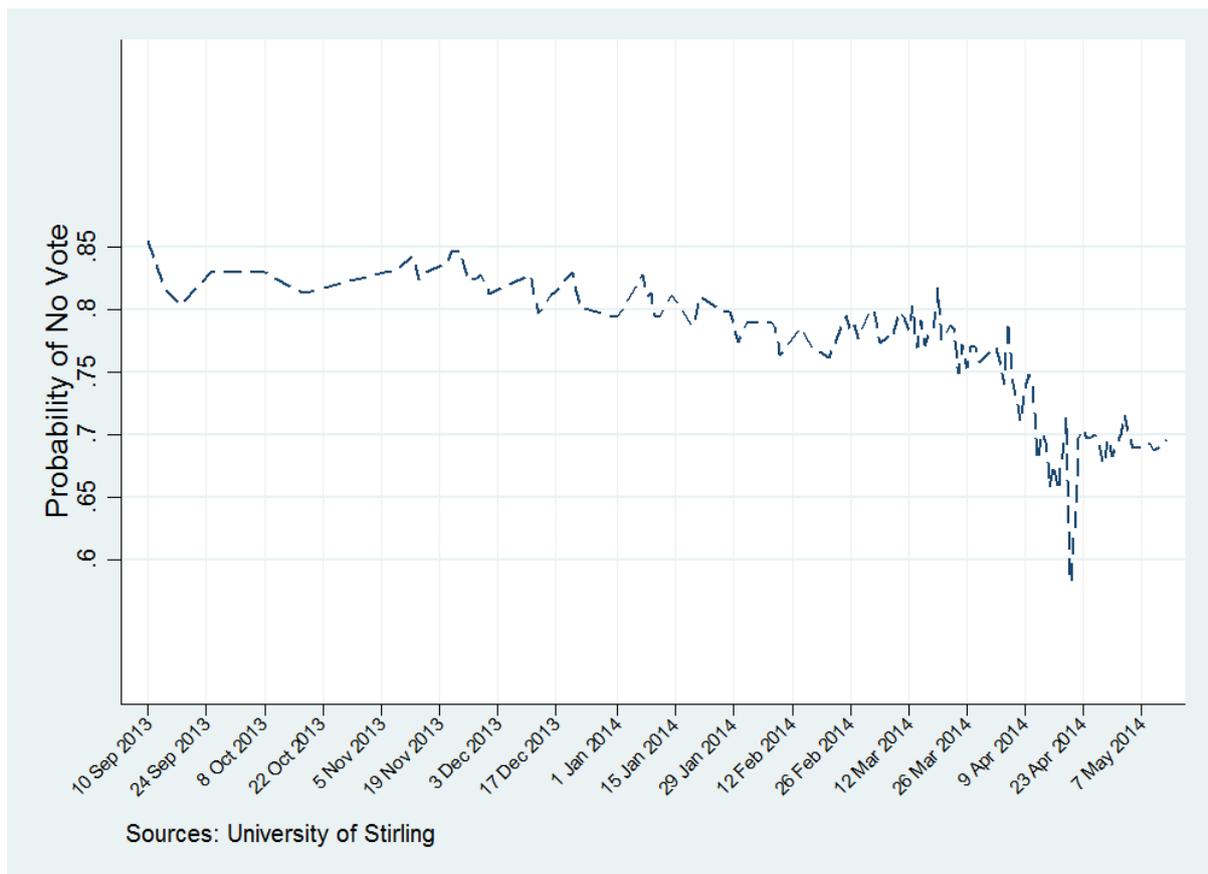
Evidence for the Scottish Referendum

This suggests that market-based prediction provides a serious alternative to opinion polls in predicting political contests. So what does the market say about the outcome of the Scottish referendum?

The evidence used in this study has been retrieved from the history of the odds offered on “No” and “Yes” votes from 23 bookmakers between Sep 2013 and May 2014. These were averaged and a time series constructed of the implied probabilities of both “No” and “Yes” outcomes. This involved averaging the odds from the bookmakers and then adjusting these averages to account for their margin, which, on average, is about 7 per cent. This process gives probabilities for “Yes” and “No” referendum outcomes that sum to one.

The market-prediction of the probability of No vote is shown in Figure 1. The probability of a “Yes” vote is simply 1 minus the probability of “No”. The probability drifted slowly downwards between Sep 2013 and Mar 2014, averaging just over 80 per cent during that time. Since then, it has been more volatile, dropping below 0.6 before increasing again to around 0.7 at the beginning of May 2014. At the end of the period, (the second week of May), the market estimate is that there is around a 7 in 10 chance of a “No” vote.

Figure 1: Market-based probability of “No” vote in Scottish Independence Referendum



Are there trends which suggest that the probability is moving in a particular direction? Recalling how the EMH is constructed suggests that this is not a very sensible question to ask. If there were predictable trends, then these would be exploited by rational gamblers. The probability is the market’s best guess of the likely outcome, given currently available information. Existing information cannot be used to improve the market forecast.

