


Essays on the behaviour of political and financial markets



Thomas James Auld

Churchill College

Faculty of Economics

University of Cambridge

This dissertation is submitted for the degree of Doctor of Philosophy

July 2022

Declaration

This thesis is the result of my own work and includes nothing which is the outcome of work done in collaboration except as declared in the preface and specified in the text. It is not substantially the same as any work that has already been submitted before for any degree or other qualification except as declared in the preface and specified in the text. It does not exceed sixty thousand words.

An early version of the electoral model in Chapter 1 was submitted as a dissertation for the MPhil in Economics Research at the University of Cambridge in 2017. A theoretical model linking the price of the British pound to betting markets for the Brexit referendum was also included in the MPhil dissertation. There are some similarities in that theoretical model and the one presented in Section 2.2 in this thesis. This is particularly the case for the discussion of the assumptions of the model: Section 2.2.1. However, the model presented here applies to all financial assets, not just the British Pound. A Later version of the Mphil dissertation was published with Oliver Linton, Professor of Political Economy and Director of Research, Faculty of Economics, University of Cambridge, in Auld and Linton (2019). Professor Linton largely acted in a supervisory role. 90% of Auld and Linton (2019), and the work mentioned above, is my own work. Chapter 1 is an even later, improved, version of earlier work. Most significantly, it is corrected for a major data error present in the Mphil thesis and Auld and Linton (2019). The theoretical asset pricing model of Section 2.2 also relies on fewer assumptions.

Tom Auld

Cambridge, July 2022

Essays on the behaviour of political and financial markets

Thomas James Auld

Abstract

This thesis considers the behaviour and relationships between financial and prediction markets around elections.

We begin by reviewing the literature. There are many small studies of individual elections and events, particularly of the 2016 UK European Union referendum. However, no studies that consider multiple events, nor present theories that apply in a general setting, are found. We believe this is a gap in the literature.

Chapter 1 begins with a study of the Brexit referendum. Using a flexible prior and Bayesian updating, we demonstrate a major violation of semi-strong market efficiency in both the betting and currency markets on the night following the vote. It appears that it took a full three hours for prices to reflect the information contained in the publicly available results of the referendum.

Chapter 2 presents a model linking the prices of financial and binary options in the prediction markets in the overnight session following an election. Starting from basic assumptions we find that prices in both markets should be cointegrated. Under risk neutrality the relationship is linear. However, departures from this assumption result in a non-linear cointegrating relationship. We test the theory on three recent elections. Strong support for the theory is found for two events. The linear cointegrating model fits the data from the night of the EU referendum remarkably well. However, departures from risk neutrality are needed to explain the behaviour observed on the night of the 2016 US presidential election.


Chapter 3 considers pricing relationships in the weeks and months leading up to an election. Again using economic assumptions, we derive a relationship between asset price returns and changes in the prices of betting market binary options linked to an election result. This model is extended to equities using the ubiquitous Fama–French 5 factor model. The result is a 6 factor characteristic model,

where the additional factor is related to political risk. We test the model on six recent elections. Using daily data, strong support is found for the theory for four events and weak evidence for one. The remaining election does not appear to be informative for asset prices. Interesting relationships are also uncovered between firm characteristics and political sensitivity. This is achieved by exploring the political factor loadings of the different equities under study.

The main contributions of this thesis are, one, using a flexible Bayesian approach to demonstrate that, without a shadow of a doubt, any ‘bubble’ in opinion for remain continued well into the night of the EU referendum, and two, presenting pricing models of prediction and financial markets that apply in general settings and have strong support in the data. We also show that on nights after elections, betting markets lead financial markets on the scale of minutes to tens of minutes. This is consistent with, and an extension of, the conclusion of the existing literature that prediction markets have superior forecasting ability. Whether or not this lead-lag relationship occurs at other times prior to political events is an open research question.

Acknowledgments

I would like to thank my supervisors Professor Oliver Linton and Professor Alexei Onatski for all the help they have given me over the last few years. I would also like to thank Duncan Hawthorne for identifying the timing issue of the Betfair data in Auld and Linton (2019). My wife, Sue, and my three children have also been exceptionally patient and understanding during this, my second, PhD. However, this thesis is dedicated to all the chancers, schemers and schmoozers who have, over the last few years, provided such rich material to study. I am of course talking about our political leaders...



Contents

List of Figures	4
List of Tables	6
Introduction	8
Literature Review	10
Political prediction markets	11
Election risk and financial markets	15
The efficient market hypothesis	17
Fama–French factor models	18
Events	19
1 Semi-strong market efficiency on the night of the EU referendum	23
1.1 Introduction	23
1.2 Electoral model	25
1.2.1 Choice of approach	25
1.2.2 Ex-Ante public information	26
1.2.3 Formation of Prior	27
1.2.4 Prior probability of Brexit	30
1.2.5 Marginal calibration	31
1.2.6 Update	34
1.2.7 Model properties	35
1.2.8 Parameter choices	36
1.2.9 Vote-share	38
1.2.10 Correlation between vote-share and turnout	41
1.2.11 A Note on model correlation and measured correlation	43
1.2.12 Missing Priors	43
1.3 Theoretical framework	43

1.3.1	Betting market	44
1.3.2	Currency market	45
1.4	Results	47
1.4.1	Market efficiency	48
1.4.2	Robustness checks	52
1.4.3	Irrationality	54
1.5	Conclusion	56
2	Election-night models	58
2.1	Introduction	58
2.2	Pricing model	59
2.3	Statistical specifications	71
2.3.1	Cointegration frameworks	71
2.3.2	Non-linear cointegration	73
2.3.3	Statistical tests	76
2.4	Results	81
2.4.1	The Brexit referendum	81
2.4.2	The 2016 United States presidential election	99
2.4.3	The 2014 Scottish independence referendum	116
2.5	Conclusion	123
3	Political factor models	126
3.1	Introduction	126
3.2	Motivation	127
3.3	Pricing model	131
3.3.1	Extension to multi-factor Fama–French model	139
3.4	Empirical specification	141
3.4.1	Significance tests	144
3.5	Results	149
3.5.1	Events	149
3.5.2	Data sources and handling	157
3.5.3	Results	161
3.6	Conclusion	175
	Bibliography	180

Appendix A: Election forecast model summary	192
A.1 Variables	192
A.2 Prior probability of Brexit	194
A.3 Prior marginal calibration	194
A.4 Update	194
Appendix B: Review of probability model in Wu et al. (2017)	196
Appendix C: Power of two-step empirical specification versus a single regression approach	199

List of Figures

1.1	Realised vote-share versus Hanretty expectation for all voting areas predicted in the EU referendum.	28
1.2	Change in UKIP + Conservative vote-share versus change in constituency turnout at the 2015 general election.	41
1.3	Prior distributions for forecast model.	49
1.4	National vote distribution evolution (Logit Normal) and model probability paths for forecast model.	50
1.5	Scatter plot GBP against the price of the Betfair contract for Brexit with the line of best fit.	51
1.6	The forecast model probability and market implied probabilities. . .	52
1.7	Model probability paths when lowering ρ_p	56
2.1	Betfair contract for ‘Remain’ on the night of the Brexit referendum.	82
2.2	Rebased financial asset prices versus Betfair contracts.	84
2.3	Cointegration errors with autocorrelation functions for Brexit. . . .	90
2.4	Bollinger Band trading strategy for GBPUSD.	93
2.5	Bollinger Band trading strategies for Brexit.	95
2.6	Possible ex-ante estimated cointegration errors and Bollinger Band strategies.	97
2.7	Ex-Ante Bollinger Band trading strategies for Brexit.	100
2.8	Betfair contract for a Trump win in the 2016 presidential election. .	102
2.9	Rebased financial asset prices versus US presidential election contracts.	103
2.10	Non-linear cointegration errors with autocorrelation functions. . . .	111
2.11	Ex-post Bollinger Band trading strategies for the 2016 presidential election.	112
2.12	Possible ex-ante estimated cointegration errors and Bollinger Band strategies for USDMXN.	115

2.13	Estimated cointegration relationships for USDMXN (presidential election) and GBPUSD (Brexit referendum).	116
2.14	Estimated cointegration relationships for (a) the 2016 presidential election and (b) the Brexit referendum.	117
2.15	Betfair price for “No” bet and GBPUSD on the night of the independence referendum.	119
2.16	Spuriously estimated cointegrating relationships between the pound and the betting markets for the 2014 independence referendum. . .	121
2.17	Betfair price for “No” bet and GBPUSD on the night of the independence referendum.	122
3.1	Difference in performance of the FTSE250 and FTSE100 indices versus the Betfair contract for Remain.	127
3.2	Prediction based on cointegration model for $\log(\text{FTSE250}) - \log(\text{FTSE100})$ following the Brexit referendum.	130
3.3	Schematic of the empirical approach.	143
3.4	Liquidity on Betfair in the months prior to our chosen political events.	150
3.5	Betfair implied probabilities.	151
3.6	Betfair Exchange market for the ‘Next President’ on 21st Nov 2020, 18 days after voting closed.	158

List of Tables

1.1	Historical UK general election turnout.	37
1.2	Turnout per EC Region in 2010 and 2015.	37
1.3	Opinion polling prior to the EU referendum.	39
1.4	Opinion polls and vote-share for the Conservatives for recent general elections.	40
1.5	Plausible parameter values.	42
1.6	The four voting areas with missing priors*.	44
1.7	How the times to predict brexit vary with more conservative parameter values (Normal Marginal).	53
2.1	Financial assets and changes for the Brexit referendum.	83
2.2	Results of the Phillips–Ouliaris Z_t and Z_α tests for cointegration.	85
2.3	Results of the bivariate Johansen max eigenvalue test.	86
2.4	Results of the multivariate Johansen max eigenvalue test.	86
2.5	Results of the trivariate Johansen constraint test for risk aversion.	88
2.6	Fitted short run parameters in the error correction model for Brexit, with tests on the short run dynamics.	89
2.7	Estimated autoregressive models for cointegration errors.	92
2.8	Bollinger Band gross trading profits for Brexit.	96
2.9	Cointegration tests for the smaller training period.	98
2.10	Ex-ante Bollinger Band gross trading profits for Brexit.	98
2.11	Financial assets and changes for the 2016 presidential election.	104
2.12	Results of the Phillips–Ouliaris Z_t and Z_α tests for cointegration.	105
2.13	Results of the bivariate Johansen max eigenvalue test.	105
2.14	Results of the multivariate Johansen max eigenvalue test.	105
2.15	Results of the trivariate Johansen constraint test for risk aversion.	106

2.16	Results of the Phillips–Ouliaris Z_t and Z_α tests for cointegration. R_t included.	107
2.17	Results of the multivariate Johansen max eigenvalue test. R_t included.	108
2.18	Fitted short run parameters in the error correction model for the Trump election and tests on the short run dynamics.	110
2.19	Estimated autoregressive models for cointegration errors.	113
2.20	Ex-Post Bollinger Band gross trading profits for the 2016 presidential election.	114
2.21	Cointegration tests for the smaller training period for the presidential election.	118
2.22	Results of cointegration tests for the Scottish independence referendum.	120
2.23	Results of linear cointegration tests for Scottish independence referendum where GBPUSD is shifted forward 30 minutes.	123
3.1	Results of the Phillips Ouliaris Z_t and Z_α tests for cointegration between $\log(\text{FTSE250}) - \log(\text{FTSE100})$ and the betting markets using daily data for the Brexit campaign period.	129
3.2	A regression of the first differences of $\log(\text{FTSE250}) - \log(\text{FTSE100})$ and the betting probability of Brexit.	131
3.3	Specifications for each political event.	153
3.4	Regression of Fama–French factor returns on ΔPB_t for political events in the testing window.	160
3.5	Regression of Fama–French factor returns on ΔPB_t for US presidential elections from the month prior to election day.	163
3.6	Results of the significance tests for γ	164
3.7	Results for weighted regressions, with the weight increasing by 5% each week.	169
3.8	Regressions of factor loadings γ against stock characteristics for different events.	172

Introduction

The research presented in this thesis is concerned with the interaction and behaviour of prediction and financial markets around political events. That political risk affects asset prices is an established concept. Similarly, the information content and superior forecasting ability of political prediction markets is a settled matter in the literature. Could the ‘wisdom of the crowds’ from both these markets concerning political risk be used to derive pricing relationships between these two very different markets? Does one market outperform the other in terms of reflecting information about an election? Also, are there any violations of market efficiency around political events and can they lead to any realistic profit opportunities? This thesis aims to answer these and other questions.

The period from 2016 onwards provided several opportunities in the Anglo-Saxon world to study financial and election markets around political events. The decision in June 2016 by the UK to leave the EU caused a political earthquake. This was followed a few months later by the shock result of the 2016 US presidential election that put Donald Trump in the white house. The years that followed were the most politically turbulent in the western world for a generation. Following the Brexit vote in 2016, political chaos ensued in the UK. There were general elections in 2017 and 2019 and, as this introduction was written, the ruling Conservative party was in the process of defenestrating their third prime minister in six years. The apparently baseless and brazen attempt of Donald Trump to overturn the result of the 2020 US presidential election, and the riots at the Capitol on January 6th that followed, tested the US Constitution to its limits. These and other events are studied in the context of testing the ideas and models in this thesis.

The thesis is organised as follows. The remainder of this introduction is taken up by reviewing areas of the existing literature relevant to the topic of the main chapters. Chapter 1 is a study on market behaviour, both political and financial, on the night of the EU referendum. Using a Bayesian electoral forecast model that updates in realtime, we demonstrate a major violation of semi-strong market

efficiency on the night after the vote. The electoral model used has been published in Auld and Linton (2019). The model is an improvement upon that of Wu et al. (2017). This paper studied the same question and concerned British Pound price action on that night. The flexibility of the prior in our model allows us to demonstrate that, without a shadow of a doubt, participants in both markets behaved irrationally as the night unfolded. This extends the existing finding that there was a ‘bubble’ in opinion that was present in the weeks leading up to the plebiscite. The analysis in Chapter 1 is also extended from GBPUSD to the betting markets. We show that the prediction markets were slightly more efficient at reflecting the referendum results than the currency market. This effect is on the order of a small number of minutes.

Chapter 2 presents a pricing model linking prices in prediction markets to financial assets that applies in the very particular circumstance of the overnight session following an election. Contrary to most of the existing literature, the model is derived from economic assumptions and common pricing restrictions taken from the asset pricing literature. We find that under suitable conditions, election and financial markets will be cointegrated. Deviations from risk neutrality lead to the presence of a non-linear term relating to risk in the cointegrating relationship. Three recent political events are studied: The 2014 Scottish independence referendum, the 2016 Brexit referendum and the 2016 US presidential election. Strong support is found for two events (the Brexit election and the Trump win). We find that weak market efficiency broadly holds although there are violations of the order of minutes to tens of minutes. This is apparently caused by betting markets leading financial markets, repeating and extending the finding of chapter 1. This demonstrates a failure of weak market efficiency on a small timescale. A realistic ex-ante trading strategy is presented for Brexit that profits from these inefficiencies. However, the success is not repeated for the 2016 presidential election. This is due to an apparent deviation from risk neutrality that was not observed on the night of Brexit. The revealed risk preferences for this election are pleasing. ‘Risk-on’ assets and currencies demonstrate risk aversion, whereas safe haven currencies and the US Treasury future demonstrate risk seeking behaviour, benefiting from a relative ‘flight-to-quality’.

Chapter 3 studies the relationship between political and financial markets over the weeks and months preceding an election. Again a model is derived from economic conditions. In fact the overnight model of chapter 2 is a special case of this model, where there is no variation in factors affecting asset prices beyond those re-

lated to the political event. The resulting model yields a relationship between asset price returns and first differences in betting markets. This is no great surprise. A small number of studies have studied this relationship, albeit only for single events and only empirically. However, our model is based on economic principles and applies in a general setting. We extend this model to equities, using the ubiquitous 5 Fama–French factors to describe the residual variance of asset prices that are unrelated to the political event. The result is a 6-factor augmented Fama–French characteristic factor model, with an additional factor being described by the betting markets and related to political risk. We test the model on six elections. Strong support is found for four, mixed results for one, and no evidence for a single election. The conclusion for the latter event is that this event was not significantly informative for stock prices. We also find evidence that betting markets become more informative as the event approaches. This idea already has some support in the literature. An exponential weighting scheme, where observations closer to the election are weighted more heavily, improves and sharpens our results. Finally an inspection of the political factor loadings reveals some pleasing relationships between firm characteristics and political sensitivity. Internationalisation of revenue was a key explanatory factor. This is consistent with the hypothesis that firms that have a greater share of off-shore sales are more able to diversify domestic political risks. We also find geographical location and nationalisation risk under an opposition win also explain differences in political sensitivity.

We end the thesis with some concluding remarks and directions for future research.

Literature Review

This thesis is primarily interested in the relationships between, and behaviour of, prediction and financial markets, around political events. There are several areas of literature that are relevant. Chapter 1 is a study of semi-strong efficiency on the night of the EU referendum. There are relevant event studies for this and other elections in the literature. However, chapters 2 and 3 present and test economically derived pricing models of political and financial markets that apply in two different settings. There are a small number of studies in the literature that present linkages of the prices in the two types of market but they are few and far between. Examples include: Manasse et al. (2020), Hanna et al. (2021), Acker and Duck (2015) and Darby et al. (2019). All of these studies are restricted to a single political event and

are not considered in a general setting. Further, the majority simply consider the discovery of an empirical relationship and present no theories based on economic principles. The sole exception to this is Manasse et al. (2020), which again is applied to only a single event (Brexit). The assumptions behind their model are not meaningfully discussed or scrutinised with any great deal of depth. Chapters 2 and 3 present pricing theories derived from assumptions and common pricing restrictions from the asset pricing literature. They apply in a general setting and are tested on several events. The fact that there is no similar work in the literature speaks to the contribution of this thesis.

This chapter will review several areas of literature and is organised as follows. Section 1 reviews prediction markets and, in particular, political prediction markets. Section 2 covers literature on the effects of political risk on financial markets. Throughout this thesis a common theme, and indeed assumption of our models, is market efficiency. Section 3 reviews the conclusions of the literature for this topic. Section 4 briefly reviews Fama–French equity factor models. These are relevant to the extension of the model derived in Chapter 3. Finally section 5 reviews studies of particular elections related to the events studied in this thesis.

Political prediction markets

Prediction markets are exchange traded financial markets for the purpose of trading on the outcome of events. The most common form is of the binary option variety where on expiry a contract pays out a fixed amount, or nothing, according to the outcome of an event. Contracts for differences also exist that are used to bet on, for example, the vote share of a political party in an election. Prediction markets include the Hollywood Stock Exchange which concerns films and other movie related events, Intrade.com where contracts are listed related to economic and current events, and the Policy Analysis Market, opened by the US Department of Defense in 2003. The latter swiftly closed following a backlash against the possibility of profiting from, and the incentivising of, terrorist events (Hanson et al., 2005). Corporate bodies have also run internal prediction markets. Notable examples include Google using them to better forecast product launches (Brodkin (2008)) and Eli Lilly to predict which drug targets had the best chance of clearing clinical trials (Polgreen et al., 2007). Election, or political, markets are prediction markets that are based on the outcome of elections. There is evidence that election betting occurred on Wall Street as long ago as 1884 (Rhode and Strumpf, 2004).

Modern examples of electronic election markets include University of Iowa’s Iowa Electronic Markets, introduced for the 1988 US presidential election, the University of British Columbia’s UBC Election Stock Market (now superseded by the Sauder School of Business Prediction Markets) and the Betfair Exchange, prices for which are used in this thesis.

There is a plethora of research on the accuracy of prediction, and election markets. These markets are a type of crowd sourcing whereby information from a large number of participants is aggregated into a single metric (the price). Accuracy should follow from the efficient market hypothesis, whereby assets prices immediately and fully discount all available information. The marginal trader hypothesis also postulates that prices will be set accurately even in the presence of bias in the majority of traders. ‘There will always be individuals seeking out places where the crowd is wrong’, Mann (2016). For a recent review of the field see Horn et al. (2014) and for political markets see Graefe (2016). We discuss notable papers in the field below.

It is intuitive to expect the clearing price of a prediction market to be equal to some weighted average of market participant’s expectations or perhaps some representative investor’s expectation. Two papers that provide theoretical underpinnings for this interpretation are Gjerstad (2004) and Wolfers and Zitzewitz (2006). It is found that both dispersion of beliefs amongst participants as well as risk aversion are required to make this conclusion. Wolfers and Zitzewitz (2006) also show that data from both American football and consumer sentiment markets are consistent with the idea that the market price is a good estimator of the mean belief. Page and Clemen (2013) explore the idea that the forecasting ability of prediction markets is negatively correlated with time to expiry of the market. They conclude that ‘Prediction markets are reasonably well calibrated when time to expiration is relatively short’, but for events far into the future ‘prices of long-term prediction markets are systematically biased towards 50%’.

Studies, many of which are based on the Iowa and UBC markets, have demonstrated the remarkable accuracy of forecasts from election markets. There is a consensus in the literature that political markets outperform other methods including polling and expert predictions. Two early papers by Forsythe et al. demonstrate the outperformance of prediction markets when comparing final prices with final polling numbers for vote shares. This is despite the presence of judgement bias amongst traders, which they do detect. Forsythe et al. (1992) studies the Iowa market for the 1988 US presidential election. They find that ‘the market worked

extremely well, dominating opinion polls’. By looking at the positions of different constituent groups they find that traders tended to place bets on their preferred candidates (indicating judgement bias). However, they conclude that prices are indeed set by the marginal trader. The implication being that even if a majority of participants are irrational and have misspecified beliefs, the existence of a small group of unbiased traders, perhaps arbitrageurs, keep prices in line. Oliven (2004) considers whether the Iowa market is skewed by biased participants. They conclude that market-making traders are more rational than price takers, the implication being that arbitrageurs are indeed making prices efficient. The analysis and conclusions are repeated for the 1993 Canadian federal election in Forsythe et al. (1995). This paper also identifies key campaign events. These include statements by the then current Prime Minister, a televised debate and the release of opinion polls. Both these studies focus on shorter-term market predictions. A latter study, Joyce E. Berg and Rietz (2008), extends the analysis to longer-term forecasts. This paper compares vote share prices from the Iowa market with 964 polls over the five US presidential elections from 1988 to 2004. They find that the prediction market is closer than the polls 74% of the time. The average error in vote share for presidential candidates in the final five days of polling is 1.2% versus 1.64% for polls. Further ‘the market significantly outperforms the polls in every election when forecasting more than 100 days in advance’. This provides evidence that election markets are not only accurate at times close to a vote but have superior explanatory power months from an election.

Challenges to political markets include the possibility of manipulation and large scale biases backed up by participants trading large amounts of capital possibly overwhelming marginal trader effects. Berg and Rietz (2006) provide evidence of political campaigns attempt to manipulate prediction market prices. There is also evidence that intrade prices were manipulated in the 2012 US presidential election for pecuniary benefit, Rothschild and Sethi (2016) and Goodell et al. (2015). However, the profit incentive can actually increase the accuracy of political markets. Hanson et al. (2006) shows that, in the Iowa market, distortion of prices can create the incentive to bet against the manipulator.

Despite the superior performance of prediction markets to forecast the outcome of elections, there have been two fairly recent notable failures. In 2016, betting markets seemingly failed to predict the UK voting to leave the European Union. Odds for leaving the EU were implied at around 25% during the final week of the campaign and bottomed out around 10% just after polls closed. This was despite

tight polling showing only a small lead for ‘Remain’ that was within the margin of error of most polls. Either the market mechanism appears to have failed to aggregate the ‘wisdom of the crowd’ for this event, or that wisdom was in fact biased. There is evidence that markets may have been distorted by a small number of market participants committing overwhelming capital to bets on Remain. For instance, Ladbrokes, a UK bookmaker, reported that most bets were for ‘Leave’ but were in small size only, Economist (2016). The much larger currency markets also seemingly appeared to fail to forecast the possibility of Brexit. This can be explained by those markets being dominated by traders who both had a preference for Remain and suffered from judgement bias (the same constituency that was placing large trades against Brexit on the betting platforms). One other hypothesis is that mass bias and a lack of diversity of opinion created a self reinforcing feedback loop, failing to create dispersion of beliefs required to make markets efficient. Wu et al. (2017) conclude that for Brexit there was ‘a herding of market participants away from fundamental value’. Fry and Brint (2017) also confirm the view that there was a ‘bubble’ in opinion for remain. They use a novel model of financial bubbles applied to polling data to make this conclusion. It has also been hypothesised that the widespread demonstration of accuracy of prediction markets prior to 2016 may have added to this reinforcement of beliefs for both the Brexit referendum and the US election that year, Gelman and Rothschild (2016). The public, commentators and bettors may have taken the implied odds as a fact, not updating beliefs sufficiently with outside information. For the presidential election, both the polls and betting markets failed to countenance a win by Trump, with betting markets putting the odds of such an outcome below 20%.¹ It appears prediction markets failed to see beyond the bad information in polls, and again there was an ‘echo chamber’ of self-reinforcement, Diepenbrock (2016).

Despite the high profile failures of 2016 the literature and weight of evidence strongly favours the reliability and superiority of betting markets over other methods. There is also a more prosaic explanation for the apparent failures seen in 2016. This is that over the many years that prediction markets have been in operation we would expect to see many such cases where events with odds of 10-25% come to pass. That two happened in the same year may just be down to chance.

¹The widely followed Nate Silver predicted the odds of Trump winning at 35%. Other experts though had odds more in line with those implied by prices in the markets.

Election risk and financial markets

The financial market implications of elections and political risk are well documented. Changes or the risk of changes in the composition of government naturally brings about changes in policy. These may relate to taxes, trade or welfare that often have direct implications for the fortunes of economies, companies, currencies and other asset prices. There are many studies that demonstrate either the effects on financial markets of election campaign periods or results. Studies most relevant to this thesis are outlined below.

One of the earliest papers that demonstrated association between elections polls and asset prices is Gemmill (1992). This studied the campaign period of the 1987 UK general election. In the paper the author derives the probability of a Conservative party win from polling data. They find that the 'FTSE100 index was very closely related to the probability of a Conservative win'. Further, in the final two weeks before the voteshare options prices showed large increases in implied volatility. This was particularly the case for two nationalisation targets (of the opposition Labour party). The increase in option volatility was despite the polls showing an improving chance of a Conservative win, apparently violating market efficiency.

We find three large multi-country studies of asset price volatility around elections. Bialkowski et al. (2008) studies 134 elections in 27 OECD countries from 1980 to 2004. Using a GARCH methodology they find that the relevant national stock exchange index volatility can easily double during the week after the election. Apparently 'investors are surprised by the election outcome'. They find several metrics relating to elections that are explanatory of the magnitude of the increase in equity volatility. These include closeness of the result, lack of compulsory voting, change in political orientation of the government and the size of, or lack of, the majority of the ruling party in any national legislature. Kelly et al. (2016) find that this uncertainty is priced in the options market. They analyse data from prices for options on either the national index, or an ETF tracking that index, for a sample of 20 countries. This is for various time periods in the range 1990-2012. They find that prices and implied volatility are higher for options that span elections. They also document spillover effects from the election country to other international markets. Pantzalis et al. (2000) is another large multi country study. This paper finds significantly positive returns two weeks prior to election dates for elections in 33 countries between 1974 and 1995. The conclusion is that as election uncertainty is resolved, prices respond positively. The largest abnormal returns are found in

two cases. The first is for less free countries where the election was won by the opposition, perhaps as the market prices those countries becoming more open. The second case is when the election is called early and the incumbent loses. Two later papers come to the opposite conclusion for US presidential elections. Goodell and Bodey (2012) consider how the Graham price to earnings (P/E) of the S&P500 index stocks, a valuation metric as well as a measure of consumer sentiment, changes during the campaign periods of US presidential elections. They find the measure *worsens* as the winner becomes clearer, according to odds seen on the Iowa market (that is, as uncertainty reduces). They conclude that for the US, arguably a very free country, ‘during presidential election seasons, the market discovers its distaste for the economic policies of the likely winner’. This is ‘consistent with efforts to ingratiate voters acting to dis-ingratiate the market’. The analysis is extended in Goodell and Vähämaa (2013) for the five presidential elections from 1992 to 1998. They consider the effects on the VIX, a measure of implied volatility of options on the S&P500 expiring in under one month. They find the VIX is positively associated with positive changes in the probability of the winner. This ‘indicates that the presidential election process engenders market anxiety as investors form and revise their expectations regarding future macroeconomic policy.’

Brander (1991) and Bernhard and Leblang (2006) study the 1988 Canadian election. This was shortly after the implementation of the Canada-US FTA.² The FTA was widely expected to increase trade between the two countries, being positive for the stock market. However, the opposition Liberal party was opposed to the agreement. Their leader, John Turner, threatened to ‘tear it up’ if they won. Prices on the Toronto Stock Exchange (TSE) were found to be significantly positively related to polling numbers for the Conservative party during the campaign period. Similarly for Australian elections from 1998 to 2010 Smales (2016) finds that polling *uncertainty* is positively related to implied volatility in core financial markets. The greatest effects are seen in base materials stocks, likely due to the introduction of the Mineral Rent Resource Tax, a surcharge on company profits earned from mining non-renewable resources.