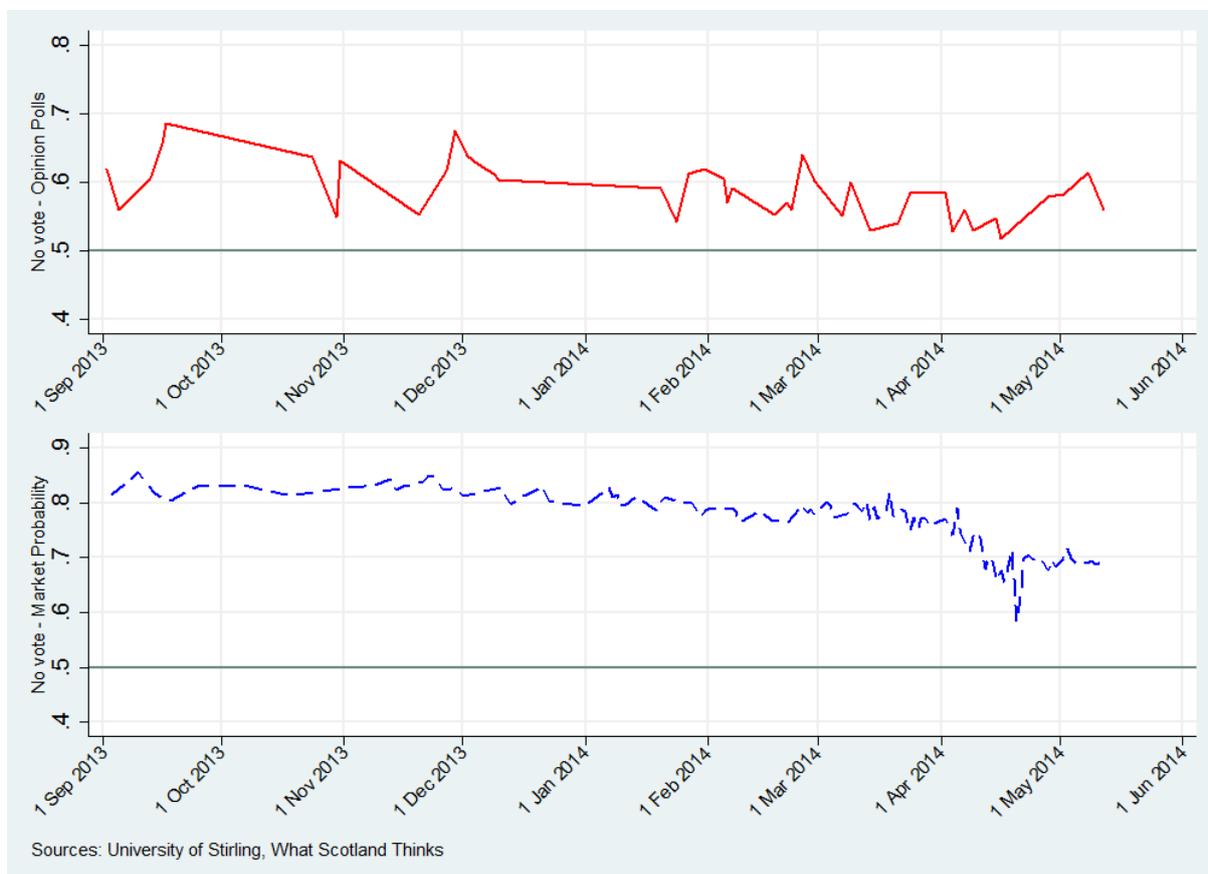
Of course there are caveats. Expectations may be unreliable in thin markets – where trades are relatively infrequent. Thin markets may also be affected by relatively large bets being placed on one side or the other. To cover their exposure to risk, market makers (bookmakers) are likely to increase their margins in thin markets. Our data suggest that bookmakers’ margins have been stable at around 7 per cent for the last few months, as have the derived probabilities. This does not suggest that the market is particularly thin. In 2011 and 2012, when relatively few bets were being made, fluctuations in the probability of both the “Yes” and “No” outcomes were much greater.

Nevertheless, it is difficult to establish whether one or other sides of the campaign is trying to manipulate the market. In so doing, of course, they risk losing their money by moving the odds away from an unbiased estimate of the outcome. This was quite common in US politics: betting on a favoured candidate to shorten his/her odds might influence voting behaviour, if voters, other things being equal, have a tendency to support favourites.

Opinion Polls

As indicated in the previous discussion, the obvious comparators to market-based predictions are the regular opinion polls that are testing the independence question. Figure 2 compares the market based probability with the opinion poll data drawn from the “What Scotland Thinks” website. It should be emphasised, that particularly at this stage of the campaign, these two time series measure different things. The opinion polls are drawing data on *current* voting intentions. We know that these might change if individual’s perception of the economic costs and benefits of independence change Bell, Delaney and McGoldrick (2014). The market-based probability is an estimate of the likely outcome of the referendum. The comparison is interesting in that it may reveal the extent to which the opinion polls are affecting expectations of the referendum result.