Several studies demonstrate the financial market implications of the decision by the UK to leave the European Union in 2016. Aristeidis and Elias (2018) study the effects of Brexit on 43 developed and emerging stock markets. Using bivariate copulas with time-varying correlations they divide the sample into three peri-

²Following the inclusion of Mexico in 1994 the CUSFTA became the North American Free Trade Association in 1994.

ods: pre-referendum, post-referendum and after the triggering of Article 50. They demonstrate ‘immediate financial contagion’. However, the ‘shock and uncertainty were very limited’ as most markets bounced back within a few days. Alkhatib and Harasheh (2018) consider the effects of the Brexit vote on ETFs listed on the London market. The paper calculates abnormal returns for the 10 days before the vote, the day of the vote, and ten days after the vote. Contrary to the conclusions of Aristeidis and Elias (2018), they find significant positive abnormal equity returns on the day of the vote, as well as the period afterwards. This may be due to differences in the size of the event windows used in the two different studies. Alkhatib and Harasheh (2018) also find similar positive abnormal returns for a US Treasury ETF, emerging markets and gold. They attribute the performance of these assets to effective diversification and hedging strategies. Sultonov and Jehan (2018) study the effects of Brexit and 2016 US presidential effects on the Japanese Yen and stock market. They use three 4.5 month periods. The first is pre-EU referendum. The second is the 4.5 months between the EU-referendum and the presidential election and the third is the 4.5 months following Trump’s win. The individual returns series of the Yen and Nikkei are shown to be statistically significantly different in the three periods. Using a dynamic conditional model, the authors also observe differences in the series correlations for the different periods too. It appears that both the surprise political results of 2016 had profound effects on the financial markets of Japan. Caporale et al. (2018) use long memory methods to investigate whether there were any persistent effects of the UK’s decision to leave the EU. They consider the implied volatility of both the FTSE100 index and the British Pound relative to other major currencies. Using an event study methodology and splitting the samples into pre- and post-referendum periods they find an increase in the degree of persistence of volatility of GBP versus all major currencies except the Japanese Yen.

The efficient market hypothesis

The efficient market hypothesis (EMH) by Fama (1965) states that the prices of financial assets immediately reflect all available information. Therefore investors cannot make above average returns, except by chance. There are various forms of the hypothesis. In the weak form, financial prices instantaneously reflect all *market* information; in the semi-strong form, prices instantaneously discount all *publicly available* information; in the strong form, prices instantaneously reflect all

information *both public and private*, including privileged information available to insiders. Many authors, e.g. Malkiel (2003), argue that the EMH does not imply that pricing is perfect or that mispricings never occur, just that mispricings are random and it is not possible to systematically profit from them in advance.

There are various behavioural explanations that attempt to explain why the EMH may not hold. For a recent comprehensive review see Huang et al. (2016). One such theory is that of investor inattention. See Hirshleifer et al. (2013), Hirshleifer et al. (2009) and Hou et al. (2009) for examples, as well as DellaVigna and Pollet (2009) where the authors show that earnings announcements on a Friday take longer for the market to react to. Other behavioural explanations include anchoring and systematic overconfidence. Another idea presented in Caballero and Simsek (2016) postulates that a study of any anomalies of the EMH require an analysis of the presence or absence of any arbitrage process that may exist to bring prices rapidly back to the “correct” value. There have been many opinions and studies published on the EMH and no consensus exists as to its validity. For a recent review, see Gabriella Titan (2015).

Fama–French factor models

The mainstay of equity asset pricing models is the capital asset pricing model (CAPM) of William Sharpe (1964) and John Lintner (1965). Indeed Sharpe received a Nobel Prize for the model in 1990. The model explains the cross-sectional returns of equity prices. It has also been widely used to evaluate the performance of managed portfolios and estimate firm’s cost of capital, Fama and French (2004). However, two empirical findings were found that were inconsistent with CAPM. These are the size effect (Keim (1983), Banz (1981)) and the value effect (Rosenberg et al. (1985)), both of which have been found to be explanatory factors in the cross section of company price returns. The 3 factor model of Fama and French (1993) augmented CAPM with factors that proxy size and value. Subsequent evidence found that this model may be incomplete. It was found that growth in investment affected average returns adversely, Titman et al. (2004) and Anderson and Garcia-Feijóo (2006). Separately Novy-Marx (2013) documented that profitability was strongly correlated with average returns. Fama and French extended the 3 factor model to include proxies for these two effect, resulting in a 5 factor model, Fama and French (2015).

We find two studies in the literature that directly relate political prediction market prices to the ideas of CAPM and factor models. These are both studies of the 2014 Scottish independence referendum. Acker and Duck (2015) find that the residuals of an estimated CAPM model are significantly positively related to several proxies for the probability of a vote to remain as part of the UK. One of the proxies is a weighted sum of the Betfair exchange odds for ‘No’. The authors believe that uncertainty with respect to the independence referendum was indeed reflected in market prices. Darby et al. (2019) is a similar study of equities listed on the LSE that were headquartered in Scotland. They find that uncertainty betas help predict cross-sectional returns. They again conclude that ‘heightened political uncertainty was priced during the period surrounding the referendum’.

Events

There are several studies in the literature that relate to the elections we apply our models to in this thesis. Many of these relate to Brexit. However, we find only one example of an economically derived relationship between prediction markets and asset prices. This is in Manasse et al. (2020). A simple portfolio model is built for currencies. This implies that currencies are cointegrated with betting prices for the period running up to the 2016 UK Brexit referendum. Under risk neutrality they find a linear cointegrating relationship but risk aversion leads to the presence of a risk factor related to uncertainty. This leads to a non-linear term of the betting market appearing in the cointegrating relationship. We build a similar model in chapter 2 but which only applies to the overnight session following an election. Manasse et al. (2020) find that currency and Betfair data for Brexit does not reject a null hypothesis of their model. However, we do not believe the assumptions behind their model are valid. For it to be so, one would have to believe that the only determinant of the GBPUSD price in the weeks and months preceding the Brexit vote is the result of that vote. We believe this is not plausible. News and information beyond that relating to the referendum, including US economic releases, are likely to affect the British Pound and United States Dollar exchange rate in the period under study.

Wu et al. (2017) investigate the real time response of the exchange rate to the announced vote outcomes on the night of the EU referendum. Their conclusion is that the ‘Brexit result could have been predicted with high confidence under realistic conditions’. Examining social and psychological factors as well as Betfair

data prior to the vote, the authors conclude that the mispricing ‘indicates both generic inefficiency and a specific inertia / durable bias in the market similar to herding during bubbles’. The paper also examines trading behaviour in the pound around the announcement of specific results.

Another study that used betting data in the campaign period for the Brexit referendum is Hill et al. (2019). They analyse both the sensitivity to changes in betting odds for leaving the EU as well as the return on the day after the result for UK firms. Using a cross-sectional regression they then identify factors that affected the political sensitivity. They find that the most explanatory factor for the political risk of a firm is internationalisation, particularly in relation to activities outside the European Union. The relationship between internationalisation and referendum sensitivity is negative. This is consistent with a hypothesis that such firms are in a better position to diversify the domestic policy risks of Brexit.

Hanna et al. (2021) also use Betfair data to analysis how changes in the betting odds for ‘Leave’ influence financial markets for the Brexit referendum. They consider the period from January 2016 to the date the referendum was resolved (the early hours of June 24). Using high frequency data for trades on the Betfair exchange they regress short-term returns of GBPUSD and major UK and European stock indices on changes in betting prices during stock market opening hours. They find that changes in the odds for leave cause prices of UK equities and the pound to fall in the following 5 minutes. They also find some spillover effects into EU equity prices. However, the slopes of the regression appear somewhat small when compared to the sensitivities estimated in other studies, including in this thesis. They show that a 1% increase in the probability of the UK leaving the EU causes a 0.004% drop in the price of UK stocks and a 0.006% depreciation of the pound relative to the dollar. This suggests a difference of only 0.4% and 0.6% in UK equity prices and the currency between a ‘Leave’ and ‘Remain’ outcome. The magnitude of these effects are negligible, particularly when compared to the actual market reaction following the shock vote to leave the EU.

The authors Clark and Amen have devised a method to estimate the probability of the outcome of an election using the shape of implied volatility surfaces from foreign exchange options. This is effected by fitting a bi-modal mixture distribution related to two possible outcomes. In Clark and Amen (2017), GBPUSD options are used to estimate the odds of referendum outcomes for both the 2014 Scottish independence referendum and Brexit. For the former they find ‘asset price distributions which are consistent with the observed post-referendum spot price data

after the event'. For Brexit they find evidence of significant tail mass in the volatility surface related to the possibility of a vote to leave the EU. Clark and Amen (2019) apply the method to other recent elections. These are the 2016 US presidential election, using USDMXN, the 2017 French election, using EURUSD, and the 2017 UK general election, using GBPUSD. For the presidential election, their calculated probabilities of a Trump win are within the range 24-55%, higher (and arguably better) than betting markets and experts. They do not find evidence for the mixture distribution for the French election, and for the UK general election they find some probability mass in the 1.25-1.275 range for GBPUSD. They argue this is consistent with the immediate fall of the pound to 1.275 after the surprise loss of the Conservative party's majority.³

Wall et al. (2017) uses prediction market data to find significant events during the 2015 campaign for the Scottish independence referendum. Betfair data is used to control for polling shocks and isolate campaign effects. They find that the second leader's debate, between the SNP's leader Alex Salmond and Alistair Darling, the head of the 'No' campaign, was the most significant event of the period. The debate initiated a surge in support for 'Yes'.

Goodell et al. (2015) provide an interesting study of betting markets during the 2012 US presidential election. As betting odds are market prices, they are naturally integrated of order unity. They are also bounded between zero and one. This can lead to some mathematical complexity. Models of bounded Brownian motion do exist (see Carr (2017) and Taleb (2018) for an application to election forecasting). However, in this thesis we prefer to think of betting odds as standard random walks that cease to exist when they hit the boundary; that is, when the election result becomes known. Either way, Goodell et al. (2015) demonstrate that two betting markets for this election (Iowa and Intrade.com) are cointegrated. Further, they show that these markets are also cointegrated with the forecast of Nate Silver, arguably the most followed political pundit in the US, if not globally. They show causation of both Iowa and the expert to prices on Intrade. It appears the latter follows the former. They confirm the findings that both Intrade prices appear to have been manipulated and that betting markets were superior to polls at forecasting the election. The most surprising finding though is that the 'three series consistently differed in the degree of optimism in an Obama victory'. This suggests some degree of segmentation between the users of each betting market.

³This author has his doubts about this conclusion given the small size of the move in GBPUSD, and the fact that it was very short lived.

This is perhaps not surprising as Iowa is an academic platform with a trading cap whereas Intrade is a commercial betting venture.

Chapter 1

Semi-strong market efficiency on the night of the EU referendum

Note to examiners

An early version of the electoral model in this chapter was submitted as a dissertation for the MPhil in Economics Research at the University of Cambridge in 2017. A later version was published with Oliver Linton, Professor of Political Economy and Director of Research, Faculty of Economics, University of Cambridge, in Auld and Linton (2019). Professor Linton largely acted in a supervisory role; 90% of that paper, and the work in this chapter is my own work. The material below is a later, improved, version of earlier work. Most significantly, it is corrected for a major data error present in the Mphil thesis and Auld and Linton (2019). Earlier versions incorrectly assumed the data supplied by Betfair had timestamps in BST when in fact they were in GMT. This significantly changed the conclusions of the work. All work and data presented below is converted to BST timestamps correctly.

1.1 Introduction

Were currency and prediction markets efficient overnight on 24 June 2016 as the results of the United Kingdom European Union membership referendum were announced?

This question is important as the EU referendum was one of the great political shocks of 2016. The results of the vote itself provide for a unique period in market history for which both financial and prediction market efficiency can be studied.

The night is a special event for a number of reasons. Firstly, referendums are rare events with no similar votes in history for market participants to base expectations.¹ There was also a strong prior belief that the UK would vote to remain in the European Union. This provided fertile ground for inefficiencies and behavioural biases to arise. Secondly, the EU referendum results were the only information to affect the market during the hours of the night. There were 382 different voting areas and results were announced and widely distributed at different times. This represented a drip-feeding of information to the market for a period of a few hours. Thirdly, there are two markets to study: a prediction market in the Betfair betting market and the pound dollar currency market.

The efficient market hypothesis (EMH) holds that financial markets immediately reflect all available information in prices. If this is true, investors cannot receive above market returns except by chance. The weakest form of the hypothesis relates only to historical price information. Opinion is split on whether this form holds. Stronger forms of the EMH relating to both fundamental (semi-strong) and private information (strong) also exist. Most studies conclude that the stronger forms of the EMH do not hold. For the night of the EU referendum, one existing working paper (Wu et al. (2017)) concludes that the pound market was slow to reflect the information contained in the vote results and hence the EMH in the semi-strong form did not hold. Regarding prediction markets, there is a consensus in the literature that prediction markets provide better estimates of future events than experts, and that the predictions of such markets are useful in a variety of situations.

This chapter makes a number of contributions: This is the first high frequency study in the literature comparing a prediction market with a financial market. We agree with Wu et al. (2017) that the EMH in semi-strong form did not hold in the currency market during the night of the referendum. The use of a flexible Bayesian prior suggests that without a shadow of a doubt market participants behaved irrationally on the night. However we pull and twist the prior we cannot come up with a description of prior beliefs that match price behaviour and are in the slightest way plausible. This finding is an extension of the hypothesis that there was a ‘bubble’ in opinion polling for ‘Remain’ present prior to the day of the vote.

¹There have only been two other UK-wide referendums. The first, the European Communities membership referendum held in 1975, would be of little use for inferring voting patterns today. The other, on an unrelated subject, was the Alternative Vote referendum in 2011 and had a turnout of only 42.2%, opposed to a typical figure of 60–70% for general elections

We also demonstrate that market failure was present in the the prediction market as well as in the British Pound. Further, we show that the Betfair market was slightly more efficient than the sterling futures market, which provides some support for the view that prediction markets yield useful predictions. Small sample inference is required to predict Brexit early on in the night of the vote and we improve upon earlier prediction methods by using a rigorous Bayesian approach that is valid for small samples. Finally, we show that investors both appeared to have overconfident ex-ante beliefs, and that they updated those beliefs inconsistently as information arrived in the form of the results of the referendum.

The remaining sections of this chapter are organized as follows. In section 2 we present the electoral model updating methodology, which employs Bayesian machine learning. section 3 presents a simple theoretical approach to extracting realtime implied probabilities from the two markets (Betfair contracts and Sterling futures) under standard economic assumptions. In Section 4 we present our empirical results. Section 5 concludes.

1.2 Electoral model

In this section we present a Bayesian model for calculating a probability of Brexit that updates throughout the night as the results of the vote are announced. It is based on the update of a joint copula prior distribution over the area vote results. A summary of the model is presented in Appendix A.

1.2.1 Choice of approach

We seek a model that provides a probability of Brexit using public information available at times throughout the night. This information includes the vote-share and turnout of different voting areas that have announced. We model the problem as follows:

There are n constituencies with fixed sizes which we label $s_1, s_2, \dots, s_n \in \mathbb{N}$. Suppose $p_i, q_i \in [0, 1]$ are the proportion of voters in favour of leaving the EU and the turnout percentage in constituency i , ordered by time. Then the proportion of the national vote is:

$$p_N = \frac{\sum_{i=1}^n p_i q_i s_i}{\sum_{i=1}^n q_i s_i} \quad (1.2.1)$$

and the event of leaving the EU occurs when $p_N > \frac{1}{2}$. The probability of Brexit at time t , say $\Pr(\text{BREXIT})_t$, is thus $\Pr(p_N > \frac{1}{2} | p_1, q_1, \dots, p_t, q_t)$. As p_N is based on the results announced so far, $(p_1, q_1, \dots, p_t, q_t)$, as well as those yet to be announced, $(p_{t+1}, q_{t+1}, \dots, p_n, q_n)$, this probability can be calculated from the conditional distribution $(p_{t+1}, q_{t+1}, \dots, p_n, q_n) | (p_1, q_1, \dots, p_t, q_t)$.

If t is large one could attempt to form an asymptotic approximation to the conditional distribution above. This would have the advantage that it relies on as few assumptions as possible. Wu et al. (2017) take this approach. They form a linear regression of results announced so far on the expectation of those results, and use the resulting linear equation to predict unannounced results. $\Pr(\text{BREXIT})_t$ is then computed by sampling from the asymptotic Gaussian distribution of the OLS estimators and computing $(p_{t+1}, q_{t+1}, \dots, p_n, q_n)$ and thus p_N for each realised sample. Their approach is discussed extensively in Appendix B. The OLS estimators will only have a Gaussian distributions in large samples and the standard error estimator relies on having correctly specified heteroskedasticity. Robust errors can be used but they are likely to be severely biased in small samples. Unfortunately this is precisely the situation we want to apply the model: when there are few vote results available during the early part of the referendum night.

We use a Bayesian framework whereby we construct a joint prior distribution for $(p_1, q_1, \dots, p_n, q_n)$ valid before any results have been announced. It is then a simple mathematical matter to form $(p_{t+1}, q_{t+1}, \dots, p_n, q_n) | (p_1, q_1, \dots, p_t, q_t)$ from the conditional distribution of the joint prior. This is a superior approach in some ways to a large sample method as distributions are exact no matter how small the data available. Of course it relies on having a defensible valid prior. However, we argue that, given the considerable information available prior to the announcement of results on voting intentions and the demographic make up of different areas, using a prior is actually more appropriate in this application. It is more rational to make ex-ante assumptions about the distribution of $(p_1, q_1, \dots, p_n, q_n)$ than make to make no assumptions at all.

1.2.2 Ex-Ante public information

Prior to the vote there was widespread polling data on the voting intentions. For example, a poll was conducted by YouGov which was published shortly after voting closed at 10pm (YouGov (2016)). This showed the vote-share for Leave at 48.38% with a standard error of 3%. Further, the psephologist Professor Chris Hanretty


published a blog post before the vote that gave expectations of the vote-share for very nearly all constituencies in the referendum. We will make extensive use of the dataset available on this blog.

Professor Chris Hanretty is a Professor of Politics at Royal Holloway, University of London. Professor Hanretty not only conducts academic research but works with a large number of media organisations to help anticipate the likely results of elections and consults for a variety of companies. Prior to the vote he published a blog entitled “The EU referendum: what to expect on the night” (Hanretty, 2016). The blog post contained expectations and 90% confidence intervals (CIs) for vote-share based upon publicly available data for all but 4 of the 382 voting areas. This work was covered extensively by the media, including Bloomberg, and was also followed up in a piece Professor Hanretty wrote for the *Observer* called “How the EU Referendum Result Will Emerge in the Hours after Polls Close”. Stating that market participants were aware of these priors is not a controversial comment.

The priors were based on a panel data analysis of the British Electoral Study (BES) from 2015 and demographic results at the local authority level. Firstly the priors are calculated directly from the BES, and secondly, a uniform swing is applied to each area to bring the results in line with polling information available on the date of Hanretty’s publication (7 June). The data available on the blog has not been subject to peer review. However, Professor Hanretty has published an article with other co-authors that compares methods for estimating constituency opinion from national survey samples using the General Election of 2010 as an example (Hanretty et al. (2016)). The methods that were peer reviewed in this paper were precisely those used to generate the results on the blog post. Further, a regression of the realised result in the referendum on Hanretty’s predictions yields an R^2 of 0.859 and a correlation of 0.927. Figure 1.1 shows the corresponding regression and scatter plot for the EU referendum.

1.2.3 Formation of Prior

Hanretty provided both expectations and confidence intervals for the area vote-shares. This effectively gives us the first two moments of the marginal distribution for the variables (p_1, \dots, p_n) . As such, we can fit any 2-parameter family of distributions to this information and have an exact form for at least the marginal distributions for half of the variables in the prior. Forming a reasonable marginal prior for the turnout variables (q_1, \dots, q_n) is an easier problem. There is plenty



The red line is the best linear predictor.

Figure 1.1: Realised vote-share versus Hanretty expectation for all voting areas predicted in the EU referendum.

of information about voter turnouts from previous elections. The exact choices of marginal forms and information used is discussed further in sections 1.2.5 and 1.2.8.

Write the $2n$ variables being the values of vote-share and turnout as the vector $r = (p_1, \dots, p_n, q_1, \dots, q_n)^\top$, ordered chronologically. Given marginal CDFs for the unknown variables $F_1(r_1), \dots, F_{2n}(r_{2n})$ the question is now how to form a joint prior. Fortunately Sklar's Theorem provides a prescription to do just this. The theorem states that any multivariate distribution which has continuous marginals (which is natural to assume in this application) can be expressed in terms of a unique copula. The copula defines the dependence structure between the variables and is defined as:

$$C(F_1(t_1), \dots, F_{2n}(t_{2n})) = \Pr(r_1 < t_1, \dots, r_{2n} < t_{2n})$$

In short the copula is a joint CDF of the quantiles of the random variables. In defining our copula there are a variety of ways to proceed. For instance, we may wish to choose a uniform copula which can be argued to make no assumptions on the dependence of at least the quantiles of the unknowns in our problem. However, it is natural to assume that there will be a high degree of association between both the vote-shares and turnouts of different voting constituencies in the EU referendum and we actually would prefer to assume some structure. Further, we are motivated by computational tractability. Earlier we noted that we will require conditional distributions of the joint prior. Given the high dimensional nature of our problem (>750 variables) the use of a sampling method to evaluate the conditional copula would be expensive and an analytic form for the conditional copula is sought. Luckily, there is an off the shelf multivariate copula that provides both a natural way to encode any ex-ante expected association between variables and has a simple conditional form: the Gaussian copula. This is in fact a special case of the students-t family of copulas (being that obtained in the limit of the degrees of freedom parameter, ν , going to infinity). A finite ν would produce a copula with fatter tails and greater tail association. However, the Gaussian copula is in fact the copula that yields the highest entropy of all copulas that have the first two moments of the copula specified. Thus, if all we want to prescribe the first two moments only of the quantiles of the variables, then this copula provides the least additional information amongst all copulas that we could choose.

The Gaussian copula with correlation matrix Σ_0 is defined as:

$$\Pr(r_1 < t_1, \dots, r_{2n} < t_{2n}) = \Phi_{\Sigma_0}(\Phi^{-1}(F_1(t_1)), \dots, \Phi^{-1}(F_{2n}(t_{2n})))$$

where Φ_{Σ} is the CDF of a multivariate normal distribution with zero mean and correlation matrix Σ . There is the possibility of choosing a hierarchical form for the prior by choosing a randomly specified correlation matrix. For instance the Inverse-Wishart conjugate family of prior distributions is used in Bayesian inference, being conjugate to the normal likelihood. However, we would prefer to include information about the likely association between our variables in our problem. This could be specified by a prior over a covariance matrix in some random way but, again, computational tractability would be compromised for no benefit over using a fixed covariance matrix. We proceed by fixing the correlation matrix to include

two factors. This is parameterized by a specifying constant correlations between vote-shares, between turnouts and between vote-share and turnout. This leads to the following form for Σ_0 :

$$\Sigma_0 = \begin{pmatrix} \Sigma_p & \Sigma_{pq} \\ \Sigma_{pq}^T & \Sigma_q \end{pmatrix},$$

where

$$\begin{aligned} \Sigma_p &= (1 - \rho_p) I_n + (\rho_p) i_n i_n^\top \\ \Sigma_q &= (1 - \rho_q) I_n + (\rho_q) i_n i_n^\top \\ \Sigma_{pq} &= \rho_{pq} \times [(1 - (\rho_q \rho_p)) I_n + (\rho_q \rho_p) i_n i_n^\top]. \end{aligned}$$

The correlation matrices Σ_p , Σ_q , and Σ_{pq} are all of size $n \times n$ and represent the dependence between area vote-shares (p_1, \dots, p_n) , turnouts $(q_1, \dots, q_n)^\top$ and correlation between vote-share and turnouts, respectively. The relevant correlations are the same between areas and are ρ_p , ρ_q and $\rho_{pq} \in (-1, 1)$.

1.2.4 Prior probability of Brexit

Given specifications for the dependency Σ_0 and marginal distributions

$$\{F_{r_1}(r_1), \dots, F_{r_{2n}}(r_{2n})\}$$

the prior probability of Brexit is

$$\Pr(BREXIT)_0 = \Pr\left(\frac{\sum_{i=1}^n p_i q_i s_i}{\sum_{i=1}^n q_i s_i} > \frac{1}{2}\right), \quad (1.2.2)$$

where:

$$\begin{aligned} (p_1, \dots, p_n, q_1, \dots, q_n)^\top &= r \\ \Phi^{-1}(F_1(r_1)), \dots, \Phi^{-1}(F_{2n}(r_{2n})) &\sim N(0, \Sigma_0). \end{aligned}$$

As there is no analytical form for the integral in equation (1.2.2), a sampling method is required for evaluation. This is easy to achieve by sampling from the Gaussian multivariate distribution with the relevant correlation matrix, passing

those samples through $\Phi(\cdot)$ to give the quantiles, then passing those quantiles through the relevant inverse marginal distribution $F_r^{-1}(\cdot)$ to yield the collections of samples in the space of vote-shares and turnouts. The vote-share for Brexit p_N can be evaluated for each sample and the proportion with $p_N > 0.5$ gives the implied probability of Brexit.

1.2.5 Marginal calibration

Hanretty provided expectations and 90% CIs for the marginals. These CIs are implied by responses of the BES coupled with Local Authority demographic data. They do not take into account the uncertainty of the national vote. There may be an argument that the distributions that Hanretty supplies as part of his analysis of panel and local authority data are asymptotically normal. No such argument can be made once the uncertainty of the national vote-share is taken into account. We interpret the prior as an expression of a degree of belief about the possible values that vote-shares could take. As such, we are not constrained by the normal distribution.

To calculate the expectations of the marginals, suppose that μ_H is the vector of expectations provided by Hanretty. Given an expected level for the national average vote-share μ_N , then expectations for each area i , denoted μ_{p_i} , can be formed by applying a fixed uniform shift, α_N to μ_{H_i} , where μ_{H_i} is the expectation provided by Hanretty.²

We calibrate the marginal distributions by targeting a subjective level of the prior variance of the national vote-share σ_p^2 equal to a generous estimate of what we think it could be, say $\hat{\sigma}_p^2$. Either way, we add a constant variance σ_N^2 to each marginal variance and adjust σ_N^2 to achieve the result. We convert the CIs as if the marginal distribution were normal:

$$(\sigma_H^2)_{p_i} = ((90\% \text{ Confidence Interval})_i / (2 \times 1.645))^2,$$

$(\sigma_H^2)_i$ is now the unadjusted variance implied for the i 'th voting area. The national vote-share is uncertain and treating this variation as independent of the idiosyncratic variances implied by Hanretty's values leads us to add to each area variance a constant variance, say σ_N^2 ³:

²Note, Hanretty himself forms μ by applying a uniform shift to the priors he calculates from the BIS and census data to agree with polling data at the time of his publication

³We are not implying that area vote-shares are independent of the national vote-share, just that the idiosyncratic variation implied by Hanretty's study of survey respondents and local authority

$$\sigma_{p_i}^2 = (\sigma_H^2)_i + \sigma_N^2$$

The prior mean for p_N is below 50% so the part of the national vote-share distribution that lies above 50%, which is the model probability of Brexit, will be greater given higher variance σ_N^2 .

For turnout we base area level expectations, say ν , on historical general elections. However we choose the variances for each area to be equal and labelled by σ_ν^2 .

We consider the marginal distributions of the following types: Normal, Logit-Normal, Beta and Logit Student with Location and Scale. In each case the hyperparameters of the marginal distributions are set so that the mean and variance are equal to μ_{p_i} and $\sigma_{p_i}^2$ respectively, for each area i . This is outlined in the following sections.

1.2.5.1 Normal

$$r_i \sim N(\mu_i, \sigma_i^2)$$

$$F_i(r_i) = \Phi\left(\frac{r_i - \mu_i}{\sigma_i}\right)$$

This distribution is computationally cheap as it does not require the evaluation of the CDF or inverse. The prior is actually simply a multivariate Gaussian. However, a disadvantage is that it does not restrict the random variables p_i, q_i to $[0, 1]$.

1.2.5.2 Logit Normal

$$\text{logit}(r_i) \sim N(\mu_i, \sigma_i^2)$$

$$F_i(r_i) = \Phi\left(\frac{\text{logit}(r_i) - \mu_i}{\sigma_i}\right)$$

data is independent of the variation of the national vote-share number. Our assumptions are not even that strong; as we apply a uniform shift to variances we are simply implying that the *difference* in the variations of the marginal distributions of individual areas is the same as the difference implied by Hanretty's study.

We can apply the logit transform to the variables to convert them to the real line, and then assume a normal distribution under that transformation. Again, a multivariate normal distribution is implied for the transformed variables, as there is no analytic solution to the moments a sampling method could be used. However, we will use a simple transformation which was found in practice to make no difference to the results as follows:

$$\mu_i = \text{logit}^{-1}(\mu_i)$$

$$\sigma_i = \frac{[\text{logit}(\mu_i + \sigma_i) - \text{logit}(\mu_i - \sigma_i)]}{2}$$

Note that as the logit function is only symmetric around the value 0.5, the above transformations will not (quite) preserve the differences in expected values of vote-share. However, in practice the differences were not found to be meaningful.

1.2.5.3 Beta

$$r_i \sim \text{Beta}(\alpha_i, \beta_i)$$

$$F_i(r_i) = I_{r_i}(\alpha_i, \beta_i) = \int_0^{r_i} t^{\alpha_i-1} (1-t)^{\beta_i-1} dt$$

The Beta distribution is restricted to $[0, 1]$ and is parameterized by two shape parameters α and β . For a mean of μ_i and variance of σ_i^2 , the shape parameters will be equal to:

$$\alpha_i = \left(\frac{1 - \mu_i}{\sigma_i^2} - \frac{1}{\mu_i} \right) \mu_i^2$$

$$\beta_i = \alpha_i \left(\frac{1}{\mu_i} - 1 \right)$$

1.2.5.4 Logit Student with location and scale

$$\text{logit}(r_i) \sim t_\nu(\mu_i, \sigma_i^2)$$

$$F_i(r_i) = \frac{1}{2} + \left(\frac{\text{logit}(r_i) - \mu_i}{\sigma_i} \right) \Gamma\left(\frac{\nu+1}{2}\right) \times \frac{{}_2F_1\left(\frac{1}{2}, \frac{\nu+1}{2}, \frac{3}{2}, -\frac{(\text{logit}(r_i) - \mu_i)^2 / (\sigma_i^2 \times \nu)}{\sqrt{\pi \times \nu} \Gamma\left(\frac{\nu}{2}\right)}\right)}{\sqrt{\pi \times \nu} \Gamma\left(\frac{\nu}{2}\right)}$$

The Student's t-distribution has a higher kurtosis than the normal distribution and including it enables us to study priors with greater fourth moments for a given variance. The mean is simply μ whereas the variance, for $\nu > 2$, is $\frac{\nu}{\nu-2} \times \sigma$. We set the parameters in a similar way to the logit normal marginal, but with scale parameter σ_i

$$\sigma_i = \frac{\nu - 2}{\nu} \times \frac{[\text{logit}(\mu_i + \sigma_i) - \text{logit}(\mu_i - \sigma_i)]}{2}$$

The excess kurtosis is ∞ for $\nu \in (2, 4]$ and $\frac{6}{\nu-4}$ for $\nu > 4$. To explore the implications of infinite fourth moments (for $\text{logit}(r_i)$ not r_i) we set $\nu = 3$.

1.2.6 Update

Calculation of the conditional distribution is most easily done by a re-ordering of the variables,

$$\tilde{r} = (p_1, q_1, \dots, p_n, q_n)^\top.$$

Then

$$\Phi^{-1}(\tilde{F}_1(\tilde{r}_1)), \dots, \Phi^{-1}(\tilde{F}_{2n}(\tilde{r}_{2n})) \sim N(0, \tilde{\Sigma}_0),$$

where for generic x , \tilde{x} indicates a similar re-ordering of the rows and columns of x .

To calculate the conditional distribution of the remaining variables after t results have been announced, partition $\tilde{\Sigma}_0$ into four block matrices as follows:

$$\tilde{\Sigma}_0 = \begin{pmatrix} \tilde{\Sigma}_{t,t} & \tilde{\Sigma}_{t,\setminus t} \\ \tilde{\Sigma}_{\setminus t,t} & \tilde{\Sigma}_{\setminus t,\setminus t} \end{pmatrix},$$

where: $\tilde{\Sigma}_{t,t}$, $\tilde{\Sigma}_{t,\setminus t}$, $\tilde{\Sigma}_{\setminus t,t}$ and $\tilde{\Sigma}_{\setminus t,\setminus t}$, are $2t \times 2t$, $2t \times 2(n-t)$, $2(n-t) \times 2t$ and $2(n-t) \times 2(n-t)$ matrices respectively. The multivariate conditional Gaussian copula is also a Gaussian copula. Given the observations $p_1, q_1, \dots, p_t, q_t$, write the data as:

$$\tilde{x}_t = \left(\Phi^{-1}(\tilde{F}_1(\tilde{r}_1)), \dots, \Phi^{-1}(\tilde{F}_{2t}(\tilde{r}_{2t})) \right).$$

Then

$$\Phi^{-1}(\tilde{F}_{2t+1}(\tilde{r}_{2t+1})), \dots, \Phi^{-1}(\tilde{F}_{2n}(\tilde{r}_{2n})) \mid \tilde{r}_t \sim N(\tilde{\Pi}_{\setminus t}, \tilde{\Sigma}_{\setminus t})$$

where:

$$\begin{aligned}\tilde{\Pi}_{\setminus t} &= \tilde{\Sigma}_{\setminus t, t} \tilde{\Sigma}_{t, t}^{-1} \tilde{x}_t \\ \tilde{\Sigma}_{\setminus t} &= \tilde{\Sigma}_{\setminus t, \setminus r} - \tilde{\Sigma}_{\setminus t, t} \tilde{\Sigma}_{t, t}^{-1} \tilde{\Sigma}_{t, \setminus t}.\end{aligned}$$

As $p_1, \dots, p_t, q_1, \dots, q_t$ is now known, the model probability of Brexit can be computed via:

$$\begin{aligned}\Pr(\text{BREXIT})_t &= \\ \Pr \left(\sum_{i>t} \left[p_i - \frac{1}{2} \right] q_i s_i > \sum_{i\leq t} \left[\frac{1}{2} - p_i \right] q_i s_i \mid p_1, q_1, \dots, p_t, q_t \right). \quad (1.2.3)\end{aligned}$$

1.2.7 Model properties

An advantage of the model is that it provides for closed form updates to the posterior distributions of the parameters, as the Gaussian copula has a conditional distribution. This avoids the need for a Markov Chain Monte Carlo sampling technique to calculate the integral in equation (1.2.3). This would be particularly arduous given the large number (764) of variables involved. An alternative copula, with well understand and closed form conditional distributions, is the Student's t copula, Ding (2016). The conditional distribution is also a Student's t copula and will have fatter joint tails than the Gaussian.

The model describes the dependence structure of the unknowns in a parsimonious way with a relatively few parameters. This leads to limitations. A better description of the correlations of the variables could be found by using more factors, such as the demographic ones used by Hanretty to calculate the published marginal statistics. Heterogeneous correlation coefficients would likely follow. However, we postulate that our model will capture the main swing to the Leave vote. We will perform robustness tests using some limited forms of heterogeneous parameters to test this hypothesis.

We now comment on the expected qualitative impact of the model as parameters change. For purely independent vote-share results ($\rho_p, \rho_q = 0$), convergence will solely be due to the results as they come in and the distributions of the yet to be announced results will not be affected. For higher values of ρ_p , convergence will be

faster. It is expected that the value of ρ_p and the variance σ_N^2 will have the greatest effect on the speed at which predictions change. The effect of ρ_q and ρ_{pq} on the model probability are effectively second order. Given that we will be setting ρ_{pq} as negative and turnout was above expectations, there will be some small second order effects from changes in the other parameters. Lowering ν or σ_v^2 will slow the speed of convergence whereas lower ρ_q and lower $|\rho_{pq}|$ will speed convergence, but these effects should be very small.

1.2.8 Parameter choices

There will be a degree of subjectivity involved in the setting of the hyperparameters of the prior, particularly as the referendum was a one-off event with little historical precedent. We use both general election and polling data to inform our choice of parameter values.

1.2.8.1 Turnout

National turnout There were reports of high turnout on the day of the vote itself (Gutteridge, 2016). We will use national turnout for general elections as a guide, but note that the Scottish independence referendum had an unprecedentedly high turnout of 85%. The general election turnout figures since 1945 are shown in Table 1.1. The average is 66.9% (6.7%) and for the last three elections the average is 64.2%. We use 67.6% which is the three-election average weighted upwards by half the six-election standard deviation. A reasonable range of expectations would be 65–70%.

Area turnout Voting regions for the EU referendum were not the same as the constituencies used for general elections. However, the EC categorizes both the 381 voting areas in the referendum (excluding Gibraltar) and the (most recently 650) general election constituencies by 12 region codes. This enables us to make a more granular estimate of turnout per area v_i than simply assuming a uniform expectation. We use average turnout for each region for the 2010 and 2015 general elections as outlined in Table 1.2. Similar to the means of the expected vote-share per region, (ν_1, \dots, ν_n) can be uniformly shifted to achieve the required expected national turnout.

Table 1.1: Historical UK general election turnout.

Election Year	England	Wales	Scotland	N. Ireland	UK
2015	65.8%	65.7%	71.1%	58.1%	66.1%
2010	65.5%	64.7%	63.8%	57.6%	65.1%
2005	61.3%	62.6%	60.8%	62.9%	61.4%
2001	59.2%	61.6%	58.2%	68%	59.4%
1997	71.4%	73.5%	71.3%	67.1%	71.4%
1992	78%	79.7%	75.5%	69.8%	77.7%
Standard Deviation UK, 1992–2015					6.7%
1987	75.4%	78.9%	75.1%	67%	75.3%
1983	72.5%	76.1%	72.7%	72.9%	72.7%
1979	75.9%	79.4%	76.8%	67.7%	76%
1974 Feb	79%	80%	79%	69.9%	78.8%
1974 Oct	72.6%	76.6%	74.8%	67.7%	72.8%
1970	71.4%	77.4%	74.1%	76.6%	72%
1966	75.9%	79%	76%	66.1%	75.8%
1964	77%	80.1%	77.6%	71.7%	77.1%
1959	78.9%	82.6%	78.1%	65.9%	78.7%
1955	76.9%	79.6%	75.1%	74.1%	76.8%
1951	82.7%	84.4%	81.2%	79.9%	82.6%
1950	84.4%	84.8%	80.9%	77.4%	83.9%
1945	73.4%	75.7%	69%	67.4%	72.8%

Table 1.2: Turnout per EC Region in 2010 and 2015.

Region	2015 Turnout	2010 Turnout	Average Turnout
East	67.5%	67.6%	67.6%
East Midlands	66.5%	66.8%	66.6%
London	65.4%	64.5%	64.9%
North East	61.8%	61.1%	61.4%
North West	64.3%	62.3%	63.3%
Northern Ireland	58.1%	57.6%	57.8%
Scotland	71.0%	63.8%	67.4%
South East	68.6%	68.2%	68.4%
South West	69.5%	69.0%	69.2%
Wales	65.7%	64.8%	65.2%
West Midlands	64.1%	64.7%	64.4%
Yorkshire and The Humber	63.3%	62.9%	63.1%

Turnout variance Instead of setting turnout variance by region we will simply use the same level for every area and use the standard deviation figure for the last six general elections: 6.7%.

1.2.8.2 Turnout correlation by area

Given elections in time periods $t = 1, \dots, T$ and turnouts q_{it} , if we have predictions in advance for q_{it} , \bar{q}_{it} , then we can model the prediction errors $\Delta q_{it} = q_{it} - \bar{q}_{it}$ as being due to a national error η_t and individual error terms ϵ_{it} i.e,

$$\Delta q_{it} = \eta_t + \epsilon_{it}, \quad \eta_t \sim N(0, \sigma_\eta^2), \quad \epsilon_{it} \sim N(0, \sigma_\epsilon^2), \quad \text{cov}(\epsilon_t, \eta_{it}) = 0$$

Then for $i \neq j$, ρ_q is given by

$$\rho_q = \text{corr}(\Delta q_{it}, \Delta q_{jt}) = \frac{\sigma_\eta^2}{\sigma_\eta^2 + \sigma_\epsilon^2}, \quad i \neq j.$$

A regression of 2015 constituency turnout on 2010 turnout yields a coefficient of determination of 0.734 which provides evidence for simply using the turnout of the last election as the prediction. We do so. σ_η^2 is simply the variance of the national turnout ($6.7\%^2$). Calculation of σ_ϵ^2 requires looking at errors at the constituency level for each separate election. As there was the fifth constituency boundary review in 2008 we can form no easy prediction for area turnout for the 2010 election because constituencies changed. We simply use the 2015 election to estimate σ_ϵ^2 with predictions provided by the 2010 election. This results in estimates of $\sigma_\epsilon = 3.0\%$ and $\rho_q = 0.835$.⁴ In the absence of any other estimate or information pertinent to likely voting habits, this is what we use.

1.2.9 Vote-share

1.2.9.1 Area vote-share

An expected national vote-share μ_N is required. Polls with samples in the week preceding the referendum are shown in Table 1.3, along with seven polls of polls. For general elections, exit polls measure how people declare they have voted on the day itself at a selection of particular, secret, polling stations. They are much more accurate than any pre-election polling (Curtice et al., 2011), due to the fact that there is no measurement error of respondents. There was no exit poll for

⁴Estimates based on sample moments are consistent due to the Law of Large Numbers.

Table 1.3: Opinion polling prior to the EU referendum.

Date(s)	Remain	Leave	Undecided	Remain Lead	Organisation
22 Jun	55%	45%	--	10%	Populus
20–22 Jun	51%	49%	--	2%	YouGov
20–22 Jun	49%	46%	1%	3%	Ipsos MORI
20–22 Jun	44%	45%	9%	1%	Opinium
17–22 Jun	54%	46%	--	8%	ComRes
17–22 Jun	48%	42%	11%	6%	ComRes
16–22 Jun	41%	43%	16%	2%	TNS
20 Jun	45%	44%	11%	1%	Survation/IG Group
18–19 Jun	42%	44%	13%	2%	YouGov
16–19 Jun	53%	46%	2%	7%	ORB/Telegraph
17–18 Jun	45%	42%	13%	3%	Survation
Polls of Polls					
23 Jun	52%	48%	--	4%	What UK Thinks: EU
23 Jun	50.6%	49.4%	--	1.2%	Elections Etc.
23 Jun	45.8%	45.3%	9%	0.5%	HuffPost Pollster
22 Jun	46%	44%	10%	2%	Number Cruncher Politics
23 Jun	48%	46%	6%	2%	Financial Times
22 Jun	51%	49%	--	2%	The Telegraph
23 Jun	44%	44%	9%	0%	The Economist
				2.0%	Average Poll of Polls

Source: Wikipedia (2016).

the referendum as it was a one-off election. There was, however, a poll on the day conducted by YouGov which was published shortly after voting closed at 10 pm YouGov (2016). This poll measured how people voted versus how those same individuals reported their voting intention the preceding day. The result was a demographically weighted result of 48.38% which was broadly in line with recent polls. As we consider this to be the most accurate poll, we set $\mu_N = 48.38\%$.

1.2.9.2 Variance of area vote-share

Variances are chosen by shifting those implied by Hanretty by a constant amount σ_N^2 so that a generous estimate of the variance of the national vote-share σ_p^2 results. Table 1.4 shows that the average error in opinion polls from the prior week in the last six general elections was 2.66%. General election polling is a well-researched

Table 1.4: Opinion polls and vote-share for the Conservatives for recent general elections.

Election	Average Poll (Prior Week)	Result	Error
2015	34%	37.8%	-3.80%
2010	35%	36.9%	-1.58%
2005	31%	33.2%	-2.20%
2001	31%	31.7%	-0.92%
1997	30%	30.7%	-0.27%
1992	37%	41.9%	-4.46%
$\hat{\sigma}_\epsilon$			2.66%

field with plenty of historical precedent and would provide far too confident a figure. As our aim is to produce a prediction based on a conservative prior, we set $\sigma_p = 5\%$.

1.2.9.3 Correlation between voting areas

We analyse general election data in a similar manner to section 1.2.8.2 to estimate ρ_p . We examine the variation at the constituency versus the national level of the Conservative party vote-share to inform our choice of ρ_p . Predictions for the constituency level vote are based on applying the implied swing from opinion polls from the week prior to each election, to the level of the last election. See Table 1.4 for these polling results and for the results of the last six elections. Complications arise due to Westminster constituency boundary reviews in 1995, 2005 and 2008. These reviews change the number of constituencies and their composition of voters. They occur periodically in order to remove variations in the number of electors in each area, and have tended to favour the Conservatives (Rallings et al. (2008)). This is a well understood problem and the website Electoral Calculus (Baxter (2017)) publishes implied election results for elections preceding a review to enable ready comparison; we use these implied figures.

The implied standard error of σ_η using the data in Table 1.4 is 2.66%.⁵. Relying on the last six elections, the constituency level error calculation yields an estimate of 4.18% for σ_ϵ , implying a correlation $\rho_p = 0.288$. However, our constituency level errors are probably estimated at too high a level as better predictions for constituency level results exist although we do not have ready access to them. For this

⁵As the model implicitly assumes a mean of zero this is the square root of the average of the squares of the error, not the sample variance.




Figure 1.2: Change in UKIP + Conservative vote-share versus change in constituency turnout at the 2015 general election.

reason, the value of σ_ϵ^2 is likely estimated too high and ρ_p too low. Consequently, this value of ρ_p will be treated as a lower bound.

The correlation ρ_p is likely to be the parameter with the largest effect on how quickly the model prediction will converge to the true result. The implied correlation coefficient, as estimated, appears to be stable. Using only the last three elections results in an estimate of 0.324. The largest estimated value of σ_ϵ (5.69%) in any single election for the last six was in 1997, which was (unsurprisingly) also the largest error in the national vote-share. If we combine this with $\sigma_\eta = 2.66\%$, $\rho_p = 0.18$ results. This is an artificially low estimate and parameter values below this level are highly unlikely.

1.2.10 Correlation between vote-share and turnout

There were conflicting reports concerning the probable impact of turnout, even within the same newspaper on the day of the result (Gutteridge, 2016, Foster, 2016).

Table 1.5: Plausible parameter values.

Parameter	Description	Value	Range
ν	National Turnout	67.6	65–70%
σ_ν	Turnout Error	$(6.7\%)^2$	–
ρ_q	Area Turnout Correlation	0.835	–
μ_N	Expected National Vote-share	48.38	–
σ_p	National Vote Error	5%	–
ρ_p	Area Vote Correlation	≥ 0.288	≥ 0.18
ρ_{pq}	Area Vote and Turnout Correlation	–0.361	–

Predicting how unexpected turnout affects results requires a successful prediction of whether the difference in turnout is attributed to Leave or Remain supporters. This is a difficult problem. It was well understood in advance that younger voters, who are less likely to vote, would favour Remain, and that Brexit supporters were reported in surveys being more than twice as likely to vote as Remain ones (Twyman, 2016). YouGov were widely quoted as suggesting the relationship between turnout and Brexit vote-share would be negative (see Foster, 2016 for an example).

We take a quantitative approach based on general elections. Due to boundary changes, we are restricted to studying only the 2015 general election as no implied turnouts for elections preceding a boundary review are available. We proxy support for Brexit at the 2015 general election by using combined vote-shares of UKIP and the Conservatives (the parties with supporters most likely unsympathetic to the EU). We regress the swing of the combined UKIP and Conservatives vote-share against the change in turnout at the constituency level for the 2010-2015 elections. Figure 1.2 shows the regression, which results in a statistically significant correlation coefficient of -0.361 . This is indeed negative, agreeing with YouGov, and we use this value. However, we note that there may have been beliefs that this parameter may have been of different sign due to conflicting reports in the media.⁶

⁶The actual correlation observed on the night between vote-share surprises and turnout surprises at the area level was around -0.1 . Although turnout was higher than expected, as was the vote for leaving the EU, at the area level, turnout was even higher on average, for areas with lower support for Brexit

1.2.11 A Note on model correlation and measured correlation

The model correlation parameters ρ_p , ρ_q and ρ_{pq} are not strictly the correlations of the random variables $r = (r_1, \dots, r_{2n})$. They are the correlations of Φ^{-1} of the quantiles being $(\Phi^{-1}(F_1(r_1)), \dots, \Phi^{-1}(F_{2n}(r_{2n})))$. When the marginal is normal, $\Phi^{-1}(F_i(r_i))$ will be a linear function of r_i and the correlations will be identical. When $\Phi^{-1}(F_i(\bullet))$ is non-linear, a simulation or other method would be strictly required to convert them. However, we omit this step as it is not significant in practice.

1.2.12 Missing Priors

Of the 382 voting areas of the referendum, Hanretty failed to publish priors for four areas. These are listed in Table 1.6. The four areas are:

1. Gibraltar: This makes up a tiny 0.05% of the electorate, was the first area to announce, and had overwhelming support for Remain (Reyes (2016b,a)). As the population is so distinct from that of the rest of the UK, the result is not informative. We therefore take it as given and do not include it in the model.
2. The Isles of Scilly and Isle of Anglesey make up only 0.11% of the electorate and are simply ignored.
3. Northern Ireland consists of about 1.26 million voters in a total electorate of roughly 46.5 million. We use opinion polls for the mean and a standard deviation equal to that of the average of the other areas. We use a poll published on June 20 (Shapiro (2016)) that showed Remain 11% ahead, or 9% higher than the rest of the UK at that time. We therefore set the mean equal to $\mu_N - 9\%$.

1.3 Theoretical framework

In this section we present a simple framework to analyse the efficiency of the betting and pound markets overnight on the Brexit vote. This will be used in the results section to evaluate whether the markets fully reflected the information contained in

Table 1.6: The four voting areas with missing priors*.

Area	Declaration (Actual / Expected)	Electorate (%)
Gibraltar	23:36:33 / 00:01	24,119 (0.05%)
Isles of Scilly	00:49:42 / 00:01	1,799 (0.004%)
Isle of Anglesey	02:18:00 / 02:30	51,445 (0.11%)
Northern Ireland	04:37:00 / 04:00	1,260,955 (2.71%)

*Source: Hanretty (2016)

the vote as they were announced, or if there was a delay. The general approach is to evaluate an implied electoral probability from the sterling and betting markets and compare them with that from the electoral model.

1.3.1 Betting market

The expected utility of the binary options traded on the Betfair exchange that pay out £1 in the event of a vote to leave the EU is

$$u(£1) \times \mathbb{P}_t^B(\text{Brexit}),$$

where $\mathbb{P}_t^B(E = 1)$ is the aggregate belief of participants in the betting market of the probability of Brexit at time t and $u(\cdot)$ is the Bernoulli utility function. If the price of the contract is PB_t then equality of expected utility implies

$$u(PB_t) = u(£1) \times \mathbb{P}_t^B(\text{Brexit})$$

\Rightarrow

$$\mathbb{P}_t^B(\text{Brexit}) = \frac{u(PB_t)}{u(£1)}$$

This is a non-linear increasing function of the price of the contract PB_t . As is intuitively obvious, the odds are zero when the price is zero, and unity when the price is £1. Making the further stronger assumption of risk neutrality implies a linear utility function and

$$\mathbb{P}_t^B(\text{Brexit}) = PB_t.$$

This recovers the well known result that the price of a binary option is the risk neutral probability of the event under which the option pays out. Under semi-strong market efficiency this price should equal the true probability of Brexit occurring.

1.3.2 Currency market

We need to work a little harder to extract an implied probability of leaving the EU from the currency markets. To do so we first write down the uncovered interest parity relationship for cable. This relationship holds under perfect mobility of capital and risk aversion.

$$\mathbb{E}_t(GBP_T) - GBP_t = i - i^*$$

where $i - i^*$ is the interest rate differential from t to T between the British pound and United States dollar. However, for our application, t is in the overnight session and T is the time the result becomes known. The latter is in the early hours of the morning so no interest is expected to be earned. The RHS of the above relationship vanishes. This recovers the results that under risk neutrality and perfect mobility of capital the overnight price of the pound is equal to the expectation of the price at a time directly after the announcement of the referendum result. To extract an implied probability, we apply the total law of expectation.

$$\begin{aligned} GBP_t &= \mathbb{E}_t(GBP_T) \\ &= \mathbb{P}_t^{GBP}(\text{Remain}) \times \mathbb{E}_t(GBP_T | \text{Remain}) \\ &\quad + \mathbb{P}_t^{GBP}(\text{Brexit}) \times \mathbb{E}_t(GBP_T | \text{Brexit}), \end{aligned}$$

where $\mathbb{P}_t^{GBP}(\cdot)$ is the aggregate belief of investors in the currency market of the probability of an event at time t . Write $p_0 = \mathbb{E}_t(GBP_T | \text{Remain})$ and $p_1 = \mathbb{E}_t(GBP_T | \text{Brexit})$. Since ‘Brexit’ and ‘Remain’ are mutually exclusive events $\mathbb{P}_t^{GBP}(\text{Remain}) = 1 - \mathbb{P}_t^{GBP}(\text{Brexit})$ and

$$GBP_t = p_0 + \Delta p \times \mathbb{P}_t^{GBP}(\text{Brexit}), \tag{1.3.1}$$

where $\Delta p = p_1 - p_0$. Rearranging yields

$$\mathbb{P}_t^{GBP}(\text{Brexit}) = \frac{GBP_t - p_0}{\Delta p}.$$

This is of course a linear function of the pound price. However, to evaluate the probability $\mathbb{P}_t^{GBP}(\text{Brexit})$, p_0 and p_1 are required. Ex-ante this will not necessarily be known (although estimates from commentators were available). Our approach will be to evaluate an estimate ex-post from the data. Under weak market efficiency the probabilities of Brexit implied by both the betting markets must be equal. Thus $\mathbb{P}_t^{GBP}(\text{Brexit}) = \mathbb{P}_t^B(\text{Brexit})$. Further, under risk neutrality, we can replace this probability in equation 1.3.1 with PB_t . Thus

$$GBP_t = p_0 + \Delta p \times PB_t,$$

that is, the pound and betting markets are linearly related. This relationship is explored more rigorously in the next chapter where weak market efficiency is a key assumption. We are currently considering semi-strong market efficiency. For now we note that if that holds then weak market efficiency must also hold. For the purposes of this chapter we will borrow the linear relationship derived from weak market efficiency to form estimates of the conditional expectations of the pound given either ‘Remain’ or ‘Brexit’. These expectations are formed from linear regression. As the two variables in the relationship are prices, they are non-stationary. Regression will only be valid under cointegration. However, if this does hold then estimates are super-consistent. Again, this is studied in much more depth in the next chapter but suffice to say we do find cointegration holds.⁷ Our approach appears valid. We thus evaluate the probability of Brexit from the financial markets as

$$\mathbb{P}_t^{GBP}(\text{Brexit}) = \frac{GBP_t - \hat{p}_0}{\hat{\Delta p}}, \quad (1.3.2)$$

⁷We also find that on the night of the EU referendum there do not appear to be any significant deviations from risk neutrality.

where $\hat{\Delta}p$ and \hat{p}_0 are the estimated slope and intercept of a regression of GBP_t on PB_t for the period under study.

1.4 Results

To produce results we used the following data sources.

GBPUSD future price

The GBPUSD future price traded on the Chicago Mercantile Exchange was used rather than the spot price. There are multiple exchanges where the spot trades and aggregation could be prohibitively difficult. It is well known that spot and futures prices for foreign exchange are extremely well correlated and are effectively contemporaneous on time scales of under one second when both markets are open. The data was downloaded from Bloomberg and timestamps of trades were reported to an accuracy of one second. Note that the futures contract was closed between 10pm and 11pm on 23 June which was before the announcement of results.

Betfair data

The betting website Betfair listed two contracts. These were traded on Betfair's exchange platform which acts as a limit order book. The first paid out £1 in the event of Brexit, the other paid £1 in the event of Remain. The sum of the prices of the contracts did not deviate sufficiently from £1 to enable a profitable arbitrage. Betfair supplied all trades with timestamps of one second granularity in both contracts between 10 pm on 23 June to 5 am on 24 June. We convert all prices in the Remain contract to a synthetic price in the Brexit one by subtracting from £1. There were 182,534 trades in the window we study. £51,016,907 in total was matched during this 7-hour window. This compares with 88,246 trades in the GBP future during this time with a total notional traded of around \$5.5Bn. Although the futures market is considerably larger in notional traded terms than the betting market, the Betfair contracts moved by around 90% of their price whereas the pound moved around 10%.

Timing results

The earliest confirmed time for each voting area from three sources were used. The first source is the Press Association, the second the time that the returning officer in each area reported to the Electoral Commission and the third was the time Bloomberg published the result on it's real-time news feed.⁸

Electoral Probability Model results were generated in Matlab, sampling from the relevant multivariate distributions to evaluate model probabilities. Where calibration was required, we found that a simple gradient descent algorithm was adequate.

1.4.1 Market efficiency


Figure 1.3 shows the results for calibrating the model prior to the generous standard deviation figure of 5% using the parameters from Table 1.5. The kurtosis of the distributions is lower for the normal model, due to the logit mapping squeezing the distribution. The logit t-distribution has higher kurtosis, as expected. Both logit distributions are highly significantly different to the normal distribution, as shown by the Jarque–Bera (JB) test pValue. All marginals had a very generous prior 90% confidence interval of around 16.5% width. Figure 1.4 shows the evolution of the model forecast as the night progressed. The results of all marginals are almost identical except for the higher kurtosis logit-t marginal, which had a surprisingly quicker rate of convergence. The logit-t model predicted Brexit at 1:23:34am on the 12th result with 95% accuracy, and on the 18th result at 1:45am with 99% accuracy. The other models took until 1:43:46am (15th result) and 2:03am (33th result) to get to that level of certainty respectively. This compares with the BBC projecting Brexit at 4:39:32am.⁹ The Betfair market took until 04:21:00 am to imply 99% probability.

We now apply the approach presented in section 1.3 to extract implied probabilities for Brexit from the two markets. For the betting contracts this is trivial. The probability is simply the price of the 'Brexit' bet. For the pound market, a regression is required. The markets do appear to be moving very similarly. Figure 1.5 is a scatter plot of the pound future price against that of the betting contract for data taken every one minute in the overnight session. The regression gives an intercept of $\hat{p}_0 = 1.347$ and a slope of $\hat{\Delta}p = -0.172$. We use equation 1.3.2 to

⁸Only 67 of the 382 results were published in real-time by Bloomberg

⁹According to our Bloomberg scrape.

Forecast Model Prior



Marginal	Kurtosis	90% CI	99% CI	JB p
Normal	3.0	(40.1%, 56.7%)	(35.4%, 61.4%)	> 0.5
Logit Normal	2.9	(40.2%, 56.6%)	(35.9%, 60.1%)	< 0.001
Beta	3.0	(40%, 56.6%)	(35.5%, 61.3%)	0.29
Logit t	3.6	(40%, 56.6%)	(34.4%, 62.4%)	> 0.5

95% result	99% result
15 (1:43:46)	33 (2:03:00)
15 (1:43:46)	33 (2:03:00)
15 (1:43:46)	33 (2:03:00)
12 (1:23:34)	18 (1:45:00)

Figure 1.3: Prior distributions for forecast model.

Forecast Model Path






Figure 1.4: National vote distribution evolution (Logit Normal) and model probability paths for forecast model.