Figure 2: Market-Based Probability and Opinion Polls: Scottish Referendum Sep 2013-May 2014



First, it is clear that the market estimate of the probability of a No vote is higher than the proportion of voters that may vote No². This is not surprising. Even if the “No” vote in the opinion polls exceeded the “Yes” vote by a small margin, it would be rational for a high proportion of bets to be placed on a “No” outcome, provided that those betting on this outcome thought the chance that this margin would not disappear by the time of the referendum.

² “Don’t know” voters have been eliminated from the calculation of the proportion of “No” voters. This implicitly assumes that the “Don’t knows” will either not vote or vote “Yes” or “No” in the same proportions as the existing support for these outcomes.

Second, market outcomes tend to be more stable than the opinion polls. Differences in the way that samples are drawn by polling organisations, sampling errors, differential delays in the publication of results etc. may explain some of this variability in the opinion polls.

Third, a succession of polls that showed an increase in support for the Yes campaign during March and April 2014 may have caused a significant reduction in the odds for a “No” outcome in the referendum. From mid-March to mid-April, an increase in the value of bets on a “Yes” outcome caused the implied probability of a “No” outcome to fall from 0.8 to 0.6. By early May, it had returned to 0.7. Whether this fluctuation was a response to the increase in support for “Yes” observed in the opinion polls is difficult to test statistically, partly because, unlike the betting odds, opinion polls are not observed at regular intervals.

Fourth, Bell, Delaney and McGoldrick (2014) have established previously that those likely to vote “Yes” in the referendum are more willing to take risks. For example, we know that in response to the following question:

Imagine you had won £1,000 in a lottery. Almost immediately after you collect your money you receive the following offer from a reputable bank: if you invest with them, there is a chance to double your money in two years. However, it is as likely that you could lose half the amount you invest. What part of the £1,000 would you choose to invest with the bank?

Yes voters are, on average, willing to wager £198, while the average No voter is only willing to gamble £153. Might this finding suggest that the probability of a No vote is systematically underestimated because a relatively high proportion of Yes voters are willing to gamble, given their greater willingness to accept risk? There are a number of implicit assumptions that need to be examined before such a conclusion could be drawn. First, being willing to take risks does not automatically imply that one is more likely to gamble. The prevalence of gambling is not simply determined by attitudes to risk, but is also influenced by a variety of social and economic factors. These may mask any effect of different aptitudes to risk among Yes and No voters. Second, even if Yes voters were more willing to gamble, there is no guarantee that they would gamble on the outcome of the referendum. Other gambles may be more attractive than voting on the referendum. Third, even if willing to gamble on the referendum, it does not follow that Yes voters would bias the outcome. In any gamble, the effect of any group with an interest in the outcome depends on their willingness to accept a loss where their interest in the outcome conflicts with their assessment of the likely result. There is no evidence on gambling patterns on the referendum among either Yes or No voters so it is not possible to draw any conclusion on this issue.

Conclusion

The use of prediction markets in relation to electoral outcomes is well-established. Their accuracy relies on there being some well-informed, dispassionate traders/gamblers who wish to profit from their trades/bets. Though these conditions do not always hold, prediction markets have a good record in forecasting future outcomes.

This paper has calculated the probable outcome of the Scottish referendum based on the latest prediction market data (up to May 2014). The latest data suggest that there is a 70 per cent chance of a No vote (approx.1/2 in terms of odds).

The paper also compares the market probabilities with opinion polls and speculates that changes in these polls, which measure voting intentions at a particular point in time, may have an effect on the prediction market for the Scottish referendum.

Other evidence suggests that Yes voters are more willing to take risks than No voters. While this might affect the prediction market for the outcome of the referendum, there is no evidence that this is the case.

References

Bell, David N. F., Delaney, Liam and Michael McGoldrick (2014) "Citizen Preferences for Constitutional Change in Scotland" accessed at: <http://esrcscotecon.files.wordpress.com/2014/05/individualpreferences.pdf>

Berg, Joyce, et al. "Results from a dozen years of election futures markets research." *Handbook of experimental economic results* 1 (2001): 486-515.

Rhode, Paul W., and Koleman S. Strumpf. "Historical presidential betting markets." *The Journal of Economic Perspectives* 18.2 (2004): 127-141.

Erikson, Robert S., and Christopher Wlezien. "Are political markets really superior to polls as election predictors?." *Public Opinion Quarterly* 72.2 (2008): 190-215.

Leigh, Andrew, and Justin Wolfers. "Competing Approaches to Forecasting Elections: Economic Models, Opinion Polling and Prediction Markets*." *Economic Record* 82.258 (2006): 325-340.

Rhode, Paul W., and Koleman S. Strumpf. "Historical presidential betting markets." *The Journal of Economic Perspectives* 18.2 (2004): 127-141.

Roll, Richard. "Orange juice and weather." *American Economic Review* 74.5 (1984): 861-880.

Tziralis, George, and Ilias Tatsiopoulos. "Prediction markets: An extended literature review." *The journal of prediction markets* 1.1 (2012): 75-91.

Wolfers, Justin, and Andrew Leigh. "Three tools for forecasting federal elections: lessons from 2001." *Australian Journal of Political Science* 37.2 (2002): 223-240.

Wolfers, Justin, and Eric Zitzewitz. Prediction markets. No. w10504. National Bureau of Economic Research, 2004.