Note: The blue line is the best linear predictor and is $GBP_t = 1.347 - 0.172 \times PB_t$

Figure 1.5: Scatter plot GBP against the price of the Betfair contract for Brexit with the line of best fit.

evaluate an estimated probability for Brexit from the pound market. Figure 1.6 graphs the implied probabilities of the electoral forecast model with those from the markets from 11pm. Clearly market probabilities are nowhere close to those of the forecast model, as would be expected under semi-strong market efficiency. There is a large lag between the forecast and the market. The forecast model leads both the Betfair price and the implied pound probability. The average horizontal distance on this plot between the relevant lines is 172 minutes for forecast-Betfair and 178 for forecast-pound (6 minutes slower on average). It appears that the fundamental information (forecast) led the markets by around 3 hours. Semi-strong market efficiency did not hold. Further, it appears that the betting markets led the currency market by a very small number of minutes.

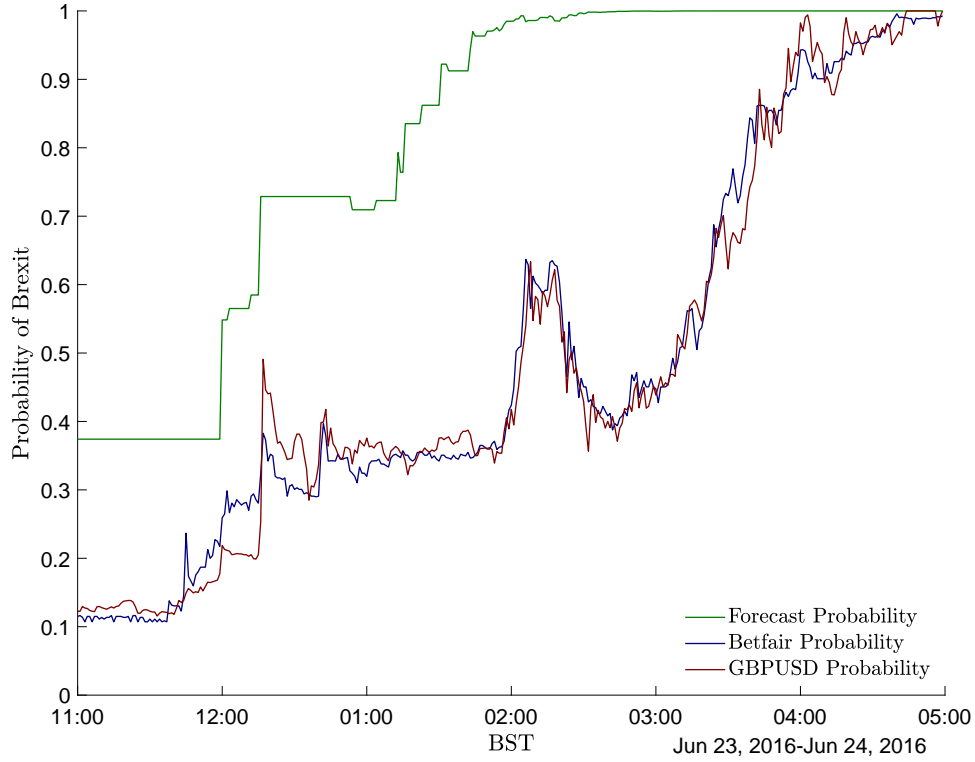


Figure 1.6: The forecast model probability and market implied probabilities.

1.4.2 Robustness checks

We now consider the sensitivity of the electoral model to changes in the model parameter changes. Table 1.7 shows the 95% Brexit prediction times for various parameter changes. No changes are observed when making all parameters except σ_p and ρ_p more conservative by roughly 5% (row 1). Increasing only σ_p by 5% or reducing ρ_p by 5% (rows 2 and 3) reduces the speed of convergence a little with the 95% prediction coming on the 16th as opposed to 15th result. Lowering ρ_p further to the very conservative lower bound of 0.18 slows the result significantly but still predicts Brexit with 99% probability about an hour and a half before the Betfair market. Making the correlation between vote-share and turnout at the area level ρ_{pq} both zero or of opposite sign increases the speed of convergence a little, with the 95% probability now coming on the 14th result. This suggests, although there is some subjectivity in the hyperparameter choices, the conclusions of violations of markets efficiency are both robust to sensible changes as well as

Table 1.7: How the times to predict brexit vary with more conservative parameter values (Normal Marginal).

Result	σ_P	ρ_p	ρ_q	ρ_{pq}	ν	σ_ν	95% result ^a
1 ^b	5%	0.288	0.877	-0.379	64.2%	6.4%	15 (1:43:46)
2 ^c	5%	0.274	0.835	-0.361	67.6%	6.7%	15 (1:43:46)
3 ^d	5.3%	0.288	0.835	-0.361	67.6%	6.7%	16 (1:44:57)*
4 ^e	5%	0.18	0.835	-0.361	67.6%	6.7%	34 (2:05:00)*
5 ^f	5.3%	0.288	0.835	0	67.6%	6.7%	14 (1:34:21)*
6 ^g	5.3%	0.288	0.835	0.361	67.6%	6.7%	14 (1:34:21)*

^a*Indicates a change from initial values

^bResult 1 changes all parameters except ρ_p and σ_p .

^cResult 2 changes ρ_p only.

^dResult 3 changes σ_p only.

^eResult 4 sets ρ_p at the lower limit of our plausible range, 0.18.

^fResult 5 sets $\rho_{pq} = 0$.

^gResult 6 changes the sign of ρ_{pq} .

completely changing the sign of ρ_{pq} . Further, we observe that changes to σ_p and ρ_p only meaningfully affect the model to a large degree, which is as expected.

Heterogeneous correlations

As noted, the electoral model is limited by having homogeneous correlation coefficients. To test this limitation we ran the model with a modified vote-share correlation matrix Σ_p . Correlations were set so that they were higher for areas with similar expected vote-share, according to the following formula:

$$(\Sigma_p)_{ij} = A - B \sqrt{|\mu_i - \mu_j|},$$

where A and B are constants chosen so that the average correlation is the usual value of $\rho_p = 0.288$ and the minimum value was 0.1. This form, although crude, is justified as areas with closer expectations are more likely demographically similar and so have a greater degree of dependence. Running this model for normal marginals resulted in 95% and 99% probabilities of Brexit at 1:15 am (10th result) and 1:17 am (11th result), respectively. This was considerably faster than the homogeneous case, suggesting that departures from homogeneity in the model would quicken convergence and lead to a strengthening of our conclusions.

1.4.3 Irrationality

The results described so far suggest that investors were exceptionally slow in discounting the information contained in the vote results. However, we have made various choices in the choice of prior and the associated hyper-parameters, any of which can be argued with. We now discuss the exact restrictions we have made in the choice of prior. This is to establish that there really was irrational behaviour on the night of the vote and that our results are not due to the particular form of the prior chosen.

The first constraint we made was that the joint prior of the area voting percentages and turnout is continuous. This is not a controversial assumption. Once that is accepted Sklar's theorem implies there is a copula form for the prior. The choices we make then are restricted to the copula, and the marginals. We have demonstrated that changing the *shape* of the marginal distributions does not meaningfully change the speed of convergence. As we discussed in section 1.2.7, increasing the variance of the prior, σ_p^2 , decreases the speed of convergence. However, it is difficult to argue that, with a 16.5% wide 90% confidence interval in the national vote-share for Brexit, σ_p is too small. Further, this variance is also much larger than that suggested by the likelihood of Brexit implied by the markets. With regard to the copula, we have chosen a simple two-factor correlation matrix with hyper-parameters chosen based on recent UK general elections. The particular form of the correlation matrix *can* be argued with. However, once that choice has been made, using a normal copula cannot. This is because a normal distribution maximises the entropy given the covariance. That is, it is an uninformative prior given the second moment constraint. So it appears that the main restriction we have made is with the parametrisation of the correlation matrix. One could arguably improve the form of the matrix, by perhaps including factors relating to different geographical areas or demographic factors. However, as discussed in section 1.2.7, the overwhelming determinant of the speed of convergence of the forecast is set by the average correlation between different areas vote-shares, ρ_p , (and we have also demonstrated robustness to some limited heterogeneity in correlations). Thus, the main restriction we have made is with regards to the magnitude of the parameter ρ_p . It is possible that choosing a smaller value for this parameter would lead to different conclusions.

To demonstrate irrationality, we now ask if it is possible that investors simply believed that the different vote area results had a low association and that they

were in fact consistently updating misspecified ex-ante beliefs? To answer this question we attempt a crude calibration of the two important model parameters to market behaviour. σ_p measures ex-ante uncertainty in our model through the standard error of the prior. We set this so as to “agree” with the probabilities expressed in the market prior to the realisation of results. ρ_p models the degree to which different area vote-shares are associated, through the correlation in the Gaussian copula prior. We explore how different values of this parameter affect the path of our model probability. Calibration of initial Brexit probability to $P_0 = 25\%$ yields an initial standard deviation of 2.4% (2.7% for logit). This is lower than recent general election polling errors. Further this is lower than the reported standard error of the YouGov poll on the day released at 10pm. Thus investors appeared to have less uncertainty about the outcome of the vote than that assumed if YouGov had modeled population turnout characteristics perfectly. This in itself appears to indicate irrationally overconfident investors. To explore the degree of association between area votes assumed by consistent investor behaviour, we lower ρ_p while setting the variance as above and fixing the other parameter values. Figure 1.7 shows the results and suggests that a model roughly consistent with the markets would necessitate $\rho_p \in (0, 0.01)$. This is simply implausible; it would assume that different area results were close to being perfectly independent whereas constituencies in general elections have a fairly stable correlation of around 0.3. It appears that either ex-ante beliefs were irrational and we updated rationally, or the ex-ante belief was rational but updating was irrational. Given Bayes law is a mathematical fact of life, the conclusion is that investors in both the betting and currency markets displayed irrational behaviour on the night of the EU referendum.

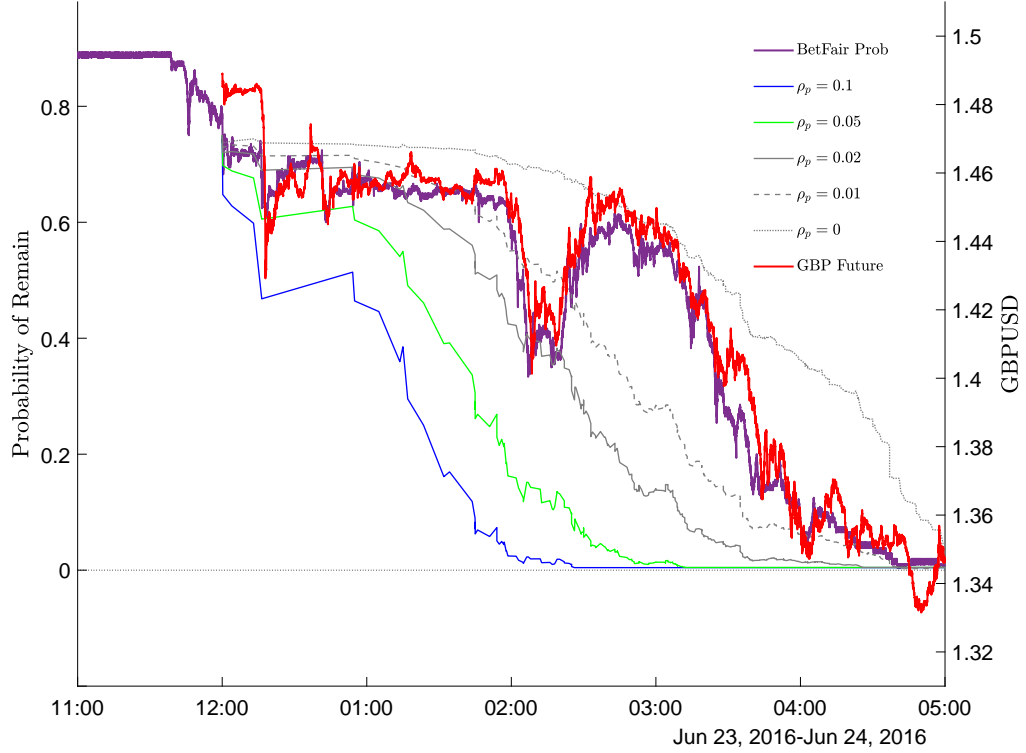


Figure 1.7: Model probability paths when lowering ρ_p .

1.5 Conclusion

This chapter examined the efficiency of the Betfair and GBPUSD market as the results of the United Kingdom European Union membership referendum were announced. This event provided a unique opportunity to study the interaction between a flow of information, a prediction market and a financial market, where there was a sole, public determinant of prices. Other work has identified the pound market as being inefficient during the period under investigation but we were able to answer questions about the efficiency of the prediction market.

We have presented a rigorous Bayesian realtime model of the probability of Brexit for the period under consideration. This is based on a copula that is not constrained to normal distributions. The Bayesian method improves upon earlier estimation methods as it does not rely on any asymptotic properties of estimators for small samples. We also demonstrate robustness of results to changes in the prior. The conclusions of the model are as follows:

1. Not only was the currency market informationally inefficient so too was the betting market. Both markets violated semi-strong EMH on the night of the vote, taking around 3 hours to reflect publicly available information.
2. The delay in the betting markets reacting to the vote was slightly smaller than that of the currency market, of the order of a very small number of minutes.
3. The observed mispricing is inconsistent with any plausible prior and flow of information. Investors were irrationally overconfident about the degree of uncertainty in the result. Further, an assumption of consistent behaviour implies a simply unbelievable degree of association between the votes of different areas of the UK (very close to perfect independence of the results of different constituencies).

Our results suggest that market participants suffered a behavioural bias as the results unfolded. It appears that traders and gamblers simply could not believe that the UK was voting to leave the EU. These results are consistent with, and complement, previous work that identified a ‘bubble’ in beliefs in the weeks and months leading up to the vote. Any future possible UK referendum on this subject will present an opportunity to study whether this inefficiency persists or whether efficient behaviour is exhibited, possibly due to the publication of this and other studies.

Chapter 2

Election-night models

Note to examiners

An early version of a theoretical model linking the price of the British pound to betting markets for the Brexit referendum was included in a submission for a dissertation for the MPhil in Economics Research at the University of Cambridge in 2017. There are some similarities in that theoretical model and the one presented in this chapter. This is particularly the case for the discussion of the assumptions of the model, Section 2.2.1. However, the model presented in this chapter applies to all financial assets, not just the British Pound.

2.1 Introduction

How do financial and political betting markets behave in the hours after an election? Are there any special behaviours that can be observed at only these times? In this chapter we attempt to answer these and other questions.

It can be argued that the overnight hours following many elections are very special times. It is possible that financial assets are uniquely determined by the outcome of those elections and nothing else. During these typically overnight hours there is an absence of economic information that may usually inform financial prices. For many elections the results are drip fed throughout the overnight session as results are announced from different voting areas and constituencies. For example, for US presidential elections there are 51 simultaneous vote counts in 50 states

plus the special district of Washington DC.¹ These counts are typically broadcast live on TV and on the Internet. For the 2016 UK Brexit referendum there were 382 different areas that announced local results at various times starting from around midnight.

We begin this chapter by creating a theoretical framework linking financial and political markets that applies only during these special times. The key assumption is that the only information that has a persistent effect on asset prices is that related to the results of a political event. We then go on to test the implications of this theory to some real world elections in the results section. We also test whether fleeting deviations from this theory can be used to make profitable trading strategies, as well as whether or not price action early on in the night can be used to predict behaviour in the latter hours. Finally we consider whether the political or financial markets were more efficient at discounting the political information. We end the chapter by concluding our results.

2.2 Pricing model

In this section we present a pricing model that links the prices of political and financial markets during the overnight results session following an election.

Notation

E	Binary political event indicator $\in \{0, 1\}$
T	Time at which event is realised
p_t	Price of financial asset at time $t < T$
PB_t	Price of betting contract paying out £1 when $E = 1$ at time $t < T$
$\mathbb{P}_t^B(\cdot)$	Belief of a representative investor in the betting market at time $t < T$
$\mathbb{P}_t^f(\cdot)$	Belief of a representative investor in the financial market at time $t < T$
p_1	Expected value of financial asset conditional on $E = 1$
p_0	Expected value of financial asset conditional on $E = 0$

¹The 2020 presidential election is a unique case where, given various legal challenges and recounts, results took days to conclude.

Scenario

First we outline an overnight scenario. We consider a time period overnight after the vote has closed, but up to and including the result becoming known. Suppose that the outcome of the election can be represented in a *binary* fashion (for example a yes/no referendum). Further suppose there is a healthy liquid prediction market for this event where binary options trade that payout on the outcome of the given event. Such contracts exist and examples include bets that paid £1 in the event of the UK voting to leave the EU in 2016 as well as Donald Trump winning the Presidential election. These bets are effectively contingent claims. Say $E = 1$ if the event occurs and $E = 0$ otherwise. Write PB_t as the price of the contract that pays out £1 when $E = 1$. The outcome of the event is realised at time T . Further, suppose there is a financial asset with price p_t at time $t < T$.

No Shocks condition

We build our model beginning with a key assumption. This is that there are no persistent shocks to the price of financial assets beyond those related to the probability of the outcome of the political event, during the overnight interval. We discuss this assumption in the case of the 2016 Brexit referendum and the pound dollar exchange rate GBPUSD. Simply put, this condition states that the only fundamental determinants of the GBP price on the night of the vote were the results of the vote, and that those results only affect the price through their effect on the probability of Brexit, $\mathbb{P}(E = 1)$. Any changes in price that are not related to $\mathbb{P}(E = 1)$ are stationary and will disappear quickly. Over a longer period of time there are certainly other determinants of the GBP price. For example unanticipated economic information relating to the health or otherwise of the US or UK economies would normally affect the exchange rate. Other examples of such information may even be related to the UK leaving the EU. If after the vote, new information arises regarding the nature of the trading relationship between the EU and the UK, then the exchange rate could be affected. However, on the night of the vote our assumption is that none of this matters. The only thing affecting financial prices is whether or not the UK leaves.

The statement that the only news affecting prices that night was related to the referendum is not controversial. Indeed, there were no major economic releases or other significant news events. The advance Econoday economic calendar, Econoday (2016), listed the final market related news releases on the 23 June 2016 as the US

New Home Sales Report at 10:00 am Eastern Time (ET) and the first one for 24 June (beyond the referendum) as Durable Goods Orders at 10 am ET. The calendar wrote regarding the 24th: “In a rare and potentially powerful wildcard, the markets will react to the Brexit outcome”. The authors of Wu et al. (2017) come to a similar conclusion and describe the circumstances as “a natural experiment” with “near perfect conditions” to study such a situation where there is a single determinant of prices.

The condition that only the *probability* of Brexit affects the GBPUSD price is more subtle. For instance, it could be believed that the vote share for leaving the EU has an affect on financial prices, over and above the decision to leave the EU. This perhaps could be due to a “harder” Brexit. However, the NS condition means that only the binary result of the vote, and not any particular form of trading relationship following a vote, was on the minds of investors on the night of the vote. Whether there is a 50.1% or 99.9% vote share for Brexit, the exchange rate will be the same. Despite there being intense scrutiny of the negotiations between the EU and the UK on the terms of withdrawal in the years following the vote, the term “hard Brexit” only first appeared several months after the referendum at the Conservative party conference in October 2016. The condition that the only thing affecting prices on the night of an election is the binary result of that election, is one we rely on in general to formulate our pricing model.

We operationalise the assumption as follows: For $t < T$ the non-stationary asset price p_t is a function only of $\mathbb{P}_t(E = 1)$. Write this function as $F(\cdot)$. Then

$$p_t = F(\mathbb{P}_t^f(E = 1)) + \varepsilon_t \quad (2.2.1)$$

where ε_t is mean-zero, serially uncorrelated and stationary and $\mathbb{P}_t^f(\cdot)$ represents the belief of a representative investor in the financial market. At the resolution of the election, $t = T$, E becomes either zero or one with certainty. Thus $\mathbb{E}_t(p_T) = \mathbb{E}_t(F(E))$ since $\mathbb{E}_t(\varepsilon_T) = 0$. The power of the NS assumption is that the expectations of p_T conditional on the outcome of the election E do not vary. This as

$$\begin{aligned} \mathbb{E}_t(p_T|E = 0) &= F(0) \\ \mathbb{E}_t(p_T|E = 1) &= F(1) \quad \forall t < T. \end{aligned} \quad (2.2.2)$$

Applying the total law of expectation to $\mathbb{E}_t(p_T)$ yields

$$\begin{aligned}
\mathbb{E}_t(p_T) &= \mathbb{P}_t^f(E = 0)F(0) + \mathbb{P}_t^f(E = 1)F(1) \\
&= F(0) + \mathbb{P}_t^f(E = 1) \times [F(1) - F(0)] \\
&= p_0 + \mathbb{P}_t^f(E = 1) \times \Delta p
\end{aligned} \tag{2.2.3}$$

where

$$\begin{aligned}
p_0 &= F(0) = \mathbb{E}(p_T|E = 0) \\
p_1 &= F(1) = \mathbb{E}(p_T|E = 1) \\
\Delta p &= p_1 - p_0.
\end{aligned}$$

Thus the NS assumption implies that the asset price at time T (‘in the morning’) is expected to be p_0 with probability $\mathbb{P}_t^f(E = 0)$ and p_1 with probability $\mathbb{P}_t^f(E = 1)$.

Weak Market Efficiency (EWMH)

The betting and financial markets are clear examples of segmented markets, having very different participants. However we add the conditions that investors in both sets of markets have identical information sets and very similar beliefs. The assumptions are consistent with those that lead to weak form market efficiency (EWMH). If we write the belief of the probability of $E = 1$ of a representative participant in the prediction market as $\mathbb{P}_t^B(E = 1)$ then a strict application of this condition means that $\mathbb{P}_t^B(E = 1) \equiv \mathbb{P}_t^f(E = 1)$. It is unrealistic though that assessments of the probability of the political event (or indeed market efficiency) will hold exactly at all times instantaneously. We thus allow errors in the relationship,

$$\mathbb{P}_t^f(E = 1) = \mathbb{P}_t^B(E = 1) + \eta_t \tag{2.2.4}$$

where η_t is the difference between the assessments of the probability of $E = 1$ by the representative investors in each market.

Results about the election arrive throughout the night of the election. Each result permanently affects the probabilities $\mathbb{P}_t^f(E = 1)$ and $\mathbb{P}_t^b(E = 1)$. These are non-stationary. They are also bounded. We model them as integrated $I(1)$ series with absorbing boundaries at zero and unity. Once the boundary is hit, the election has concluded. A strict application of equivalence of beliefs $\Rightarrow \eta_t \equiv 0$. This is unlikely. We distinguish between equivalence of beliefs holding in the long-term (long-term being a few hours here) and not holding at all. If deviations η_t quickly decay, efficiency broadly holds. This occurs when $\mathbb{E}(\eta_t) = 0$ and η_t is weakly dependent². If errors persist and η_t is non-stationary, then the condition does not hold at all. This introduction of the error, η_t , will allow for testing of our model later on.

Substituting equation 2.2.4 into equation 2.2.3 results in an expression for expected financial prices conditional on the error in beliefs η_t and dependent on beliefs of the representative investor in the betting market:

$$\mathbb{E}_t(p_T|\epsilon_t) = p_0 + \mathbb{P}_t^B(E = 1) \times \Delta p + \epsilon_t \quad \mathbb{E}(\epsilon_t) = 0 \quad (2.2.5)$$

where $\epsilon_t = \Delta p \times \eta_t$.

Risk neutral pricing in betting markets

Investors in financial assets are very likely to have correlated exposure to shock election results, leading to the presence of additional election related risk-premia. It is likely that a risk-premia existed for assets that have exposure to say, Scotland voting to leave the UK, the UK voting to leave the EU, or Donald Trump being elected in 2016. However, it is not clear that this is the case at all for bettors in election markets. There is indeed evidence that in the case of Brexit there were very many smaller bettors who favoured leaving the EU, trading against larger pro-EU participants, Economist (2016).

We assume risk neutral pricing in betting markets. This condition formalises the idea that in these markets there are investors present who will take on election risk for ever decreasing edge. This implies that the risk-neutral probability of $E = 1$ is indeed the assessed probability $\mathbb{P}_t^B(E = 1)$. As the time for payout of any bets made on the night of the election is very short, any discount factor is effectively

²Weak dependence ensures that any deviations away from zero are transient and do not persist.

unity. Thus the condition reduces to $\mathbb{P}_t^B(E = 1) = PB_t$. Then equation 2.2.5 can be updated with $\mathbb{P}_t^B(E = 1)$ replaced with PB_t . The expectation of financial asset prices can now be related to *prices* in the betting markets:

$$\mathbb{E}_t(p_T|\epsilon_t) = p_0 + \Delta p \times PB_t + \epsilon_t \quad \mathbb{E}(\epsilon_t) = 0. \quad (2.2.6)$$

The overnight session is ‘short’

Our model applies only to the hours overnight following an election vote. The no shocks condition necessarily implies a short time period, as the model only holds whilst the political event dominates the news cycle. Our model so far has resulted in an expression for the expectation of the morning price of a financial asset, given the price of a binary option in the betting market. We need a way to relate financial prices during the night at time t , p_t to this expectation. We assume a short period so that no interest can be earned and no equity dividends or bond coupons are paid. The expected appreciation of the financial asset overnight is $\mathbb{E}_t(p_T) - p_t$, as there are no other cashflows. Since no interest can be earned prices must adjust so that

$$\text{Risk premium} = \mathbb{E}_t(p_T) - p_t. \quad (2.2.7)$$

Although we assume risk neutral pricing in betting markets we do not do so in financial markets. The above risk premium is the excess expected return risk-averse investors require to hold the financial asset. During the night the ‘morning’ price p_T is expected to be p_0 with probability $1 - PB_t$ and p_1 with probability PB_t . Nothing beyond political event odds affects prices (an implication of the no shocks assumption). Another way we think of the overnight session as being ‘short’ is that any risk premium present is entirely due to the political event and nothing else. Risk factors beyond this particular election will not contribute to the LHS of equation 2.2.7, as the overnight period is too short for them to be earned. For example, investors cannot earn an additional return by holding, say, high beta stocks versus low beta stocks. Premia can only be earned due to the overnight volatility of the asset, and that volatility is solely caused by the uncertainty of the election result. This is not to say that risk premiums cannot change due to the result of the election (for example Mexican exporters would likely require higher expected returns if Trump were elected versus Clinton in 2016), just that this is

already factored into the RHS of equation 2.2.7 via $\mathbb{E}_t(p_T)$ through the expected morning price under Trump³.

Turning now to deriving an expression for the risk premium, we note that it must vanish at $E = 0$ or 1 , or equivalently $PB_t = 0$ or 1 , as at that point the price is certain. Any second order approximation to the premia must then be of the form $\lambda.PB_t.(1 - PB_t)$ where λ is a constant⁴. This closed form can also be recovered by assuming that the risk premium is proportional to the variance of p_T , which is also proportional to $PB_t.(1 - PB_t)$. (The variance is $(\Delta p)^2 \times PB_t.(1 - PB_t)$). Equation 2.2.7 now reduces to

$$PB_t.(1 - PB_t) \propto \mathbb{E}_t(p_T) - p_t.$$

Rearranging and combining the above with the expression in equation 2.2.6 yields our pricing model

$$p_t \approx p_0 + \Delta p.PB_t + \lambda.R_t + \epsilon_t \quad \mathbb{E}(\epsilon_t) = 0 \quad (2.2.8)$$

where

$$R_t = PB_t.(1 - PB_t).$$

We impose one final condition on the parameters of the model. In the absence of risk neutral pricing in financial markets, the relationship, as written in 2.2.8, could have a non-monotonic form. This occurs when the term related to risk is sufficiently large. For example when $\lambda < 0$ and $p_1 < p_0$ there is an expected depreciation under $E = 1$. If λ has sufficiently large magnitude then a minima can occur in $PB_t \in (0, 1)$. It does not seem plausible that p_t would be lower than the worst case scenario p_1 for intermediate odds, when there is still a chance of the night ending at p_0 . Monotonicity is ensured by the condition $|\lambda| \leq |\Delta p|$.⁵

Our pricing model is then summarised as

³ $\mathbb{E}_t(p_T) = \mathbb{P}_t(CLINTON)\mathbb{E}(p_T|CLINTON) + \mathbb{P}_t(TRUMP)\mathbb{E}(p_T|TRUMP).$

⁴Any quadratic form which vanishes at the endpoints $PB_t = 0$ and $PB_t = 1$ can be written in this way.

⁵The gradient of the cointegration relationship is $\Delta p - \lambda + 2\lambda.PB_t$ so is $\Delta p \pm \lambda$ at the endpoints. The condition $|\lambda| \leq |\Delta p|$ ensures the gradient does not change sign.

$$p_t = p_0 + \Delta p.PB_t + \lambda.R_t + \epsilon_t \quad \mathbb{E}(\epsilon_t) = 0, \quad |\lambda| < |\Delta p|. \quad (2.2.9)$$

Note that:

1. ϵ_t represents errors in beliefs between the betting and currency market that are inconsistent with typical conditions for weak market efficiency. If efficiency broadly holds then it is stationary and 2.2.9 describes a non-linear cointegrating relationship between financial asset prices and betting market prices. This can be tested.
2. When there is risk neutral pricing in financial markets $\lambda = 0$ and equation 2.2.8 describes standard linear cointegration between the prices in the two markets.
3. The $\lambda.R_t$ term is an adjustment to the price due to political risk. It results in a political risk premium (either positive or negative depending on the sign of λ). It is greatest when uncertainty is highest ($PB_t = 0.5$) and vanishes when the election outcome becomes known ($PB_t = 0, 1$).
4. As noted in Manasse et al. (2020), changing odds of the political event affect financial prices in two ways. The first is via the direct effect on the probability of yielding p_0 versus p_1 . The second is via the effect on the risk premium.

Alternative derivation: Mean-variance preferences

An alternative approach to imposing a ‘short’ overnight session to derive an explicit expression for the risk premium is outlined in Manasse et al. (2020). This assumes an explicit set of preferences for investors in the financial markets. The conditions that there are no cashflows during the session as well as the quadratic form of the risk premia are effectively baked into the utility function.

Proceed by considering a representative investor who chooses between holding a proportion ω of her wealth in the financial asset and the rest in a risk free asset. Standard mean-variance preferences are assumed so that the investor maximises:

$$U(w) = \omega \cdot [\mathbb{E}_t(p_T) - p_t] - \frac{r}{2} \omega^2 \sigma^2,$$

where r is the coefficient of absolute risk-aversion and σ^2 is the portfolio variance. The first term is the expected appreciation of the asset from a time t overnight before the full results of the event are apparent, and time T , the time at which E is realised. The second term is a penalty (under risk aversion) for holding the risky financial asset and is proportionate to the risk aversion coefficient r and the portfolio variance $\omega^2\sigma^2$.

Firstly we note that the expected appreciation $\mathbb{E}_t(p_T) - p_t$ must be positive for a risk averse investor to hold any of the risky financial assets. When this is the case the first order condition is

$$\mathbb{E}_t(p_T) - p_t = \omega r \sigma^2. \quad (2.2.10)$$

The portfolio share reduces with increased risk aversion r and asset variance σ^2 . Taking the supply of the financial assets S as fixed. Clearing of the financial market implies that $\omega W = S$, where W is the total available wealth of investors (assumed to be greater than S and also fixed). Then $\omega = s$ where $s = S/W$. σ^2 can be evaluated and is the variance of $p_0 + X \times \Delta p$ where X is a Bernoulli random variable with probability $\mathbb{P}_t^f(E = 1)$. This is $(\Delta p)^2 \cdot \mathbb{P}_t^f(E = 1) \cdot (1 - \mathbb{P}_t^f(E = 1))$. Thus the time varying risk premium can be written as

$$\begin{aligned} \lambda_t &= -\lambda \pi_t (1 - \pi_t) \\ \lambda &= -rs (\Delta p)^2 \\ \pi_t &= \mathbb{P}_t^f(E = 1). \end{aligned}$$

Substituting the above into equation 2.2.7 along with the approximation $\pi_t = \mathbb{P}_t^f(E = 1) \approx PB_t$ yields an expression for the expected appreciation of the financial asset overnight,

$$\mathbb{E}_t(p_T) - p_t = -\lambda R_t.$$

This can be combined with the previous expression for $\mathbb{E}_t(p_T|\epsilon_t)$, equation 2.2.6, to recover the cointegration model.

Discussion

Our pricing model relies on a number of restriction in its derivation. The model is encapsulated in equation 2.2.9. This describes a linear combination of the financial asset price with a non-linear function of the betting market binary option price as stationary. Any tests of the model are necessarily joint tests of the model's restrictions. Nonetheless, we can make some statements about how different failures of the model relate to the various restrictions imposed.

The fact that ϵ_t is stationary and mean zero comes from two restrictions. The first is the NS assumption. This states that only changes to the odds of the event, $\mathbb{P}(E = 1)$, persistently affects financial prices. This may not hold due to two reasons. The first is that non-election related news is changing asset prices. For instance if aliens invaded the planet during the night of the Brexit referendum we would expect stocks to sell off regardless of $\mathbb{P}(E = 1)$! The second is that there is information contained in the election results beyond $\mathbb{P}(E = 1)$ (and not perfectly correlated with $\mathbb{P}(E = 1)$) that informs the financial price. This could be for example due to a higher vote share for one outcome, or the geographical distribution of votes. Either way then ε_t in equation 2.2.1 is not mean zero. $\mathbb{E}_t(\varepsilon_T|E)$ will not vanish in the conditional expectations in equation 2.2.2 and then contributes to the RHS of equation 2.2.9. The second restriction which imposes the stationarity condition is similarity of beliefs of investors in the two markets, equation 2.2.4. Testing the stationarity of ϵ_t is thus a joint test of the NS assumption and weak market efficiency (although equivalent beliefs of investors in both markets is not the same as market efficiency, it is one of the conditions that lead to it). However, in our empirical work we will interpret deviations from $\epsilon_t = 0$ as deviations from equivalence of beliefs and thus weak market efficiency (EWMH). This is as ex-post, we know that aliens did not invade the planet (or war broke out or something else) and we can argue that the election really did dominate prices on the nights in question.

Equation 2.2.9 also specifies the shape of the contemporaneous function of prediction market binary options and financial asset prices that are stationarity. This shape derives from the restrictions we place on preferences of market participants. These are, one, risk neutral pricing in prediction markets, and two, mean-variance preferences of financial investors (when the latter is not assumed we show that the shape of the relationship must be approximately that of equation 2.2.9 through a second order approximation). Relaxations of these conditions lead to changes in

the shape of the contemporaneous relationship, but not the time series properties of ϵ_t . We do not provide a formal proof, but if, for example risk averse pricing is present in the betting markets then the market price PB_t will be a contemporaneous monotonic function of the belief of a representative investor $g(\mathbb{P}_t^b(E = 1))$, rather than equal to it. The result will be that the betting price PB_t will occur in a non-linear function in the RHS of equation 2.2.6 and thus in the pricing relationship. The conclusion that we make (again without formal proof) is that tests on the *shape* of the relationship in equation 2.2.9 are really tests on the restrictions we make about the preferences of investors.

Finally we comment on the condition that the overnight session is short. We claim that this is a weak condition. Firstly, it implies there are no cash flows. This is easy enough to verify for any particular asset and setting but if there were a cash flow then a performance price series that adjusts for it could be trivially constructed. Secondly, this condition means that no interest or risk premium (beyond that related to election uncertainty) can be earned. If interest, or an additional risk premium is present then the magnitude earned overnight would be proportional to the time left to the morning, $T - t$. There would be an additional term, say $\delta(T - t)$ cropping up in the difference of p_t and $\mathbb{E}_t(p_T)$ and thus on the RHS of the pricing equation 2.2.9. However, the magnitude of this term would be tiny. For example, if a relatively large annualised premium of 3.65% were being earned throughout the night this term would vary by a fraction of a basis point (there are 365 days in a year and the overnight session is a fraction of a day!). We thus ignore this possibility and assume this condition holds exactly.

Extension to multiple assets

We now consider the case when there are n assets indexed by i . Each asset is separately cointegrated with the prediction market so that

$$p_{it} = p_{i0} + \Delta p_i \cdot PB_t + \epsilon_{it}.$$

We can also consider the vector of $n + 1$ prices where the final price is that of the political contract PB_t . Then our model becomes

$$\begin{pmatrix} p_{1t} \\ \vdots \\ p_{nt} \\ PB_t \end{pmatrix} = \begin{pmatrix} p_{10} \\ \vdots \\ p_{n0} \\ 0 \end{pmatrix} + PB_t \begin{pmatrix} \Delta p_1 \\ \vdots \\ \Delta p_n \\ 1 \end{pmatrix} + \begin{pmatrix} \epsilon_{1t} \\ \vdots \\ \epsilon_{nt} \\ 0 \end{pmatrix}.$$

Note our model says nothing about whether or not idiosyncratic deviations from the cointegrating relationships are correlated or not. In fact we may expect dependence among these deviations among economically related assets. However, EWMH does imply that these deviations will not be predictable and thus must be martingale differences.

Using matrix notation we write \mathbf{P}_t as the $n + 1$ dimensional vector of prices which include the betting contract price, \mathbf{P}_0 as the n dimensional vector of expected prices of the financial asset conditional on $E = 1$ and $\boldsymbol{\epsilon}$ the n -dimensional vectors of idiosyncratic martingale differences

$$\mathbf{P}_t = \begin{pmatrix} p_{1t} \\ \vdots \\ p_{nt} \\ PB_t \end{pmatrix}, \quad \mathbf{P}_0 = \begin{pmatrix} p_{10} \\ \vdots \\ p_{n0} \end{pmatrix}, \quad \boldsymbol{\epsilon}_t = \begin{pmatrix} \epsilon_{1t} \\ \vdots \\ \epsilon_{nt} \end{pmatrix}.$$

The system of equations can be written as

$$\mathbf{\Pi} \mathbf{P}_t - \mathbf{P}_0 = \boldsymbol{\epsilon}_t \tag{2.2.11}$$

where

$$\mathbf{\Pi} = \begin{pmatrix} 1 & 0 & \cdots & 0 & 0 & -\Delta p_1 \\ 0 & \ddots & & & 0 & \\ \vdots & & 1 & & \vdots & \vdots \\ 0 & & & \ddots & 0 & \\ 0 & 0 & \cdots & 0 & 1 & -\Delta p_n \end{pmatrix}$$

an $n \times (n + 1)$ dimensional matrix of n cointegrating relationships. The n rows of equation 2.2.11 are stationary. Thus there is again a single common trend among the asset prices. This is equal to the likelihood of the outcome $E = 1$ and is equal

to the price of the betting contract PB_t . By assuming that the only persistent determinant of asset prices is the likelihood of the outcome of the political event we see that the cointegrating rank of the $n + 1$ asset prices which include the political market must be n during the results session. This is a proposition which can be tested empirically.

2.3 Statistical specifications

We aim to test the theoretical pricing model on a selection of political events. In each case we will study the overnight price series of a betting contract with a collection of key actively traded financial assets. The key implication of the theoretical model is one of cointegration. Our assumptions imply that there exist linear combinations of the non-stationary asset prices which are stationary. The statistical phenomena of cointegration was introduced by Granger (1983) who used it to study long-run economic relationships. For our model the “long-run” relationship lasts only a few hours. It exists between political and financial asset prices and arises due to the very special circumstances of the hours directly after an election. There are three statistical frameworks in which cointegration is studied, each producing various methods by which the theory can be tested. We describe these briefly below.⁶

2.3.1 Cointegration frameworks

2.3.1.1 The regression framework

Given a multivariate process $p_t = (p'_{1t}, p'_{2t})'$ where the dimension of p_{1t} and p_{2t} are n_1 and n_2 respectively the model is

$$\begin{aligned} p_{1t} &= \beta' p_{2t} + \epsilon_{1t} \\ \Delta p_{2t} &= \epsilon_{2t} \end{aligned}$$

where $\epsilon_t = (\epsilon'_{1t}, \epsilon'_{2t})'$ is a mean zero, finite variance stationary linear invertible process. p_t is non-stationary, p_{2t} is not cointegrating and p_t is cointegrating of rank n_1 .

⁶For simplicity we omit lags of differenced variables in the following exposition.

The cointegrating vectors are the rows of $\Pi = [I_{n_1}, -\beta']$. Our pricing model naturally fits into this framework with $n_1 = n$ the number of financial assets ($n = 1$ in the single asset case), $n_2 = 1$ and p_{2t} being the price of the political betting contract PB_t . $\beta = \Delta p$ the vectors of sensitivities of asset prices to the binary political event. Engle and Granger (1987) showed that, in the presence of cointegration, OLS estimators of β are superconsistent. Stock (1987) demonstrated that the asymptotic distributions of $\hat{\beta}$ are mixed normal distributions. Tests for cointegration are based on the residuals of the regression where the null hypothesis is the presence of a unit root in ϵ_t . This null hypothesis is equivalent to *no* cointegration in p_t . Tests of an econometric theory typically proceed with the assumption that the null hypothesis supports the model under consideration and a failure to reject the null is support for the theory. This is not the case here. However, rejection of the null at some significance level indicates a rejection of no cointegration in favour of cointegration. This would indicate stronger evidence of the pricing model than a failure to reject stationarity of ϵ_t at the same significance level.

Note that deviations from the long term relationships ϵ_t are not restricted to be i.i.d. Hansen and Phillips (1990), Phillips and Ouliaris (1990) and Phillips (1991) propose efficient corrections for the long-run variance of u_t due to serial correlation. This is important in our work as financial time-series typically exhibit heteroskedasticity and short term serial correlation. Despite the natural fit of our model into this framework there are some drawbacks. Models of cointegration rank are not nested and cannot be tested, and we rely on other frameworks for this purpose.

2.3.1.2 The autoregressive framework

The model for the N dimensional multivariate process is now

$$\Delta p_t = \alpha(\beta'x_{t-1} - E(\beta'x_{t-1})) + u_t$$

where u_t are mean zero, finite variance i.i.d. errors. α and β are $N \times r$ matrices. The cointegration rank is $r \leq N$ and the space of cointegrating vectors is the space spanned by the rows of β' . The framework allows modeling of the short term dynamics t

owards the long-run relationship ($\beta'x_{t-1} = E(\beta'x_{t-1})$) through adjustment speeds α . Statistical inference in this framework is based on the Gaussian likelihood

(Johansen, 1995), and likelihood ratios are used when testing various hypotheses. The fact that models of different cointegration rank r are nested enables sequential testing of specific ranks against alternatives of higher rank. Another benefit over the regression framework is that test statistics of cointegration are independent of which variable is chosen as a dependent variable in any regression. This is not the case with the Engle–Granger framework.

2.3.1.3 The unobserved component framework

The process p_t is now given by

$$p_t = \xi \eta' \sum_{i=1}^t \varepsilon_i + u_t$$

where u_t is a mean zero, finite variance i.i.d process independent of ε_t (so that Δp_t contains an MA(1) process). The $N - r$ common stochastic trends are the elements of the vector $(\sum_{i=1}^t \varepsilon_i)$. The parameters are related to the autoregressive framework via the relationships $\xi = \beta_{\perp}$ and $\eta = \alpha_{\perp}$. Again models of different rank (or the numbers of common trends) are nested, via the rank of ξ . Whereas in the autoregressive framework models of rank r can be tested against alternatives of higher rank, here specific numbers of common trends are tested against alternatives of higher numbers of trends. The testing of rank proceeds in the “opposite” direction to the Johansen tests. Note that in our model there is a single common trend equivalent to the probability of a particular political outcome implying $N - 1$ cointegration vectors. Shin (1994) proposed a residual based test for the bivariate case (or univariate regression) where the null is one of a single common trend (cointegration). Tests of higher numbers of common trends in the general multivariate case have been proposed by Nyblom and Harvey (2000).

2.3.2 Non-linear cointegration

The theory of linear cointegration is well developed. That of non-linear cointegration less so, with many open questions. In fact it is not entirely clear how non-linear cointegration should be defined. Generalisations of the properties of short memory ($I(0)$) and long memory ($I(1)$) are required. See Escanciano and Escribano (2009) and Wang (2015) for surveys of the field.

Much of the work extending cointegration to a non-linear setting involves extending ECMs to non-linear error correcting models (NECs). This in general in-

volves replacing the linear gap function $\alpha(\beta'x_{t-1} - E(\beta'x_{t-1}))$ to some non-linear reaction function $f(\beta'x_{t-1} - E(\beta'x_{t-1}))$. The argument of $f(\cdot)$ is a stationary linear cointegrating vector of non-stationary variables. Of potentially more interest to our application is the study of nonlinear cointegrating regression models. These have been studied in both the parametric and non-parametric setting.

Turning to our particular problem we recall that the pricing model is summarised by the quadratic relationships

$$p_t = p_0 + \Delta p.PB_t + \lambda.PB_t.(1 - PB_t) + \epsilon_t$$

where p_t is the non-stationary financial asset price, PB_t is the non-stationary (bounded) price of the betting contract and ϵ_t is stationary. However, this form of the model relies on risk neutral pricing in betting markets and mean-variance preferences in financial markets. Different preferences will lead to potential different non-linear relationships

$$p_t = g(PB_t) + \epsilon_t$$

where $g(\cdot)$ is a non-linear function. Recall also that assuming risk neutral pricing in both markets yields a linear $g(\cdot)$. Karlsen et al. (2007) estimate the function $g(\cdot)$ nonparametrically using the Nadaraya–Watson estimator. The asymptotics have been worked out and the distribution of $\hat{g}(\cdot)$ is Gaussian. This opens the door to specification testing. Unfortunately rates of convergence in the non-stationary case are slower than the stationary case. $\hat{g}(\cdot)$ converges to $g(\cdot)$ at a rate of $T^{1/4}$. Other developments of this theory include Gao et al. (2009). This paper presents a bootstrap scheme and test of whether $g(\cdot)$ is of a known parametric form $g(\cdot, \beta)$. Wang and Wang (2013) extends the Kernel estimate of $\hat{g}(\cdot)$ to include ϵ_t as having a nonlinear nonstationary heteroskedastic process. This is of particular relevance for applications with financial time series.

In this chapter we are primarily interested in testing whether our theory holds. We seek extensions of the tests with a null of ϵ_t non-stationary, such as Dickey–Fuller and associated tests, to the non-linear setting. Rejection of such a test in favour of stationary ϵ_t would provide strong evidence of our theory holding. KPSS-type tests of stationarity in the nonlinear parametric case (see Choi and Saikkonen (2010)) are available. However, we find no tests in the literature with a null of no nonlinear cointegration in the nonparametric case.

Of secondary interest in our model is whether there is significant risk aversion present. This is equivalent to the cointegration relationship deviating significantly from linearity. The methods presented in both Gao et al. (2009) and Wang and Wang (2013) do provide such tests conditional on cointegration holding. However, we study periods of a few hours in an overnight session of the markets. We sample prices every minute. This leads to a T being of the order of hundreds. As convergence of $\hat{g}(\cdot)$ is of the order of $T^{1/4}$ we would not expect convergence to apply. It does not appear that the application of non-parametric methods is suitable for this chapter.

Luckily though the pricing framework developed in our model does provide a straight-forward way to proceed in the linear setting with minimal restrictions. Returning though to standard mean variance risk preferences yields the convenient form for $g(\cdot)$ of

$$\begin{aligned} g(PB_t) &= p_0 + \Delta p \cdot PB_t + \lambda \cdot R_t \\ R_t &= PB_t(1 - PB_t). \end{aligned} \tag{2.3.1}$$

R_t represents the non-stationary measure of political risk and is proportional to the variance of the Bernoulli variable E representing the political event. $g(\cdot)$ is now linear in PB_t and R_t . We can also recover this form of the cointegrating relationship using a second order quadratic approximation to $g(\cdot)$, making use of the fact that once the event has been realised ($PB_t = 0$ or 1) there is no longer political risk and the contribution to the relationship must vanish. Using this form of the relationship means we can now exploit tests within the well developed field of linear cointegration theory. Testing for significant risk preferences (and non-linearity in PB_t) is equivalent to testing whether λ is significant.

We note that R_t is a quadratic function of P_t . This means that the Granger representation theorem does not apply and so the Johansen framework is not strictly valid. There are methods in the literature that yield specification tests for non-stationary quadratic regression models (see Wagner and Hong, 2016). The application of such a test is beyond the scope of this chapter and we will rely on tests solely within the linear setting. Furthermore there is no test that has a null of no-cointegration we can find.

We argued earlier that P_t behaves as an $I(0)$ process until a value of zero or unity is achieved (at the point at which the political event is resolved). R_t will also not be stationary although not necessarily $I(1)$. We do test both R_t and its first difference ΔR_t with both the Philips–Perron and KPSS tests. For all cases presented in this chapter, test results for the risk variable R_t are consistent with an $I(1)$ variable.

Using the linear framework also allows us to easily deal with the monotonic restriction, $|\lambda| \leq |\Delta p|$. When the estimate of 2.3.1 yields a non-monotonic result with $|\hat{\lambda}| > |\hat{\Delta p}|$ we can simply re-run the regression with the non-stationary explanatory variables PB_t and R_t replaced with a single explanatory variable. This will be either $PB_t + R_t$ or $PB_t - R_t$ depending on whether PB_t and R_t have the same sign. This variable is monotonic and recovers the estimate at the boundary where $|\lambda| = |\Delta p|$.

2.3.3 Statistical tests

This chapter will make use of several statistical tests derived from the three linear cointegration frameworks. They are set out as follows.

2.3.3.1 Phillips and Ouliaris test for the presence of cointegration

We apply the Engle–Granger methodology via univariate regressions of each price against that of the betting market contract

$$p_t = p_0 + PB_t \times \Delta p + \epsilon_t.$$

The presence of a unit root in ϵ_t is tested via a residual regression. However, we use the Philips–Perron test statistics, Z_α and Z_t rather than Dickey–Fuller. These include a non-parametric adjustment to the long-run variance which is robust to misspecified serial correlation and heteroskedasticity. If the statistic is below the critical value we reject the null of a unit root in the residual and conclude that cointegration is present between the asset price and the betting contract.

2.3.3.2 Johansen max eigenvalue test

The Johansen methodology considers the error-correcting form for our multivariate price process p_t

$$\Delta \mathbf{p}_t = \boldsymbol{\alpha}(\boldsymbol{\beta}' \mathbf{p}_{t-1} - E(\boldsymbol{\beta}' \mathbf{p}_{t-1})) + \sum_{i=1}^{k-1} \boldsymbol{\Gamma}_i \Delta \mathbf{p}_{t-i} + \mathbf{u}_t.$$

This specification includes lags of $\Delta \mathbf{p}_t$. The cointegration rank is the rank of $\boldsymbol{\beta}$ and models of smaller rank are sequentially nested in models of higher rank. Johansen demonstrated that the Gaussian Maximum Likelihood for a given rank r is a simple expression based on the smallest r eigenvalues of the canonical correlation matrix relating to a reduced rank regression of the above equation. The Likelihood Ratio depends on the $N - r$ largest eigenvalues. There are two forms of rank test, both where the null hypothesis is cointegration of rank $r \leq r^*$. In the trace test the alternative is that cointegration is full rank. This is immediately discounted in our application as it is equivalent to all prices being stationary. The second form, the max eigenvalue test, has the alternative higher rank $r = r^* + 1$. This test is based on λ_{r+1} , the $(r + 1)^{th}$ smallest eigenvalue. The asymptotic distribution of the test statistic is non-standard and depends on Brownian motion. It has been tabulated in Johansen and Juselius (1990). Rejection of rank $r = r^*$ occurs when λ_{r+1} is “large”. Care will be required when choosing the lag length k as tests will be biased and inconsistent if the model is misspecified.

2.3.3.3 Johansen constraint test

Likelihood Ratio tests can be formulated for restrictions on model parameters in the Johansen framework. We use this to test whether risk aversion is significant in equation 2.2.9 (i.e. $H_o : \lambda = 0$). This is done by considering the trivariate prices

$$\mathbf{x}_t = (p_t, PB_t, R_t)'$$

and checking whether the coefficient in the last dimension (the dimension of risk) are not significantly different from zero.

2.3.3.4 Nyblom and Harvey common trends test

Nyblom and Harvey (2000) work in the unobserved component framework and consider the multivariate local level model

$$\begin{aligned}
p_t &= \mu_t + \epsilon_t, & \epsilon_t &\sim IID(0, \Sigma_\epsilon) \\
\mu_t &= \mu_{t-1} + \eta_t, & \eta_t &\sim IID(0, \Sigma_\eta).
\end{aligned}$$

Σ_η is the variance of the multivariate disturbance driving the unobserved random walks. Cointegrating vectors β satisfy $\Sigma_\eta \beta' = 0$. The rank of Σ_η is equal to the number of common trends. Tests for a given number of common trends k^* are proposed with alternatives of higher rank. Non-parametric corrections that are robust to serial correlation and heteroskedasticity are included. Test statistics are now based on the sum of the $N - K^*$ smallest eigenvalues of a matrix. When this statistic is large we reject $k = k^*$ in favour of $k > k^*$. The theoretical model presented in the previous section assumes $k = 1$. As before, we discount the situation $k = 0$ (stationary prices) but can test $k = 1$ against $k > 1$. Rejection of this null would be a rejection of our theory. Failure to reject is consistent with our theory but would give less support than a rejection of cointegration rank, $r = N - 2$ in favour of $r = N - 1$, under the max eigenvalue test.

A note on lag specification

The above frameworks require lag length choices to be made when specifying tests. For both the Phillips and Ouliaris residual regressions and the common trends test, non parametric adjustments are included that are robust to misspecified serial correlation in the disturbances. All that is required for these methods is to choose a lag truncation parameter in the estimator of the long-run variance. We use the Newey and West (1994) plug-in procedure of $\left[4. \left(\frac{T}{100}\right)^{2/9}\right]$. The Johansen rank and constraint tests are not robust to model misspecification although we note that there is some evidence that testing constraints are not unduly effected in the presence of serial correlation, Silvapulle and Podivinsky (2000). Schwert (1989) suggests choosing the lag length that minimises the AIC or BIC up to a maximum lag of $\left[12. \left(\frac{T}{100}\right)^{1/4}\right]$. The choice of lag involves a trade off between accuracy of specification and statistical power. Models of higher numbers of lags and hence parameters will necessarily be more accurate (as shorter lag lengths are nested) but will result in a loss of statistical power and we note that cointegration tests often do lack power. We start by only adding lags if they result in a statistically significantly different model, as implied by a likelihood ratio test. Lags are added to

the base (zero lag) model until the likelihood ratio test fails to reject the restricted model. As the max eigenvalue test is not robust in the presence of serial correlation, we check that the residuals are indeed not correlated (according to a Portmanteau test). If they are, then further lags are added until serial correlation is no longer present.

Validity of cointegration tests in our setting

The use of the above cointegration tests to the time period of the overnight session following an election is an unusual and innovative application of a method that is usually applied to long calendar time periods. The use in a short calendar time period requires comment. Our pricing model applies only for the period under study, which is a few hours in length. It is a but brief period in the history of the financial assets under study. However, we are not concerned with the time series properties of prices outside of this small window and do not seek to measure them. Indeed our model implies that the property of cointegration is completely unique to the overnight sample. Further, we would expect volatility to be unusually high as the results of an election were announced. If we were trying to make inferences about longer term properties of the time series we study, then we would be in trouble. We would require that the series exhibited ergodicity. This means that we could make inferences about the long term nature of the time series from a short term, but large, sample. This does not always hold for financial time series. In particular, such series exhibit highly varying and often persistent variance and non-ergodicity. The same is also true of average returns, where the fat tailed nature of the distribution makes estimation harder. Most of the cointegration literature is restrictive and only allows for ergodic time-varying variance. As it stand though, we believe an assumption of constant variance (and other properties) throughout the election night is not necessarily a bad one. The fact that behaviour will be very different to other periods does not mean that behaviour itself varies greatly throughout the night.

Cointegration tests have historically been applied to daily data. The large sample asymptotics assume a fixed time gap between observations, Δ , with the time horizon, $T \rightarrow \infty$. In our application, the validity of the standard cointegration tests relies on the use of asymptotic results for a large sample taken from a small calendar time period. We believe this is valid. We cannot increase T any further, as the hypothesis of cointegration in our model only holds for the few hours of the

night itself. We use one-minute returns leading to a respectable sample size of a few hundred observations. However, tests based on high-frequency asymptotics, where $\Delta \rightarrow 0$ with $T \rightarrow \infty$ may be more appropriate.

The stylized facts of high frequency financial time series are of time-varying variance with jumps, a leverage effect and price jumps. These properties can lead to significant distortions of the power and size of the standard 2-step residual based cointegration tests, Krauss and Herrmann (2017). Time varying variance can lead to size distortions, with a jump in the long term variance causing spurious cointegration, Noh and Kim (2003). There is an active area of research concerning testing for a unit root in the presence of time varying or non-stationary variance. Beare (2018), Cavaliere and Taylor (2007), Cavaliere and Taylor (2008), Cavaliere and Taylor (2008) and Cavaliere and Taylor (2009) provide robust solutions to this issue by performing various time transformations to the data (referred to as deflation). The other issue is that high frequency data typically exhibit large numbers of price jumps. Krauss and Herrmann (2017) document the fact that price jumps deteriorate the power of cointegration tests. Gregory and Hansen (1996), Maki (2012) and Hatemi-J (2008) provide solutions which are robust to different numbers of price jumps. However, these approaches are limited by the fact that jumps need to happen at deterministic times, which is not consistent with the stylized facts of high frequency data. As we shall see in the results section, we have no problem with the power of the standard tests, as we in general comfortably reject the null of no-cointegration in our data. However, the question of distorted size due to a time-varying variance is not completely answered with the approach in this chapter. The robustness of our results could be improved by applying more recent methods from the literature. This is beyond the scope of this chapter. However, we suggest an application of the test in Clinet and Potiron (2019) may be appropriate. This adapts the standard two-step procedure of Phillips and Ouliaris (1990) to test the null of no-cointegration in the high frequency setting, where variance is not fixed and there are an unlimited number of jumps. This uses deflated (to account for varying variance) and truncated (to accommodate an unlimited number of jumps) time-series and converges under high frequency asymptotics. The model they use is particularly applicable to the overnight results session, as we would indeed expect there to be a large number of jumps. These could correspond to the announcement of election results from different voting areas and constituencies.

2.4 Results

In this section we evaluate our theory on real world data from some political events from the last few years. The choice of events is important. For our model to apply, the results of the elections need to come during overnight hours. This rules out, for instance, the 2020 US presidential election, where various legal challenges and recounts in several states took days and even weeks to resolve. Our model also has no applicability if the result is realised instantaneously. We require a meaningful period of time for the information flow to occur. This rules out UK general elections. This is due to the very high accuracy of the exit poll. These polls measure how people declare they have voted on the day itself, at a selection of particular, secret, polling stations. They are much more accurate than any pre-election polling, Curtice et al. (2011), due to the fact that there is no measurement error of respondents. The polls are released just after votes closed and are effectively an announcement of the winner. Instead we chose three elections where we believe our model is likely to apply. In each case we consider an actively traded binary betting contract along with a collection of heavily traded financial assets. We begin our investigation with the first great political shock of 2016. This was the United Kingdom European Union membership referendum, commonly referred to as the Brexit referendum.

2.4.1 The Brexit referendum

On 23rd June 2016 the UK voted in a country wide referendum to leave the European Union. This was one of only three UK wide referendums. The first was in 1975 and involved a vote to join the European Community. This is known as the common market and is what become the European Union. The second was the United Kingdom Alternative Vote referendum which was rejected by a wide margin. The third was the Brexit plebiscite. Turnout was historically high at 72% and the result was narrow: 51.9% to leave the EU versus 48.1% to stay.

The vote was split up into a large number of voting areas (382) and each area announced at different times throughout the night as their counts were finalised. The result was unexpected. There was widespread polling data that showed a small but consistent lead to stay in the European Union ('Remain'). For example, a poll was conducted by YouGov which was published shortly after voting closed




Figure 2.1: Betfair contract for ‘Remain’ on the night of the Brexit referendum.

at 10pm, YouGov (2016). This showed the vote-share for leave at 48.4% with a standard sampling error of 3%.

There was an actively traded political market that traded both up to the day of voting and overnight as results were announced. Contracts that paid out £1 in the event of both ‘Remain’ and ‘Leave’ were listed on the Betfair Exchange market. This operates like a limit order book. There did not appear arbitrage opportunities in the exchange in that the sum of the prices of the contracts do not deviate sufficiently from £1.⁷ Around £130m was wagered in total with £50m changing hands on the night.

There was widespread belief that the country would vote to remain in the EU. Figure 2.1 shows the price action of the contract price for Remain. Voting closed at 22:00 on 23rd June.⁸ The YouGov poll on the day was released shortly afterwards. The first result released was for Gibraltar around 23:36. This was

⁷Owning both one contract for Remain and one for Leave guarantees a payout of £1.

⁸Times in this section are all quoted in British Summer Time. This was the current timezone in the UK on this date

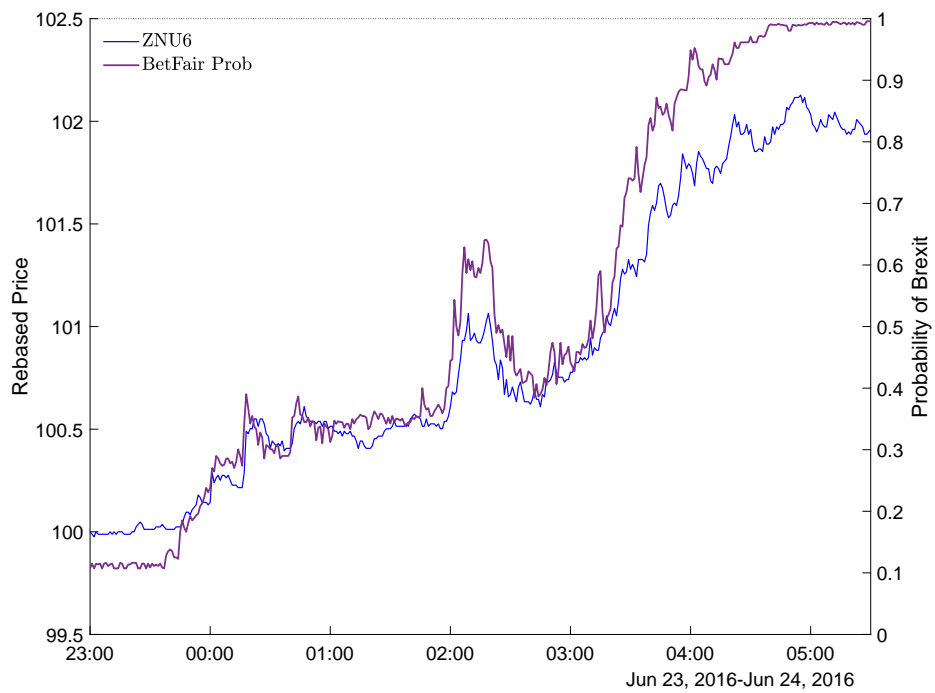

Table 2.1: Financial assets and changes for the Brexit referendum.

Symbol	Description	$p_{t=1}$	$p_{t=T}$	Δ
ESU6	E-mini S&P500 Future	2115.75	2000.00	-5.5%
ZNU6	10-Year T-Note Future	130.703	133.266	2.0%
CLU6	Crude Oil Future	50.89	47.82	-6.0%
GBPUSD	British Pound US Dollar Cross	1.5007	1.3242	-11.8%
USDJPY	US Dollar Japanese Yen Cross	106.61	100.93	-5.3%

inconsequential. Gibraltar is an overseas territory located at the southern tip of the Iberian Peninsula bordering Spain. As expected, the electorate there voted overwhelmingly in favour of Remain (96%). As can be seen, this did not affect the prices in the betting market. The risk neutral probability of remaining in the EU seen close to 90%. Meaningful results started to be announced from midnight with Newcastle upon Tyne being the first to announce. As can be seen from Figure 2.1, prices had already started to move a little against Remain from 23:45, possibly due to information leakage, or to private polling conducted by some hedge funds. There was a large quick move from 73% to 61% for Remain between 00:16 and 00:18 as Sunderland announced. The result there showed a lead for Brexit over Remain of 21% versus an expectation of around 6%. Prices moved against Remain for the next few hours, with a particular collapse around 2am which mostly recovered shortly afterwards. However, by 4am Remain was trading at under 10% and a probability for Brexit of 99% was implied at 05:21. The BBC finally projected Brexit at 5:39am and there was no doubt that the country had voted to leave the European Union.

2.4.1.1 Financial assets

We consider five financial assets for this event. Two currencies, one stock index future, one commodities future and one fixed income future. All asset classes are covered by this collection and all but GBPUSD was the most liquid leading indicator of those asset classes. GBPUSD is included as it is especially relevant for the UK and a country-specific indicator of the health of the economy. All futures were listed on the Chicago Mercantile Exchange. Foreign exchange and futures are studied as these were open for trading during the results session. Cash markets do not generally trade overnight and the UK specific FTSE100 future only opened part way through the night and so is excluded. We consider the period 23:00 BST



Note: Assets that depreciated are shown versus the Remain contract price whereas the asset that appreciated is shown versus the contract price for Brexit.

Figure 2.2: Rebased financial asset prices versus Betfair contracts.

Table 2.2: Results of the Phillips–Ouliaris Z_t and Z_α tests for cointegration.

Symbol	\hat{p}_0	$\hat{\Delta p}$	$p_{Z_t}(a = 0)$	$p_{Z_\alpha}(a = 0)$
ESU6	2134.0	-115.6	0.008***	0.003***
ZNU6	130.32	2.97	0.001***	0.001***
CLU6	51.38	-3.37	0.006***	0.004***
GBPUSD	1.5199	-0.1751	0.001***	0.001***
USDJPY	107.28	-6.15	0.001***	0.001***

The first stage regression is $p_t = \hat{p}_0 + \hat{\Delta p}PB_t + \hat{\epsilon}_t$

The residual regression is $\hat{\epsilon}_t = a\epsilon_{t-1} + \eta_t$

on 23rd June to 05:30 on 24th June and sample prices every minute. During these hours all contracts are open,⁹ the only information released to the market is that contained in the vote results and the betting markets converge to certainty.¹⁰ Table 2.1 lists the five assets with a description, their starting and ending prices and the percentage change. As can be seen four assets depreciated whilst the treasury future appreciated. This is as expected as the result was not expected and was considered negative for trade and hence the economy. Generally fixed income assets appreciate in times of economic uncertainty as both expectations of future interest rates fall and money moves out of risky assets. The Japanese yen is considered a safe asset so appreciated against the USD whilst GBP is seen as risky versus the USD and depreciated.

Figure 2.2 plots assets that depreciated against the betting contract for Remain and the treasury future is plotted against the contract for Brexit. The financial assets do appear to be moving more or less in lock-step with the betting markets. This is pleasing as this is implied by our theory. However, we now turn to statistical tests of the theory in the following sections.

2.4.1.2 Evidence for cointegration

We first test the cointegration of each asset price with the betting contract for Brexit. Table 2.2 shows the results of the Phillips–Ouliaris Z_t and Z_α tests for cointegration. Both tests reject the null of a unit root in the residual ϵ_t at the 99% level in favour of stationarity. The results of the Johansen max eigenvalue test are

⁹Note, the E-mini S&P500 Future hit a trading limit shortly after 05:30 on 24th and was thus put into auction.

¹⁰The futures were closed for an hour between 22:00 and 23:00 British Summer Time. This was straight after the vote closed but before results were announced.

Table 2.3: Results of the bivariate Johansen max eigenvalue test.

Symbol	\hat{p}_0	$\hat{\Delta p}$	$k - 1$	$p(r = 0)$
ESU6	2134.0	-121.9	1	0.024**
ZNU6	130.32	3.00	1	0.001***
CLU6	51.34	-3.48	2	0.090*
GBPUSD	1.5188	-0.1784	2	0.003***
USDJPY	107.24	-6.22	0	0.001***

The regression is $\Delta \mathbf{p}_t = \alpha(\beta' \mathbf{p}_{t-1} - c_0) + \sum_{i=1}^{k-1} \Gamma_i \Delta \mathbf{p}_{t-i} + \mathbf{u}_t$
 $r = \text{rank}(\beta)$, when $r = 1$, $\beta' = (1, -\Delta p)$, $c_0 = -p_0$

Table 2.4: Results of the multivariate Johansen max eigenvalue test.

r	h	stat	cValue	eigVal	pValue
0	1	43.32	40.96	0.106	0.027**
1	1	36.13	34.81	0.089	0.035**
2	1	31.15	28.59	0.077	0.023**
3	0	17.77	22.30	0.045	0.191
4	0	7.92	15.89	0.020	0.593
5	0	3.73	9.16	0.010	0.523

The regression is $\Delta \mathbf{p}_t = \alpha(\beta' \mathbf{p}_{t-1} - c_0) + \sum_{i=1}^{k-1} \Gamma_i \Delta \mathbf{p}_{t-i} + \mathbf{u}_t$
 $r = \text{rank}(\beta)$

shown Table 2.3. The maximum lag length $k - 1$ was chosen using the procedure outlined in subsection 2.3.3. This test also rejects a cointegration rank of 0 in favour of a single common trend for all assets, albeit at lower significance levels for the S&P500 and crude oil futures prices. During a typical trading day there will be information that would persistently affect the financial assets over and above that which affects the odds of the UK voting to leave the European Union. The fact that the null hypothesis of no cointegration is firmly rejected in favour of a single common trend in the hours after the Brexit referendum is strong evidence in favour of our theory.

We consider next the 6-dimensional multivariate price process \mathbf{p}_t which includes all five financial assets and the betting contract. Our pairwise cointegration tests show strong evidence of each financial asset being cointegrated with the betting contract. Within the space of time series, cointegration is a transitive relation. That is, if x_t is cointegrated with y_t and y_t is cointegrated with z_t , then x_t is cointegrated with z_t . If every financial asset is cointegrated with the betting market, then every financial asset is pairwise cointegrated, cointegration is of rank $n - 1$ and there is a single common trend. This is implied by our theory. We test this proposal directly using the Johansen max eigenvalue test. Using the procedure outlined in subsection 2.3.3 we chose a lag length in the error correction model of 3. Results of the tests for various cointegration ranks are shown in Table 2.4. Cointegration ranks of 0, 1 and 2 are rejected in favour of higher ranks at the 95% level. This suggests that the rank is at least 3. Our theory predicts a rank of 5. The failure to reject ranks 3 and 4 in favour of rank 5 is a rejection of our theory. It may be the case that our model does not sufficiently describe the behaviour observed after the Brexit referendum. However, it may also be the case that the Johansen's test may also have insufficient power to reject ranks of 3 and 4. As another check, we compute the test statistic used in the Nyblom and Harvey test for a single common trend. If the statistic is above a critical level the single common trend is rejected in favour of a higher number of common trends. The statistic is 0.0045. This is well below the level of rejection at the 90% level¹¹. Thus the null hypothesis, of a single common trend and cointegration rank of $n - 1$, is not rejected at any meaningful significance level. Although it would be preferable for the Johansen test to reject cointegration ranks of 3 and 4 in favour of a single common trend, the results of this

¹¹Nyblom and Harvey (2000) only compute critical values for $n \leq 4$. However, we can avoid the simulation of critical values here. We note that critical values increase with n . The published value of 0.427 for the 90% when $n = 4$ is a lower bound for higher $n > 4$.

Table 2.5: Results of the trivariate Johansen constraint test for risk aversion.

Symbol	$\hat{\lambda}$	$k - 1$	$p(\lambda = 0)$
ESU6	-2.165	0	0.913
ZNU6	-0.442	3	0.084*
CLU6	-0.718	0	0.294
GBPUSD	0.018	2	0.508
USDJPY	1.274	1	0.237

The regression is $\Delta \mathbf{p}_t = \boldsymbol{\alpha}(\boldsymbol{\beta}' \mathbf{p}_{t-1} - \mathbf{c}_0) + \sum_{i=1}^{k-1} \boldsymbol{\Gamma}_i \Delta \mathbf{p}_{t-i} + \mathbf{u}_t$

$$\mathbf{p}_t = (p_t, PB_t, R_t)'$$

$$r = \text{rank}(\boldsymbol{\beta}) = 1, \boldsymbol{\beta}' = (1, -\Delta p, -\lambda), c_0 = -p_0$$

test taken together with the pairwise tests and the common trends test is strong evidence that the pricing theory outlined in section 2.2 held during the night after the Brexit referendum.

2.4.1.3 Risk Aversion

Next we turn to the question of risk aversion. Manasse et al. (2020) derive a model of foreign exchange prices and the Betfair contracts for the months before the referendum. This is based on cointegration and includes a component related to risk. The results of the previous subsection suggest that we do need to deviate from risk neutrality to explain the high frequency price action observed during the hours when the referendum results were released. However, could a model that includes a non-zero risk aversion parameter provide a better fit? To answer this question we consider trivariate systems $x_t = (p_t, PB_t, R_t)'$ for each financial asset p_t . As described in the pricing model the third component in x_t relates to risk, and is identical to that considered in Manasse et al. (2020). With or without risk aversion our theory implies a single cointegrating vector $\boldsymbol{\beta}' = (1, -\Delta p, -\lambda)$ and two common trends PB_t and R_t ¹². We test the restriction $\lambda = 0$ in the Johansen framework with $r = 1$ using the constraint test. The results are shown in table 2.5. We do not reject the null hypothesis of risk neutrality for any asset at the 95% level and only reject the hypothesis for a single asset, the treasury future at the 90% level. The constraint test p-value for the pound of 0.51 suggests there is no evidence at all for the *GBPUSD* exhibiting risk aversion behaviour with regards to Brexit. This is in contrast to the conclusions of Manasse et al. (2020).

¹² PB_t cannot possibly be cointegrated with $R_t = PB_t(1 - PB_t)$ if it is non-stationary.

However, this author takes issue with the application of a cointegration theory for longer term periods. The basis for the model is that the only determinant of the *GBPUSD* price in the weeks and months preceding the Brexit vote is the result of that referendum. Whilst we agree that the probability of the Brexit result is a large determinant of prices, and is the only determinant in the overnight hours after the vote, we do not believe it is the only information affecting exchange rates for longer periods. Their theory assumes no other news beyond that relating to the referendum affects the British Pound and United States dollar exchange rate in the months up to the vote. For example, this would imply that any and all news about the health of the US economy, be it consumer demand, trade barriers, protectionism etc., would have no effect on *GBPUSD*. We find this implausible and postulate that including risk aversion is simply an exercise in over-fitting to avoid rejecting a model that should never have been applied on this timescale. We suggest that other models such as vector auto-regressive or factor models would be more economically justified for longer periods of time. This will be studied further in chapter 3.

Table 2.6: Fitted short run parameters in the error correction model for Brexit, with tests on the short run dynamics.

Short run regression results

$$\Gamma_1 = 10^{-3} \times \begin{pmatrix} -190^{**} & - & -3860^{***} & 13000^{***} & - & -20600^{***} \\ -3.44^{***} & - & 44.1^* & - & - & 185^* \\ 6.11^{***} & -0.312^{***} & -159^{***} & -2890^{**} & - & - \\ 0.14^* & - & - & -105^* & - & -28.9^{***} \\ 11.4^{**} & - & - & 6720^* & -116^{**} & -1640^{**} \\ - & - & - & -452^* & - & - \end{pmatrix}^1$$

Based on $\Delta p_t = \alpha(\beta' p_{t-1} - c_0) + \Gamma_1 \Delta p_{t-1} + u_t$

$p_t = (\text{ESU6}, \text{ZNU6}, \text{CLU6}, \text{GBPUSD}, \text{USDJPY}, PB_t)'$

¹ Statistically insignificant coefficients omitted

Null hypothesis	LR-Test	
	Chi-square	Prob
(ESU6, ZNU6, CLU6, GBPUSD, USDJPY)' does not cause PB_t	4.84	0.436
PB_t does not cause (ESU6, ZNU6, CLU6, GBPUSD, USDJPY)'	11.4^{**}	0.0441




Figure 2.3: Cointegration errors with autocorrelation functions for Brexit.

2.4.1.4 Deviations from long term relationships

The strong evidence for cointegration in the hours after the referendum is a pleasing result and agrees with our theoretical model. Our assumptions include that the only persistent affects on asset prices are related to the probability of voting to the leave the European Union. This gives rise to the single common trend. Another main assumption is that weak market efficiency holds. However, the presence of significant autoregressive terms in the error correction model of the Johansen framework suggests that efficiency may be less clear cut.

We can study short term deviations from the long term relationship in two ways. The first is via the coefficients of the VECM of the full multivariate system. The second is via the univariate errors of the cointegration regression. The full VECM is

$$\Delta \mathbf{p}_t = \boldsymbol{\alpha}(\boldsymbol{\beta}' \mathbf{p}_{t-1} - \mathbf{c}_0) + \sum_{i=1}^p \boldsymbol{\Gamma}_i \Delta \mathbf{p}_{t-i} + \mathbf{u}_t$$

where

$$\mathbf{p}_t = (\text{ESU6}, \text{ZNU6}, \text{CLU6}, \text{GBPUSD}, \text{USDJPY}, PB_t)'$$

It is found that an error correcting model with a single short run autoregressive term $\boldsymbol{\Gamma}_1$ is significantly different to one with no term but that higher numbers of lags are not significant. Table 2.6 shows the significant fitted values of $\boldsymbol{\Gamma}_1$ in the VECM. Deviations from the long-run relationships ($\boldsymbol{\beta}' \mathbf{p}_{t-1} - \mathbf{c}_0 = 0$) cause adjustments in both markets through the first term, $\boldsymbol{\alpha}$. The second term, $\boldsymbol{\Gamma}_1 \Delta \mathbf{p}_{t-1}$ governs short-run dynamics. Of particular interest is whether the betting markets lead the financial markets and/or vice-versa. For short run dynamics this can be tested by testing the coefficients of the coefficient matrix $\boldsymbol{\Gamma}_1$. Jointly testing the significance of the first $n-1$ values of the final column of $\boldsymbol{\Gamma}_1$ is a test of whether short run changes in the betting markets cause short run changes in financial markets. Testing the first $n-1$ values of the last row tests whether financial markets cause betting markets in the short run. LR tests for these hypotheses are also shown in table 2.6. There is strong evidence that short run changes in betting markets cause short-run changes in financial markets with the null of no causality rejected at the 5% level. However, there is no significant causation in the other direction. This is a very interesting result. Short term deviations from the long term relationship that revert occur as the betting and financial markets do not move exactly in lock-step. However, during the establishment and convergence of these deviations the betting markets are leading the financial markets. This suggests the betting markets are discounting political information more quickly (and on the scale of minutes) than the financial markets.

We now study inefficiencies via the simpler residuals of the univariate cointegrating regressions. This error is a linear combination of two asset prices (a financial asset and a betting contract) and can be traded. The return of the error can be created by holding the two assets in a ratio equal to the cointegration ratio. It is also largely risk free, at least with respect to political risk. The exposure to the result of the referendum in one asset is hedged with equal and opposite exposure in the other. Serial correlation in the cointegration error implies that deviations

from the long term relationship (long term in this context being a few hours) can be predicted. This again demonstrates violations of weak market efficiency on a short term timescale.

The estimated cointegration error is $\hat{\epsilon}_t = p_t - \hat{p}_0 - PB_t \times \hat{\Delta}p$. We consider the quantity $\hat{\epsilon}_t/|\hat{\Delta}p|$ which is the cointegration error normalised to units of the betting contract. This can be more readily compared across different assets whose prices have different magnitudes. The normalised cointegration errors with the sample Auto-Correlation Function and robust Bartlett intervals for each asset are plotted in Figure 2.3.¹³ This shows significantly positive autocorrelations out to around 20 minutes. This is consistent with deviations from equivalence of beliefs in the two markets, and hence weak market efficiency, of the order of minutes to tens of minutes on an ex-post basis. We fit autoregressive models with varying numbers of lags to the estimated error $\hat{\epsilon}_t$. Note that the constant term is fixed to zero as $E(\hat{\epsilon}_t) = 0$ by construction. Results are shown in Table 2.7. A single lag provides a good fit to the data. However likelihood ratio tests show that two lags produces a significantly different model to ones with a single lag for the S&P500 future and the pound but that those with more lags are no different. These results contradict EMH. The possibility of profiting from them systematically is explored in the next section.

Table 2.7: Estimated autoregressive models for cointegration errors.

lag	ESU6	ZNU6	CLU6	GBPUSD	USDJPY
$\hat{\epsilon}_{t-1}$	0.746*** (0.039)	0.761*** (0.029)	0.897*** (0.021)	0.675*** (0.046)	0.818*** (0.020)
$\hat{\epsilon}_{t-2}$	0.166*** (0.038)			0.166*** (0.044)	

The regression is $\hat{\epsilon}_t = \psi_1 \hat{\epsilon}_{t-1} + \dots + \psi_k \hat{\epsilon}_{t-k} + \eta_t$.
 $\hat{\epsilon}_t$ is estimated from the first stage regression.

2.4.1.5 Profit opportunities

We now turn to the possibility of profiting from mean reverting deviations to the long term relationship. Rather than focus on trading assets outright (or “naked” in trading parlance) via predictions from an error correction model, we focus on

¹³The Bartlett intervals are adjusted to allow for serial correlation in the variance of ϵ_t . Such serial conditional heteroskedasticity is expected as this is a financial time series.




Figure 2.4: Bollinger Band trading strategy for GBPUSD.

trading the cointegration error. This is effected by buying or selling the financial asset against an opposing position in the betting contract with sizes equal in ratio to the cointegrating ratio.¹⁴ This is effectively politically risk free and involves taking positions in pairs of contracts simultaneously rather than larger numbers of assets outright which be would exposing oneself to greater risk.

We use a modification of the common Bollinger Band trading signal. When used to trade non-stationary price series, Bollinger Bands are used to produce directional trading indicators. A corridor around a moving average of the price of the asset is constructed by adding and subtracting twice the sample standard deviation, σ_t , of the price calculated along the length of the preceding moving average window. For a contrarian strategy the signal will be to sell when above the upper band and to buy when below the lower band. All positions are closed out once the betting contract has converged to certainty. As the cointegration error ϵ_t has zero expectation we do not calculate the corridor around a moving average but about

¹⁴To be clear $\hat{\Delta}p$ notional of the betting contract is traded for every unit notional of financial asset exposure

zero. An investor may wish to be more aggressive when applying this strategy to a stationary error ϵ_t than to an unbounded non-stationary asset price. Also the theoretical model implies that any non-zero value of ϵ_t is a deviation from market efficiency and should be fleeting. As such, using a multiple of less than $2\sigma_t$ may be more appropriate in this context. For example, if the unconditional distribution of ϵ_t is normal then using two standard deviations would imply generating a trading signal, and hence an opening position, only around 4% of the time. Given there are only a few hours to trade this is very conservative. Fatter tails in the unconditional distribution are expected but a smaller corridor may still be needed to generate a reasonable amount of trades.

We first try this strategy in the most studied asset with respect to Brexit which is the pound. When $\hat{\epsilon}_t$ is above the upper band we sell and buy when below the lower band. Both long and short positions are closed out when $\hat{\epsilon}_t$ has converged to zero, the expected equilibrium level. Note that selling the error is equivalent to selling the pound and buying the betting contract for Remain (or selling the contract for Leave) in the cointegrating ratio. At time t where $t > \text{midnight}$, the period from 11pm, 23rd June to t is used to calculate the sample standard deviation σ_t . $\hat{\epsilon}_t/|\hat{\Delta p}|$, Bollinger Bands and trades for both 2σ and 1.5σ strategies are shown in Figure 2.4.

The Bollinger strategy appears excellent. The 2σ signals generate seven trades, six of them winning, with a total gross profit of 52p for every £1 Betfair contract traded. This is impressive as the total move in the contract is only around 90p in the whole night. The strategy is able to capture a large amount of the entire overnight move in the betting contract without apparently taking political risk. The 1.5σ strategy generates 10 trades (9 winning) of lower average profit but a greater total gross profit of 94p for every £1 of Betfair contract. Trading costs for these markets are relatively small. They are well below the order of gross profits and so net profits will still be significant.¹⁵ Figure 2.5 shows the 1.5σ and 2σ strategies applied to all assets. Results are presented in table 2.8. Again the strategies appear excellent. Across all assets the $1.5\sigma_t$ strategy has 57 trades, 55 winning, with average profit of 7.8p for every contract traded on Betfair. The 2σ strategy has 36 winning out of 38 total trades with an average profit per trade

¹⁵In terms of transaction costs, selling the pound would cost about 2–3 hundredths of a cent at that time, whereas the Betfair cost is 3–5% levied on any bets that pay out. This would slightly change the ratio of the portfolio but not significantly affect profits or these conclusions.





Figure 2.5: Bollinger Band trading strategies for Brexit.

of 10p. Trading costs and slippage should not be above 1p per contract so these profits appear real.

Table 2.8: Bollinger Band gross trading profits for Brexit.

Symbol	1.5σ			2σ		
	$\frac{nWin}{nTrade}$	Profit ^a	$\frac{Profit}{Trade}$	$\frac{nWin}{nTrade}$	Profit ^a	$\frac{Profit}{Trade}$
ESU6	11/12	70.7p	5.9p	7/8	54.8p	6.8p
ZNU6	19/19	104.9p	5.5p	10/10	77.6p	7.8p
CLU6	8/8	74.4p	9.3p	6/6	69.5p	11.6p
GBPUSD	9/10	66.2p	6.6p	6/7	52.0p	7.4p
USDJPY	8/8	127.1p	15.9p	7/7	125.4p	17.9p
Total	55/57	443.3p	7.8p	36/38	379.3p	10.0p


^aProfits are shown for every £1 contract traded on Betfair

Unfortunately, being able to apply this strategy ex-ante is not at all realistic. Firstly the trader would need to be confident weak market efficiency would ultimately hold. Secondly, and more importantly, the calculation of $\hat{\epsilon}_t$ uses the cointegrating relationship estimated ex-post using data from the whole night. Ex-ante, a successful investor would have to correctly foresee the cointegrating relationship. This is equivalent to knowing the conditional expectations $GBP_L = E(GBP_T|BREXIT)$ and $GBP_H = E(GBP_T|NO BREXIT)$, ie where the pound settles given a vote to leave the EU (and the price it would have achieved in the counterfactual remain scenario).

We investigate an application of the same strategy when the cointegrating relationship is informed by forecasts of the conditional expectation from a market commentator. Prior to the referendum in late April the investment bank JPMorgan published a forecast of precisely where the pound would be priced given a vote to leave the EU. This was 1.32, Peters, 2016. If a trader on the night used this value for GBP_L and assumed that prices were efficient at 11pm then she would evaluate GBP_H at 1.522¹⁶ and $\hat{\Delta}p = GBP_L - GBP_H = -0.2024$. This is a larger predicted drop than estimated from the ex-post regression. Simply put, JP Morgan's prediction was too pessimistic. The pound did not fall (in the overnight

¹⁶This is because at 11pm GBPUSD was at 1.5007 and the betting contract implied a 10.7% chance of leaving the EU. If 100% chance of leaving the EU results in $GBP = 1.32$ then this implies that a zero chance implies $GBP_H = (1.5007 - 0.107 \times 1.32)/0.893 = 1.5224$. Note this is similar to \hat{p}_o calculated from the ex-post cointegrating relationship shown in Table 2.2.

session) as much as expected given the positive referendum result for Brexit. The cointegration error estimated in this way does not converge to zero as the pound does not fall as much as expected. The misspecified error is shown on the left hand side of figure 2.6 with Bollinger Band strategies. The error does drift downwards. Remarkably though, and against this author’s expectations, the trading strategy is still profitable, albeit less so. The error is biased downwards and so only “buys” are executed. It appears that JP Morgan’s conditional estimate for the pound, albeit a little pessimistic, was close enough to allow a profitable strategy. Although biased, there does seem to be some mean reversion of this estimated cointegration error. The conservative natures of the Bollinger Band signals avoid losing money. However, it is not the case that conditional predictions for the other financial assets are available and so this strategy does not appear readily applicable to the other financial assets.



Bollinger Bands strategies for misspecified cointegration error when using $E(GBP_T|BREXIT) = 1.32$ (left) and when using the first part of the night to estimate (right).

Figure 2.6: Possible ex-ante estimated cointegration errors and Bollinger Band strategies.

There does seem to be one way in which ex-ante profits may have been possible for the other symbols without the foresight of conditional predictions of the asset prices. This is by using the first few hours of the night to estimate the cointegrating relationship and trading in the latter part of the night. This actually does result in profits. An investor has the most chance of estimating the cointegrating relationship if there is sufficient support of the relationship within the training set. However, waiting for too much training data reduces the window of opportunity. We settle on using data up to the point where the betting market predicts a 50%

Table 2.9: Cointegration tests for the smaller training period.

Symbol	$p_{Z_t}(a = 0)^a$	$p_{Z_\alpha}(a = 0)^a$	$p(r = 0)^b$
ESU6	0.026**	0.007***	0.004***
ZNU6	0.059*	0.014**	0.001***
CLU6	0.344	0.356	0.423
GBPUSD	0.115	0.027**	0.057*
USDJPY	0.011**	0.003***	0.008***

^aPhillips–Ouliaris tests.^bJohansen max eigenvalue test.

chance of Brexit. This occurs at 2:02am. The strategy is shown on the right hand side of Figure 2.6. The error does appear to converge to zero in the trading period. There are five trades with all but the final trade profitable. We note that the final trade also lost money when trading using the ex-post estimated relationship. It appears that the long term relationship estimated from the first few hours of the night is stable and persists into the later hours of the night.

Table 2.10: Ex-ante Bollinger Band gross trading profits for Brexit.

Symbol	1.5σ			2σ		
	$\frac{nWin}{nTrade}$	Profit ^a	$\frac{Profit}{Trade}$	$\frac{nWin}{nTrade}$	Profit ^a	$\frac{Profit}{Trade}$
ESU6	0/1	-40.4p	-40.4p	0/1	-37.8p	-37.8p
ZNU6	10/10	58.0p	5.8p	7/7	46.8p	6.7p
CLU6	7/7	64.5p	9.2p	5/5	62.2p	12.4p
GBPUSD	9/10	56.8p	5.7p	4/5	29.3p	5.9p
USDJPY	6/6	106.4p	17.7p	5/5	100.1p	20.0p
Total	32/34	245.3p	7.2p	21/23	200.6p	8.7p

^aProfits are shown for every £1 contract traded on Betfair

Turning now to the other assets, we first check the cointegration tests for the shorter training period. The relevant p-values are shown in table 2.9. Naturally these tests have less power with less data. Nevertheless, four out of the five assets have results which reject the null of no cointegration in favour of cointegration at high significance levels. The exception is the crude oil future. We note that the significance was less for this asset when testing on the whole night and conclude

that the smaller data-set generates insufficient power to reject the null. The fact that there is strong evidence of cointegration for four of the assets in this smaller training period is pleasing.

The results of applying the ex-ante Bollinger Band strategy to all assets is illustrated in Figure 2.7. Gross profits are shown in Table 2.10. For four of the assets the strategy is impressive. There are less trades due to the smaller trading period, yet profits per trade are similar to those using the ex-post cointegration relationship. The exception is the S&P500 future. A single trade is executed which unfortunately loses around 40p. This occurs as the estimated error does not converge. The cointegration relationship does not appear stable for this asset. Nevertheless, the losses of this trade are outweighed by the profits in the other contracts.

These trading results are quite remarkable. Starting from three reasonable assumptions we wrote down a theory of asset prices which implied cointegration. We find that there is, for most assets, strong evidence of that theory generated in the first part of the night. Moreover the estimated cointegration relationship is so stable that trading profits can, apparently for most assets, be generated by taking positions against deviations from those long-term relationships. Whether or not they would be realistically crystallised by an intelligent investor is debatable. We have written down a theory after the fact. Evidence does appear to have quickly emerged on the night that the theory holds. However, to execute the trades described in this chapter would require a market participant to have confidence that the theory would indeed continue to apply for the remaining hours. It would take a brave soul to do so. It is not the case that there have been large numbers of referendums or political events where our approach has been shown to work¹⁷. Whether or not these profits are realistic in practice is a philosophical point. One thing is certain though, and that is that it is far easier for me to write this apparently successful study with hindsight after the fact than actually risk my money upfront on the night!

2.4.2 The 2016 United States presidential election

We next study the second great political shock of 2016, the US presidential election, held on Tuesday 8th November. The Republican ticket of Donald Trump and Mike

¹⁷Arguably, if there had been a history of success for this strategy, then the opportunity would have disappeared due to the actions of arbitrageurs.





Figure 2.7: Ex-Ante Bollinger Band trading strategies for Brexit.

Pence, against expectations, beat the Democratic ticket of Hilary Clinton and Tim Kaine.

Under the Electoral College system the winner needs at least 270 of the 538 electors. There are 51 voting areas, 50 states plus the special federal district of Washington D.C., that each award electoral votes. Electors, for the most part, vote for the winners of the popular vote within their respective area¹⁸. The democratic candidate led in the vast majority of nationwide and swing-state polls. However the margin decreased as the election was approached. On election day Donald Trump outperformed his polls, winning all of the key battleground states of Florida, North Carolina, Ohio and Iowa. Additionally, and against all expectations, he took the three formerly Democrat “rust-belt” states of Pennsylvania, Michigan and Wisconsin. The Republican ticket’s votes were exceptionally well distributed. Donald Trump won 30 states with 306 electoral votes whereas Hillary Clinton won 20 states with 232 votes.¹⁹ This is despite Trump garnering 2.87 million less votes than Hilary Clinton.

Similar to the Brexit referendum there were multiple vote counts (51 versus 382) from different regions occurring throughout the night after voting ended. The situation is complicated further by polls closing at different times in different states. However, evolving vote counts were published in real time on all the major news networks as well as the Internet. As with the Brexit referendum the only information affecting the market that night was the vote counts and results. Similar to Brexit there were also various betting markets open and trading. Bets that paid out in the event of either a Trump or a Clinton win were widely traded. Figure 2.8 shows the price series for the Betfair contract that pays out £1 in the event of a Republican win in Greenwich Mean Time (GMT).²⁰ Between midnight and 1:00am the risk neutral odds of a Trump win varied between 10% and 20%. However, by 1:30am GMT (20:30 EST) the count of the crucial swing state of Florida was almost completed and showed a lead for Trump of 0.7% versus an expectation that Clinton would win by 0.6%. From this point on the odds for a Trump win improved as various other counts showed Trump consistently out-performing his polling. By a little after 4am the betting markets implied a Trump win with 95% probability

¹⁸Exceptions include Maine and Nebraska where electors are allocated based on a combination of the plurality of votes as well as the popular winner in each of their congressional districts. There are also typically a handful of “faithless electors” in each election who chose to vote against the candidate for whom they had pledged to vote.

¹⁹There were seven faithless electors in total; five defections from Clinton and two from Trump.

²⁰GMT is 5 hours ahead of Eastern Standard Time (EST) and is the time in London on the date of the election. All times in this section are quoted in GMT

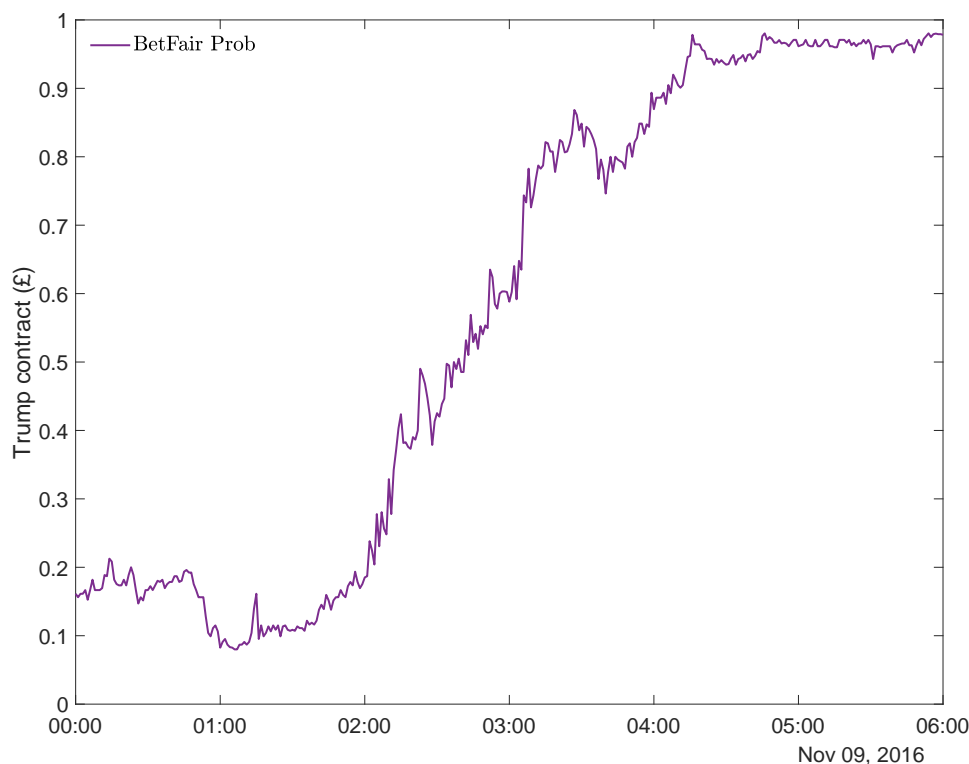


Figure 2.8: Betfair contract for a Trump win in the 2016 presidential election.

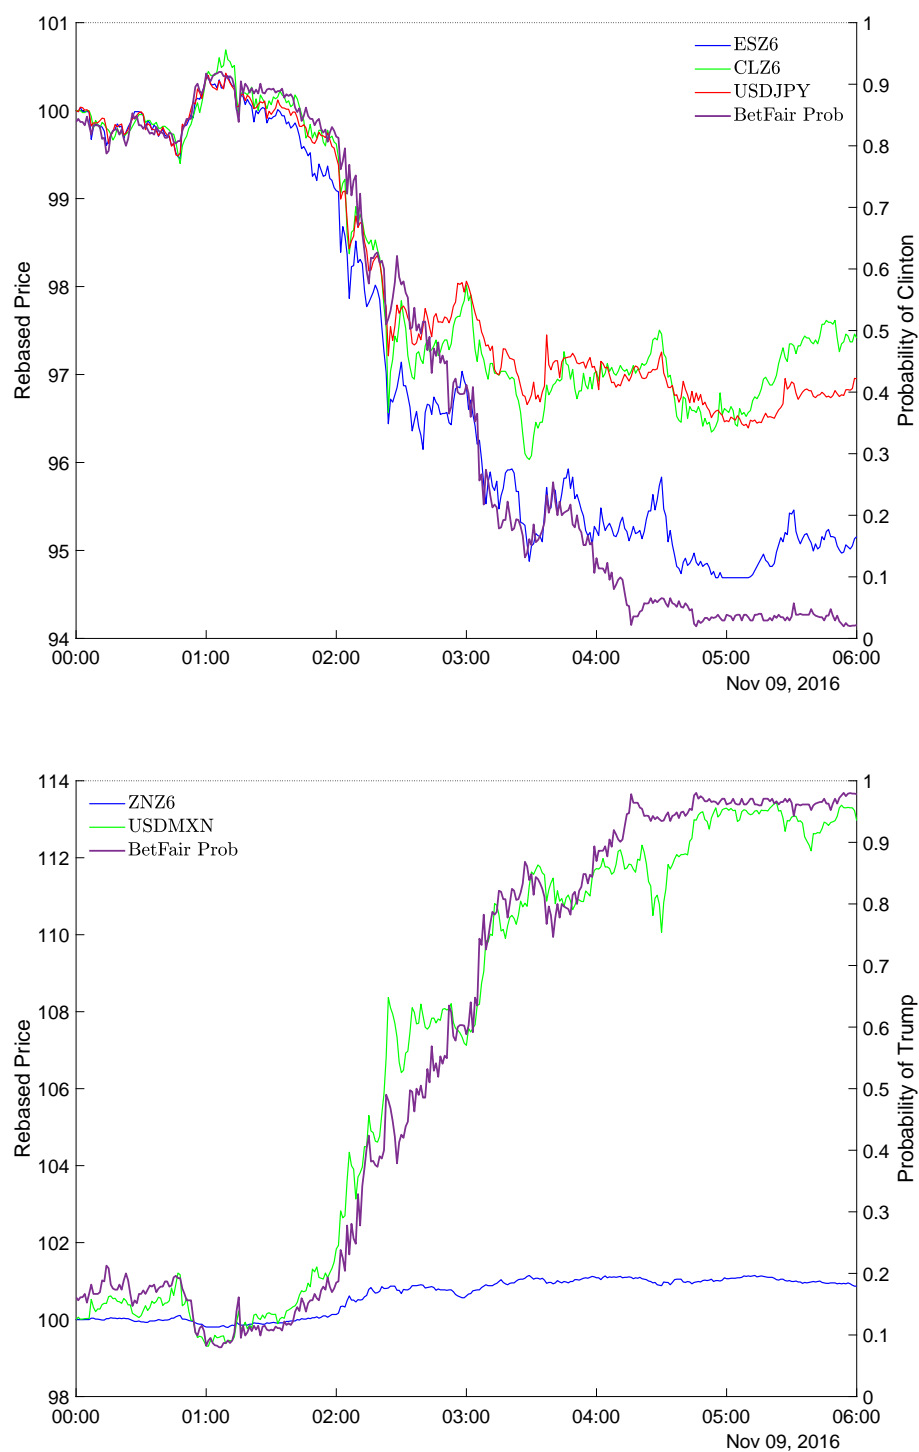
which slowly increased to 98% by 6am. At 7:50am Donald Trump made his victory speech.

2.4.2.1 Financial assets

We use a similar basket of financial assets as we did with the Brexit referendum. The exception being we swap out the UK specific GBPUSD cross and include the USDMXN exchange rate. Trump had proposed that if he won he would renegotiate or exit various trade agreements including those with Mexico. Given the dependence of Mexico's economy on trade and exports to the US,²¹ a Trump win was seen as extremely negative for that country's economy.

We consider the period midnight to 6:00am BST on 9th November and sample prices every minute. During these hours all contracts are open, the betting market almost converges and the only information released to the market is that contained in the vote results. We do not consider beyond 6am as this is the start of the trading

²¹Over 80% of Mexican exports in 2015 were to the US.



Note: Assets that depreciated are shown versus the “Clinton” contract price whereas assets that appreciated are shown versus the contract price for “Trump”.

Figure 2.9: Rebased financial asset prices versus US presidential election contracts.

Table 2.11: Financial assets and changes for the 2016 presidential election.

Symbol	Description	$p_{t=1}$	$p_{t=T}$	Δ
ESZ6	E-mini S&P500 Future	2142.25	2038.500	-4.8%
ZNZ6	10-Year T-Note Future	129.484	130.6094	0.9%
CLZ6	Crude Oil Future	44.880	43.72	-2.6%
USDMXN	US Dollar Mexican Peso Cross	18.310	20.684	13.0%
USDJPY	US Dollar Japanese Yen Cross	104.990	101.789	-3.1%

day in London. Other economic news beyond the election may be released which would invalidate our model and, either way, the result had become apparent by then. Table 2.11 lists the five assets with a description, their starting and ending prices and the percentage change. The election of Donald Trump was a shock. Not only had he pledged to renegotiate various trade deals, reducing world trade and hence the outlook for the economy, he was widely seen as unpredictable and inconsistent. The US dollar depreciated against the safe haven Japanese yen, and the oil and stock market futures depreciated too. The treasury future appreciated as would be expected in a time of increasing risks to the US economy and the US dollar appreciated a large 13% against the Mexican peso. This is as the market re-priced the very significant risks to the Mexican economy following the Trump win.

Figure 2.9 plots depreciating and appreciating assets versus the Betfair contracts that pay out £1 for a Clinton and Trump win respectively. By and large the financial markets do seem to be moving together with the betting contracts. The relationship does not look quite as established as that for the Brexit referendum with some reversal of the large falls past 5am for the S&P500 and oil futures as well as the USDJPY exchange rate. We turn to the statistical specifications and tests in the next sub-section to make robust conclusions.

2.4.2.2 Evidence for cointegration

We now consider the evidence for cointegration between the financial assets and the betting markets. Table 2.12 shows the results of the Phillips–Ouliaris regression tests for cointegration and Table 2.13 shows the results of the bivariate Johansen max eigenvalue tests. Neither the Z_α or Z_t tests reject the null of no cointegration for assets beyond the Mexican Peso whereas the Johansen test does reject the cointegration rank of 0 for all assets at either the 95% or 99% level. There

Table 2.12: Results of the Phillips–Ouliaris Z_t and Z_α tests for cointegration.

Symbol	\hat{p}_0	$\hat{\Delta p}$	$p_{Z_t}(a = 0)$	$p_{Z_\alpha}(a = 0)$
ESU6	2152.59	-126.27	0.109	0.094*
ZNU6	129.314	1.673	0.709	0.671
CLU6	44.99	-1.67	0.497	0.467
USDMXN	18.031	2.751	0.015**	0.014**
USDJPY	105.298	-4.041	0.308	0.279

The first stage regression is $p_t = \hat{p}_0 + \hat{\Delta p}PB_t + \hat{\epsilon}_t$

The residual regression is $\hat{\epsilon}_t = a\epsilon_{t-1} + \eta_t$

Table 2.13: Results of the bivariate Johansen max eigenvalue test.

Symbol	\hat{p}_0	$\hat{\Delta p}$	$k - 1$	$p(r = 0)$
ESU6	2159.49	-126.6967	6	0.006***
ZNU6	129.174	1.5993	4	0.003***
CLU6	45.21	-1.64	7	0.028**
USDMXN	17.908	2.764	5	0.024**
USDJPY	105.600	-3.981	4	0.009***

The regression is $\Delta p_t = \alpha(\beta' p_{t-1} - c_0) + \sum_{i=1}^{k-1} \Gamma_i \Delta p_{t-i} + u_t$
 $r = \text{rank}(\beta)$, when $r = 1$, $\beta' = (1, -\Delta p)$, $c_0 = -p_0$

Table 2.14: Results of the multivariate Johansen max eigenvalue test.

r	h	stat	cValue	eigVal	pValue
0	1	58.75	40.96	0.152	0.001***
1	0	27.54	34.81	0.074	0.311
2	0	18.81	28.59	0.051	0.538
3	0	12.11	22.30	0.033	0.659
4	0	6.46	15.89	0.018	0.742
5	0	2.69	9.16	0.008	0.679

The regression is $\Delta p_t = \alpha(\beta' p_{t-1} - c_0) + \sum_{i=1}^{k-1} \Gamma_i \Delta p_{t-i} + u_t$
 $r = \text{rank}(\beta)$

Table 2.15: Results of the trivariate Johansen constraint test for risk aversion.

Symbol	$\hat{\lambda}$	$k - 1$	$p(\lambda = 0)$
ESU6	-15.951	4	0.904
ZNU6	13.378	3	0.132
CLU6	-5.74	5	0.001***
USDMXN	-4.988	4	0.575
USDJPY	-7.481	4	0.010***

The regression is $\Delta \mathbf{p}_t = \boldsymbol{\alpha}(\boldsymbol{\beta}' \mathbf{p}_{t-1} - c_0) + \sum_{i=1}^{k-1} \boldsymbol{\Gamma}_i \Delta \mathbf{p}_{t-i} + \mathbf{u}_t$

$$\mathbf{p}_t = (p_t, PB_t, R_t)'$$

$$r = \text{rank}(\boldsymbol{\beta}) = 1, \boldsymbol{\beta}' = (1, -\Delta p, -\lambda), c_0 = -p_0$$

is weaker evidence for cointegration than in the Brexit case. This is because the Phillips–Ouliaris tests are robust to serial correlation and heteroskedasticity. In applying the Johansen test we do remove serial correlation from the residual through the inclusion of sufficient lags but we cannot remove conditional heteroskedasticity. Thus the Johansen test is not as robust as the Phillips–Ouliaris tests. We conclude there is strong evidence for our model in the case of USDMXN and weak evidence for the remaining assets. Turning now to the order of cointegration, Table 2.14 shows the results of the max eigenvalue tests of the 6-dimensional system which includes all financial assets and a betting contract. The evidence here for our model is weak. The tests do reject a cointegration rank of zero in favour of a single cointegration vector but higher orders of cointegration are not rejected. It may be that the tests do not have sufficient power, but compared to the Brexit case where ranks of below 3 were rejected, the data do not appear to fit our model (which implies rank 5) as well for this election. The Nyblom and Harvey test for a single common trend statistic is 0.008. This is well below the level of rejection at the 90% level so the data also do not reject our model's prediction of cointegration rank 5. We conclude that there is strong evidence that there is at least a cointegrating relationship for the Mexican Peso but that the evidence for cointegration of the other assets and hence higher orders of cointegration is mixed.

2.4.2.3 Risk Aversion

Our model that assumes risk neutrality implies a linear cointegrating relationship. There is weak evidence for the risk neutral model for four of the five financial assets considered. Could the variant of our model that includes risk aversion provide a

Table 2.16: Results of the Phillips–Ouliaris Z_t and Z_α tests for cointegration. R_t included.

Symbol	\hat{p}_0	$\hat{\Delta}p$	$\hat{\lambda}$	$p_{Z_t}(a=0)$	$p_{Z_\alpha}(a=0)$
ESU6	2170.038	-135.277	-101.664	0.001***	0.001***
ZNU6	128.873	1.900	2.570	0.036**	0.026**
CLU6	45.50	-1.94	-3.00	0.010***	0.007***
USDMXN	17.817	2.861	1.252	0.003***	0.003***
USDJPY	106.032	-4.419	-4.274	0.007***	0.007***
ZNU6 ^a	128.980	1.863	1.863	0.065*	0.039**
CLU6 ^a	45.35	-1.88	-1.88	0.036**	0.021**

The first stage regression is $p_t = \hat{p}_0 + \hat{\Delta}pPB_t + \hat{\lambda}R_t + \hat{\epsilon}_t$

^aThe first stage regression is $p_t = \hat{p}_0 + \hat{\Delta}p[PB_t + R_t] + \hat{\epsilon}_t$ so $\hat{\lambda} = \hat{\Delta}p$

The residual regression is $\hat{\epsilon}_t = a\hat{\epsilon}_{t-1} + \eta_t$

better fit to the data for the US presidential election? Again we apply the Johansen constraint test to the trivariate systems $x_t = (p_t, PB_t, R_t)'$ for each financial asset p_t to test this idea. Results are shown in table 2.15. Unlike in the Brexit case, there does appear to be evidence of non-trivial risk preferences and non-linear relationships. The coefficient relating to risk, λ , is significant and negative (as expected) for the yen and the oil future. This is interesting. It is not significant for the S&P500 future, the treasury future, or the Peso (although the linear model seems adequate to describe the USDMXN behaviour). To further explore the possibility of non-linear cointegration we re-run the Phillips–Ouliaris tests with the risk factor $PB_t \cdot (1 - PB_t)$ included as an explanatory variable in the regression. Results are shown in Table 2.16. The results are striking. The null of no cointegration is now rejected at the 99% level for four assets and at the 95% level for the treasury future. We also note that for the treasury and oil futures $|\hat{\lambda}| > |\hat{\Delta}p|$ which violate our model. A coefficient of risk with larger magnitude than $\hat{\Delta}p$ implies a non-monotonic utility function which decreases at an end point. This is likely due to interpolation to the slight pull back in the asset prices observed between 5am and 6am. We re-run the regression using the single non-linear monotonic explanatory variable $PB_t + R_t$. This also results in the null of no cointegration being rejected for both assets.

The signs of the estimated risk parameters are encouraging. For all “risk” assets that depreciated on the shock Trump win the parameters are negative, implying risk aversion. For the asset that appreciated, ZNU6, the sign is positive, indicating risk loving behaviour. This is discussed further below.

Table 2.17: Results of the multivariate Johansen max eigenvalue test. R_t included.

r	h	stat	cValue	eigVal	pValue
0	1	59.78	47.08	0.155	0.002***
1	0	32.91	40.96	0.088	0.330
2	0	27.75	34.81	0.075	0.297
3	0	23.79	28.59	0.065	0.182
4	0	10.76	22.30	0.030	0.773
5	0	8.47	15.89	0.024	0.537
6	0	3.17	9.16	0.009	0.608

The regression is $\Delta \mathbf{p}_t = \boldsymbol{\alpha}(\boldsymbol{\beta}' \mathbf{p}_{t-1} - \mathbf{c}_0) + \sum_{i=1}^{k-1} \boldsymbol{\Gamma}_i \Delta \mathbf{p}_{t-i} + \mathbf{u}_t$
 $r = \text{rank}(\boldsymbol{\beta})$

According to the Johansen test the coefficient is significant for the oil future and the dollar denominated in yen. The sign is negative implying risk aversion. To be clear, the Japanese yen is considered a “risk-off” asset. The US dollar denominated in yen is “risky” and depreciated 3.1% overnight on the shock result that Donald Trump was elected. An explanation for this behaviour is that speculators will borrow in the lower yielding yen to finance investment in assets of higher yielding currencies when risk appetite increases.²² The coefficient in the Phillips–Ouliaris regression for USDMXN is positive. This is consistent with risk aversion, as this is the price of the dollar denominated in the risky asset, the Peso. A positive value of λ for USDMXN is equivalent to a negative value in the regression of MXNUSD against PB_t and $PB_t \cdot (1 - PB_t)$ ²³. The coefficient of ESU6 is also negative whereas that for ZNU6 is positive. The positive (albeit not necessarily significant) coefficient for ZNU6 indicates risk loving preferences, which on first impressions may be surprising. However, our model incorporated risk preferences by assuming that the participants in the betting market were risk neutral and those for financial assets were risk averse. This led to a discount on asset prices proportional to the uncertainty in the outcome of the political event. Uncertainty is highest when PB_t is around the middle of its range $[0, 1]$ and furthest from the endpoints. For ZNU6 there instead appears to be a premium on the price when

²²Japanese interest rates and asset yields have historically, and at the time of the vote, been lower than elsewhere.

²³Mathematically: negative λ is equivalent to a convex cointegrating relationship and positive λ implies concavity. If USDMXN is concave, then $\text{MXNUSD} = \text{USDMXN}^{-1}$ is convex. Economically: if USD is “safe” relative to the MXN, then MXN is “risky” relative to the USD.

uncertainty is highest. This is not unexpected. US treasuries are considered the ultimate safe haven asset, and are typically bought when there is a flight to quality and risk appetite decreases. Investors sell risky assets such as commodities (CLU6), stocks (ESU6) or emerging market currencies (MXNUSD), recover any funding in lower yielding currencies (USDJPY) and invest the proceeds in treasuries. This behaviour reveals itself on the night not only by the ZNU6 appreciating on the shock result, but as a risk loving preference and a positive value of λ in the cointegrating relationship.

Finally we check the implications for the order of cointegration for the 7-dimensional system which includes the five asset prices, the betting contract, and the risk factor R_t . When risk is included in the system our model still implies a cointegration rank of five but there are now two common stochastic trends, relating to the probability of the political event, and the risk factor. Results when the Johansen test is re-run with risk included are shown in Table 2.17. Disappointingly the inclusion of the non-linear risk factor does not provide any more evidence of a higher cointegration rank. Again, only the zero rank is rejected in favour of higher ranks. The situation of the NH test is similar. The test statistic for two common trends (PB_t and risk) of 0.0029 again are well below the 90% level of rejection in favour of higher numbers of trends and a lower cointegration rank than 5. The strong evidence in favour of a non-linear cointegrating relationship between the financial and betting markets from the Phillips–Ouliaris tests is not repeated in the rank tests.

2.4.2.4 Deviations from long term relationships

As in the Brexit case, the presence of significant autoregressive terms in both the cointegrating regressions and the VECM suggests deviations from efficiency on a short term scale. There is also evidence for the 2016 presidential election that risk neutrality does not hold. The “long-term” relationship, which lasts a few hours, is now non-linear. The full multivariate system is

$$\mathbf{p}_t = (\text{ESZ6}, \text{ZNU6}, \text{CLU6}, \text{USDUSD}, \text{USDJPY}, PB_t, R_t)'$$

where R_t is the non-linear political risk measure $PB_t \cdot (1 - PB_t)$. The fitted VECM, as in the case with Brexit, has a single significant autoregressive short run matrix Γ_1 . Table 2.18 shows the significant values of this matrix. Testing for short run

Table 2.18: Fitted short run parameters in the error correction model for the Trump election and tests on the short run dynamics.

Short run regression results

$$\Gamma_1 = \begin{pmatrix} -0.196^{**} & -9.28^{**} & 4.92^* & - & - & -36.4^{***} & - \\ - & - & -0.0955^* & -0.0956^* & - & 0.502^{***} & - \\ - & - & - & -0.147^* & - & -0.426^{***} & - \\ - & - & - & - & - & 0.567^{***} & - \\ - & -0.240^* & 0.200^* & - & -0.274^{***} & -1.10^{***} & -0.893^* \\ - & - & -0.0554^{**} & -0.0500^{**} & - & -0.329^{***} & - \\ - & -0.0555^{**} & - & - & - & - & -0.277^{***} \end{pmatrix}^1$$

Based on $\Delta p_t = \alpha(\beta' p_{t-1} - c_0) + \Gamma_1 \Delta p_{t-1} + u_t$

$p_t = (\text{ESZ6}, \text{ZNZ6}, \text{CLZ6}, \text{USDMXN}, \text{USDJPY}, PB_t, R_t)'$

$R_t = PB_t \cdot (1 - PB_t)$

¹ Statistically insignificant coefficients omitted

Null hypothesis	LR-Test	
	Chi-square	Prob
$(\text{ESZ6}, \text{ZNZ6}, \text{CLZ6}, \text{USDMXN}, \text{USDJPY})'$ does not cause (PB_t, R_t)	15.7	0.108
(PB_t, R_t) does not cause $(\text{ESZ6}, \text{ZNZ6}, \text{CLZ6}, \text{USDMXN}, \text{USDJPY})'$	32.8^{***}	< 0.001

causality in this system requires considering both the risk measure R_t in addition to the betting contract price P_t . To test causation from the betting markets to financial markets, consideration of the $(n - 2) \times 2$ sub-matrix of Γ_1 , which is the first $n - 2$ values of the last two columns, is required. Testing in the opposite direction requires joint testing of the $2 \times (n - 2)$ sub-matrix which is the first $n - 2$ values of the last two rows. Table 2.18 shows results of these two tests. As was the case with the previous event, there is significant causation from the betting markets to the financial markets in the short run (at the 1% level) but no significant causation in the other direction.

The cointegrating error for this election is $\hat{\epsilon}_t = p_t - \hat{p}_0 - PB_t \times \hat{\Delta p} - \hat{\lambda} \times R_t$. Unlike in the linear case with Brexit, this error cannot be recreated simply by trading a portfolio of the betting contract with the financial asset. The non-linearity would require constant adjustment, or delta hedging. This would be difficult and expensive to achieve in practice. However, as the relationship is monotonic, if the error can be predicted, this does imply that the two assets separately can be predicted. Figure 2.10 plots $\hat{\epsilon}_t / |\hat{\Delta p}|$, the cointegration error normalised to units of the Betting contract, with the sample Auto-Correlation Function and robust Bartlett




Figure 2.10: Non-linear cointegration errors with autocorrelation functions.

intervals. As in the case with Brexit, this shows significant and positive autocorrelations out to around 20–30 minutes. The presence of long term cointegration of the markets implies market efficiency on a longer term time timescale. The fact that deviations appear predictable suggests an inefficiency on a shorter timescale. We fit autoregressive models to $\hat{\epsilon}_t$. We find, for all assets, that models with two lags are significantly different to a model with a single lag but that further lags are not significant. Table 2.19 shows the results of the regressions. This again provides evidence that EWMH was violated on a short time scale on the night of the 2016 US presidential election.

2.4.2.5 Profit opportunities

We now turn to the possibility of profiting from the apparently predictable deviations from the long term cointegrating relationship. In practice it is not possible to recreate the return of the non-linear error so we will focus on the linear error.


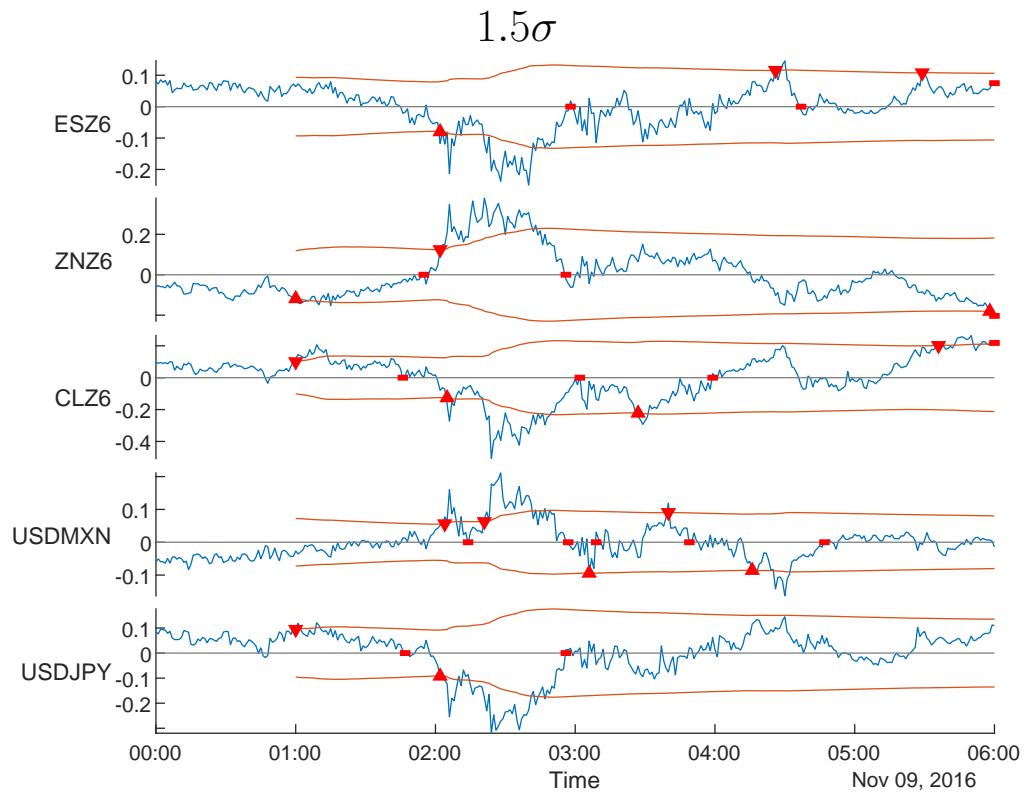


Figure 2.11: Ex-post Bollinger Band trading strategies for the 2016 presidential election.

Table 2.19: Estimated autoregressive models for cointegration errors.

lag	ESZ6	ZNZ6	CLZ6	USDMXN	USDJPY
$\hat{\epsilon}_{t-1}$	0.559*** (0.038)	0.783*** (0.034)	0.771*** (0.043)	0.629*** (0.040)	0.658*** (0.046)
$\hat{\epsilon}_{t-2}$	0.398*** (0.036)	0.196*** (0.036)	0.204*** (0.044)	0.302*** (0.041)	0.308*** (0.044)

The regression is $\hat{\epsilon}_t = \psi_1 \hat{\epsilon}_{t-1} + \dots + \psi_k \hat{\epsilon}_{k-1} + \eta_t$.

$\hat{\epsilon}_t$ is estimated from the first stage regression.

Again, this is reproduced by trading the financial asset against the political bet in a ratio equal to the linear cointegrating ratio. This is estimated from the ex-post regression and the parameters are shown in Table 2.12. We apply the 1.5σ and 2σ sigma trading strategies. Results are shown in Table 2.20 and the normalised error $\hat{\epsilon}_t/|\hat{\Delta}p|$ is shown with bands and trades in Figure 2.11. The strategies are profitable although there were far fewer trades than in the Brexit case. This may be because for four of the five assets linear cointegration does hold. The 1.5σ strategy generates 17 trades, 15 profitable, compared with 57 trades for Brexit. The 2σ strategy only generates eight trades, all profitable compared with 38 for Brexit. The errors appear to mean revert with far less frequency than for Brexit. However, each trade is more profitable, indicating a greater volatility on the night of the Trump election. As discussed earlier, recreating these profits in practice is not at all realistic. It would require an investor to correctly evaluate the conditional expectations of each asset price given the two possible outcomes of the election. The results of the ex-post strategies are at best thought experiments, and rely on the conditional expectations that have been revealed after the event. In fact applying the Bollinger strategy in this way ex-post to any two unrelated random walks is always likely to generate apparent profits. There will always be mean reversion to an apparent relationship that is estimated from the (spurious) regression. Given there is not even evidence of linear cointegration for the financial assets (beyond the Peso) the results below are likely pure fiction.

The USDMXN cross showed evidence of linear cointegration. Could profits be generated from this asset using an ex-ante estimated forecast for the conditional expectations of the asset prices? To answer this question we exploit a Reuters poll of several economists that was published on 1st November 2016, a week before the election, Cascione (2016). The prospect of a Republican win was expected to be very grave for the Mexican economy. The Reuters poll asked a question as to where


Table 2.20: Ex-Post Bollinger Band gross trading profits for the 2016 presidential election.

Symbol	1.5σ			2σ		
	$\frac{nWin}{nTrade}$	Profit ^a	$\frac{Profit}{Trade}$	$\frac{nWin}{nTrade}$	Profit ^a	$\frac{Profit}{Trade}$
ESU6	3/3	27.6p	9.2p	1/1	11.7p	11.7p
ZNU6	2/3	23.8p	7.9p	1/1	20.7p	20.7p
CLU6	3/4	51.8p	12.0p	2/2	41.7p	20.8p
GBPUSD	5/5	45.4p	9.1p	3/3	31.1p	10.4p
USDJPY	2/2	20.2p	10.1p	1/1	25.4p	25.4p
	15/17	168.8p	9.9p	8/8	130.6p	16.3p

^aProfits are shown for every £1 contract traded on Betfair

the Mexican Peso would settle in the event that Donald Trump were elected. The median forecast for USDMXN was 21.50. In the event this was a little pessimistic. The Peso bottomed out at 5:23am (GMT) with a USDMXN price of 20.77. It slightly appreciated to 20.68 by 6am. Assuming efficiency at midnight and using the Reuters median for $p_1 = 21.50$ implies $p_0 = 17.70$. Thus $\Delta p = 3.80$ and $\hat{\epsilon}_t = GBP_t - 17.70 - 3.80 \times PB_t$. This ex-ante estimated error from the Reuters forecast is plotted with Bollinger Band strategies and trades on the left hand side of Figure 2.12. The overly pessimistic forecast makes the error diverge from zero, yielding a large loss making trade in the second half of the night. For the 2σ strategy there are two profitable trades made in the first half of the night. These occur before the error diverges. These do more than offset the loss of the final trade but the total gross profits of 3.2p are likely to be negligible after trading costs. The profits of around 1p per trade compare very poorly with that of the (unrealistic) ex-post strategy (16.3p).

Next we study whether the cointegrating relationship can be estimated for the USDMXN in the first few hours and successfully exploited later in the night. The betting markets implied a 50% probability of a Trump win by 2:39am. Using the period midnight to 2:39am to estimate a linear cointegrating relationship estimates $p_0 = 17.81$ and $p_1 = 21.69$. Again the conditional expectation of USDMXN given a Trump win is overly pessimistic. It results in an estimated error that diverges from zero. This yields a terrible trade in the second half of the night as demonstrated

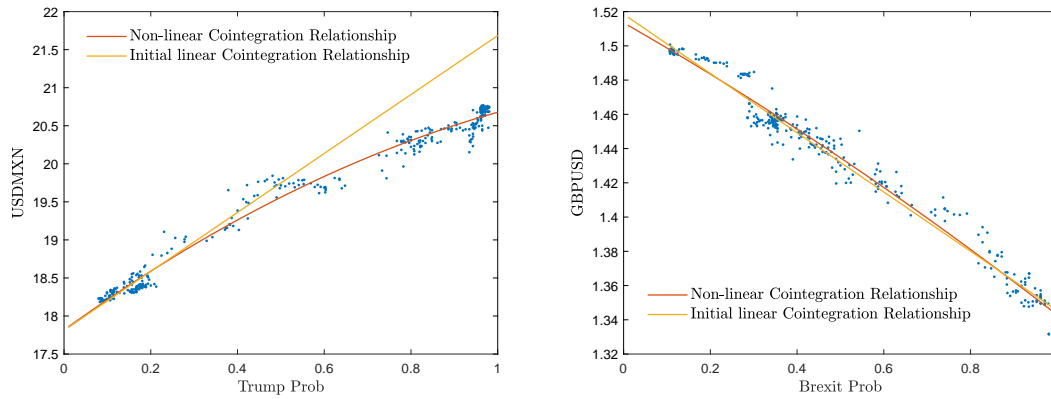


Bollinger Bands strategies for misspecified cointegration error when using $E(USD/MXN|TRUMP) = 21.50$ (left) and when using the first part of the night to estimate (right).

Figure 2.12: Possible ex-ante estimated cointegration errors and Bollinger Band strategies for USD/MXN.

by the right hand side of Figure 2.12 where the misspecified error and Bollinger bands are shown.

The results obtained for GBP/USD for Brexit, where the linear relationship estimated from the first period was stable, yielded a successful strategy. The answer as to why the performance is so different in this instance is the non-linearity and apparent deviation from risk neutrality. Figure 2.13 illustrates the difference. The GBP/USD and USD/MXN are plotted against the relevant betting contract along with both the linear relationship from the first part of the night as well as the ex-post non-linear cointegrated relationship. The non-linear plot clearly diverges from the linear plot for USD/MXN whereas there is almost no difference in the two lines for GBP/USD. There is no hope of estimating the long term non-linear cointegrating relationship for USD/MXN from the first part of the night and making money from any convergence. Figure 2.14 shows similar plots for the remaining assets for the 2016 presidential election as well as Brexit. The situation is replicated for these assets. There is a clear difference in behaviour for the Trump election. In the Brexit case the non-linear relationships are virtually identical to the linear one. This is consistent with the earlier Johansen constraint tests where there was no evidence of significant risk parameters λ . The only Brexit contract where the market does not converge to the linear cointegrated relationship is the ESU6, and the divergence is far smaller than for the presidential election. Note that this was the contract that lost money when applying the Bollinger ex-ante strategy errors. However, this does



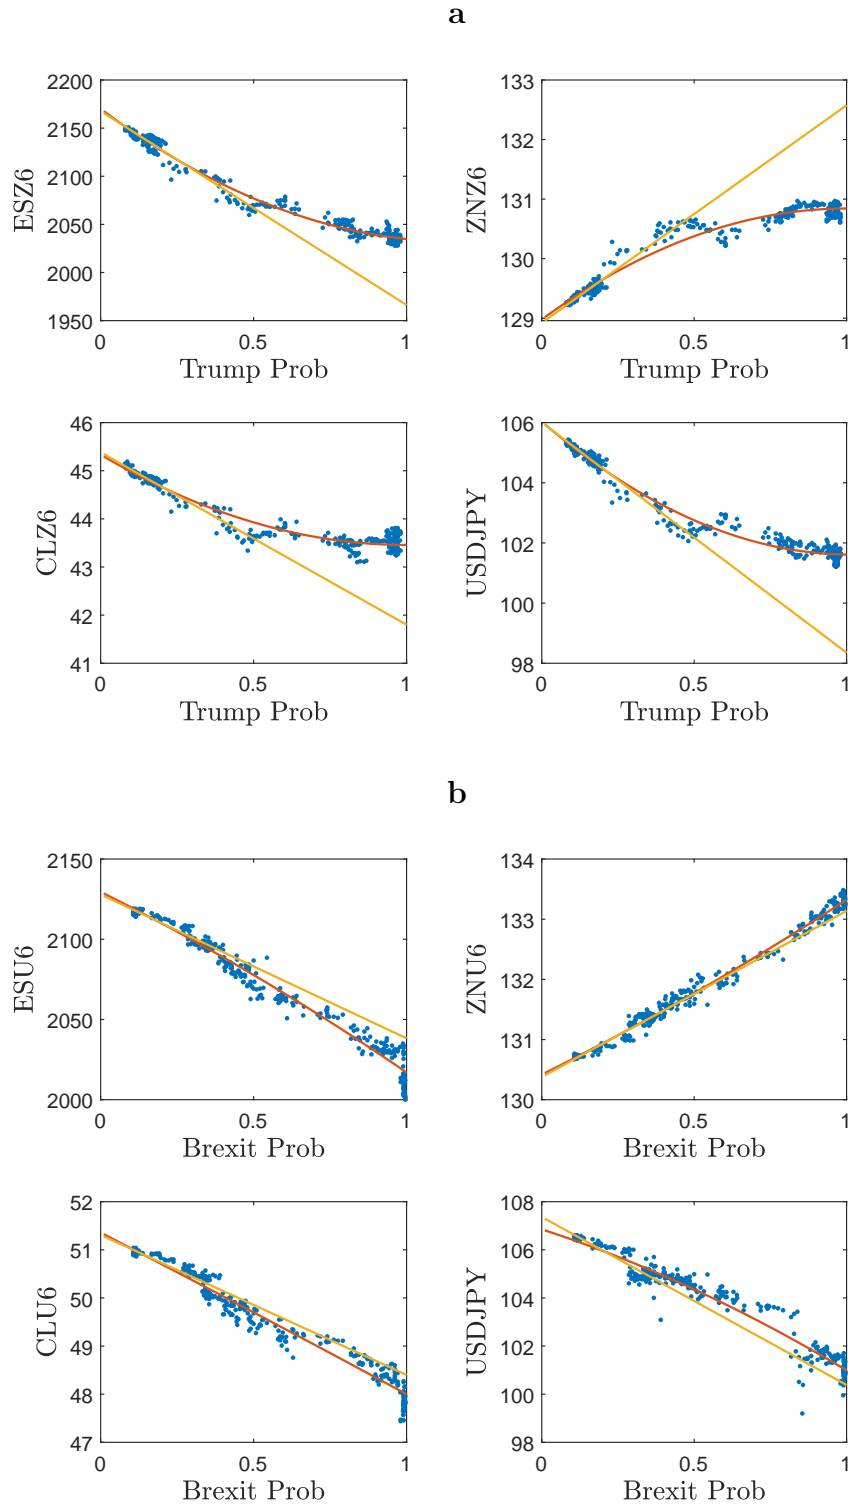
The non-linear relationship is calculated ex-post on all data whereas the initial linear relationship uses data from an initial period when event probabilities were under 50%.

Figure 2.13: Estimated cointegration relationships for USDMXN (presidential election) and GBPUSD (Brexit referendum).

not appear to be due to the presence of risk aversion but more to a slight misestimation in the linear cointegration ratio. The application of the ex-ante Bollinger strategy on the night of the Trump election would be ruinous. Moreover the data from the first period appear to indicate a linear relationship. Table 2.21 shows the pValues for the null hypothesis of no (linear) cointegration for the Phillips–Ouliaris Z_t and Z_α tests and the Johansen max eigenvalue test. Cointegration is indicated. This could lead any brave investor that profited from this strategy on the night of the Brexit referendum to trade a linear relationship that does not exist. The apparent divergence from risk neutrality after the presidential election leads to far different behaviour than that observed after Brexit, where risk neutrality appeared to hold. We conclude that there appear to be no realistic opportunities to profit from our model for the 2016 US presidential election.

2.4.3 The 2014 Scottish independence referendum

The 2014 Scottish independence referendum was held on 18th September 2014. It concerned whether Scotland should remain in the UK. The question asked “Should Scotland become an independent country?” and it required a simple majority to pass. All EU and Commonwealth citizens residing in Scotland were eligible to vote. This was the first time that an election was open to 16 and 17 year olds. Turnout, at 84.6%, was the largest for any UK election since the general election of 1910. The “No” side won with 55.3% of the vote against 44.7% for “Yes”.



Note: The non-linear relationship (orange line) is calculated ex-post on all data whereas the initial linear relationship (yellow line) uses only data from the first few hours when event probabilities were under 50%.

Figure 2.14: Estimated cointegration relationships for (a) the 2016 presidential election and (b) the Brexit referendum.

Table 2.21: Cointegration tests for the smaller training period for the presidential election.

Symbol	$p_{Z_t}(a = 0)^a$	$p_{Z_\alpha}(a = 0)^a$	$p(r = 0)^b$
ESU6	0.001***	0.001***	0.187
ZNU6	0.001***	0.001***	0.014**
CLU6	0.001***	0.001***	0.032**
USDMXN	0.002***	0.003***	0.028**
USDJPY	0.001***	0.001***	0.040**

^aPhillips–Ouliaris tests.

^bJohansen max eigenvalue test.

Voting took place from 7am to 10pm BST. Voters still queuing at the close of polls were allowed to vote. Each of the 32 local authority areas announced their results at separate times overnight. This makes this event suitable for our model as the information was drip fed to the public throughout the overnight hours.

During July and August 2014, opinion polls showed a consistent lead for “No” with the average difference being 6%, Curtice (2014). However, polls tightened at the start of September and on the 6th of that month YouGov published a poll that showed the Yes side ahead with a small lead of 2%, Dahlgreen (2014). This was the first time a poll showed the Yes side ahead. It was also the first time many commentators took the possibility of a Yes vote, and the resulting disruption to the UK’s economy, seriously. In response sterling lost around 1% against the euro and dollar and companies with links to Scotland sold off sharply, Wearden (2015). However the 6th September poll was the only poll that showed Yes ahead. The polls subsequently reverted to a Yes trailing No. The final poll released was conducted by YouGov on the day of voting. The results were released just after the close of voting and predicted an 8 point lead for No.

Figure 2.15 shows the Betfair contract that pays out £1 in the event of a No vote, along with the GBPUSD price. The betting market had previously, following the release of the 6th September poll, implied a probability of Scotland leaving the UK of as high as 35%. However, on the night of the vote the probability was much lower. From 11pm to midnight Betfair implied around a 90–92% chance of remaining. This contract started rising after 1am. The first local authority to announce their result was Clackmannanshire at 1:28am. This showed the vote for No nearly 8% ahead of Yes and in line with the poll on the day released by YouGov. The SNP, whose headline policy was independence, had achieved their highest vote

share in Clackmannanshire in the 2012 council elections and so this result was seen as very negative for the Yes campaign. The No contract on Betfair continued its ascent and shortly after this result priced a 98–99% probability of No winning the referendum. The currency was sensitive to the possibility, albeit small, of a yes vote. It rallied around a cent to the dollar as this possibility was effectively ruled out. From the point of this first result onward the result did not seem in doubt. The betting contract price remained above 98% for the most part with only a brief fall to 96% shortly after 4am which quickly reversed.

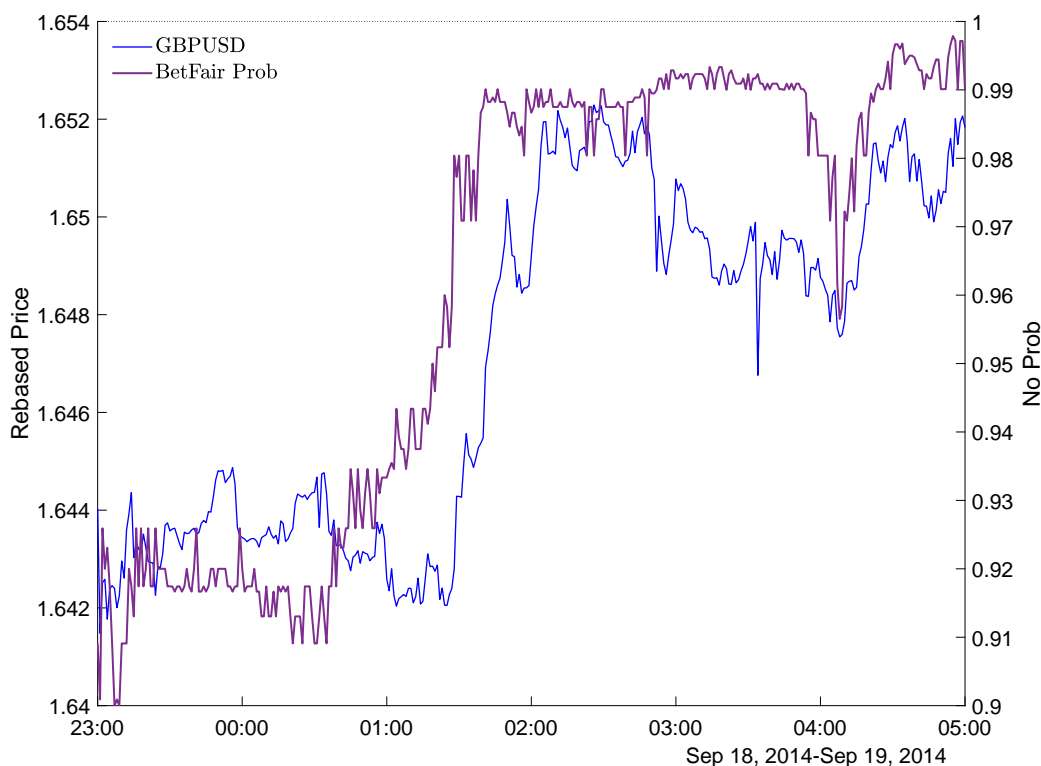


Figure 2.15: Betfair price for “No” bet and GBPUSD on the night of the independence referendum.

2.4.3.1 Evidence for cointegration

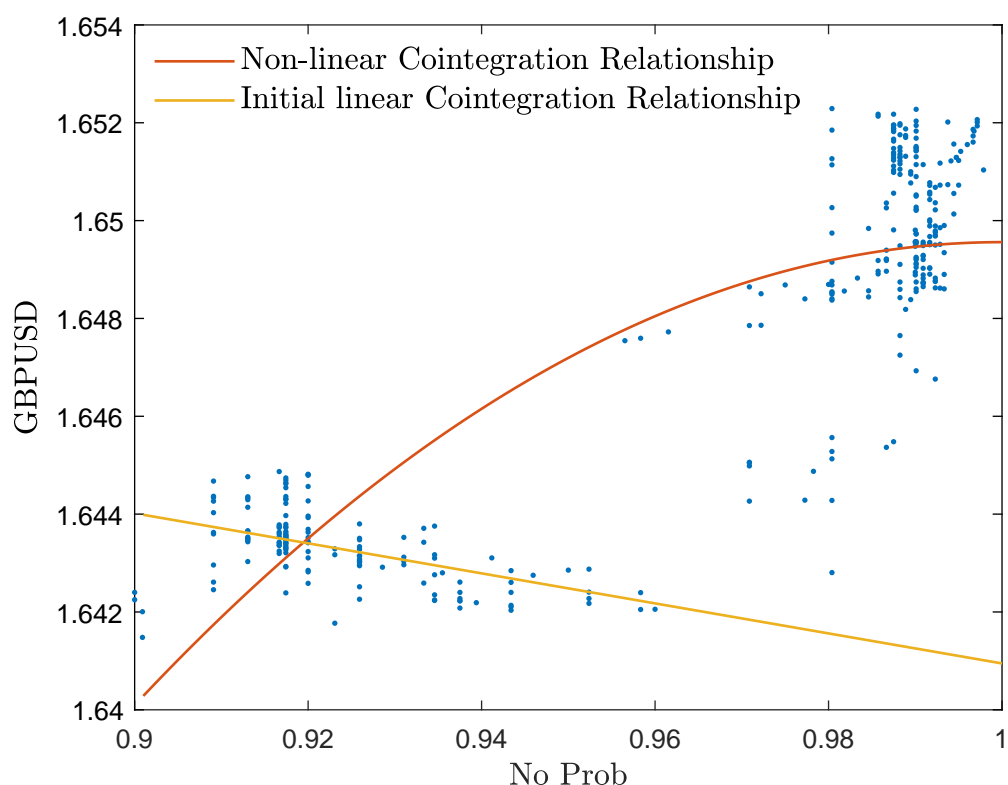
For this event we consider GBPUSD only. The currency was expected to sell off sharply in the event that the Yes camp prevailed, and recover a little if Scotland voted to stay in the UK. The price at 10pm was 1.6395, it peaked at 2:26am at 1.6523 and ended the night at 6am at 1.6476. The appreciation was modest but

this has to be considered against the smaller move in the betting market contracts of only around 10p.

In table 2.22 we show the results of various cointegration tests for GBPUSD and the betting contract. The first line of the table shows the results for linear cointegration. Both Phillips and Ouliaris tests as well as the Johansen max eigenvalue test do not reject the null of no cointegration at the 95% level. The Z_α test marginally rejects the null at the 90% level. This does not provide sufficient evidence for cointegration. The second line shows the results for the case where the risk factor is included in the cointegration regression (or in the system for the Johansen test). Here, the Phillips and Ouliaris tests do reject the null at the 99% level but the Johansen rank test does not. However, we note that the magnitude of the risk parameter λ is significantly larger than the linear parameter Δp . This violates our model as it results in a non-monotonic cointegrating relationship and utility function. The estimated value for the counterfactual conditional expectation of the pound given a Yes vote of $\hat{p}_0 = 3.18$ is nonsensical. The pound was expected to depreciate given the unlikely scenario of Scotland voting to be independent, not appreciate nearly 100%. The third and final line of the table shows the tests when the non-linear cointegration relationship is constrained to be positive with $|\lambda| = |\Delta p|$. This is effected by regressing GBP_t against the single monotonic factor $[PB_t + R_t]$. The three cointegration tests do not reject the null of no cointegration at the 90% level. Again the value of $\hat{p}_0 = 0.703$ is implausible. Although the pound was expected to depreciate significantly given a Yes vote, this value implies a huge fall of over 50%. After all, GBPUSD only sold off around 12% on the night of the Brexit referendum. Of all three models considered, the linear model seems the most realistic but there is still no compelling evidence that our model holds.

Table 2.22: Results of cointegration tests for the Scottish independence referendum.

Symbol	\hat{p}_0	$\hat{\Delta p}$	$\hat{\lambda}$	$p_{Z_t}(a = 0)$	$p_{Z_\alpha}(a = 0)$	$p(r = 0)^d$
GBPUSD ^a	1.5570	0.0937	–	0.126	0.099*	0.132
GBPUSD ^b	3.1833	-1.5309	-2.4388	0.008***	0.007***	0.161
GBPUSD ^c	0.7031	0.9464	0.9464	0.182	0.150	0.369
^a The first stage regression is $p_t = \hat{p}_0 + \hat{\Delta p}PB_t + \hat{\epsilon}_t$						
^b The first stage regression is $p_t = \hat{p}_0 + \hat{\Delta p}PB_t + \hat{\lambda}PB_t.(1 - PB_t) + \hat{\epsilon}_t$						
^c The first stage regression is $p_t = \hat{p}_0 + \hat{\Delta p}[PB_t + PB_t.(1 - PB_t)] + \hat{\epsilon}_t$ so $\hat{\lambda} = \hat{\Delta p}$						
The residual regression is $\hat{\epsilon}_t = a\epsilon_{t-1} + \eta_t$						
^d Result of the Johansen max eigenvalue test for zero cointegration rank						



Note: The non-linear relationship (orange line) is calculated ex-post on all data whereas the initial linear relationship (yellow line) uses only data from the first few hours when event probabilities were under 95%.

Figure 2.16: Spuriously estimated cointegrating relationships between the pound and the betting markets for the 2014 independence referendum.

We investigate the implied cointegration models in a little more depth. Figure 2.16 shows estimated cointegration relationships and market data for the night. The yellow line is the linear relationship found from regression using the first few hours when the betting probability of No was below 95%²⁴. This is *decreasing* which makes no sense. In this first period the betting contract only varied between 90 and 95p, and for the most part was no greater than 93p. Similarly the pound moved within a roughly 0.3 cent range. These were small noisy moves of no consequence as no information was released. The non-linear monotonic relationship fitted ex-post on all data is at least increasing. However, it is a poor fit to much of the data and explains the failure to reject the null of no cointegration in the non-

²⁴The probability of a vote for Scotland to remain in the UK rose above 95% at 1:21am, shortly before the Clackmannanshire announcement.

linear regression. Deviations from this relationship do not appear stationary nor converge. These insights provide further evidence that our pricing model does not apply on the night.

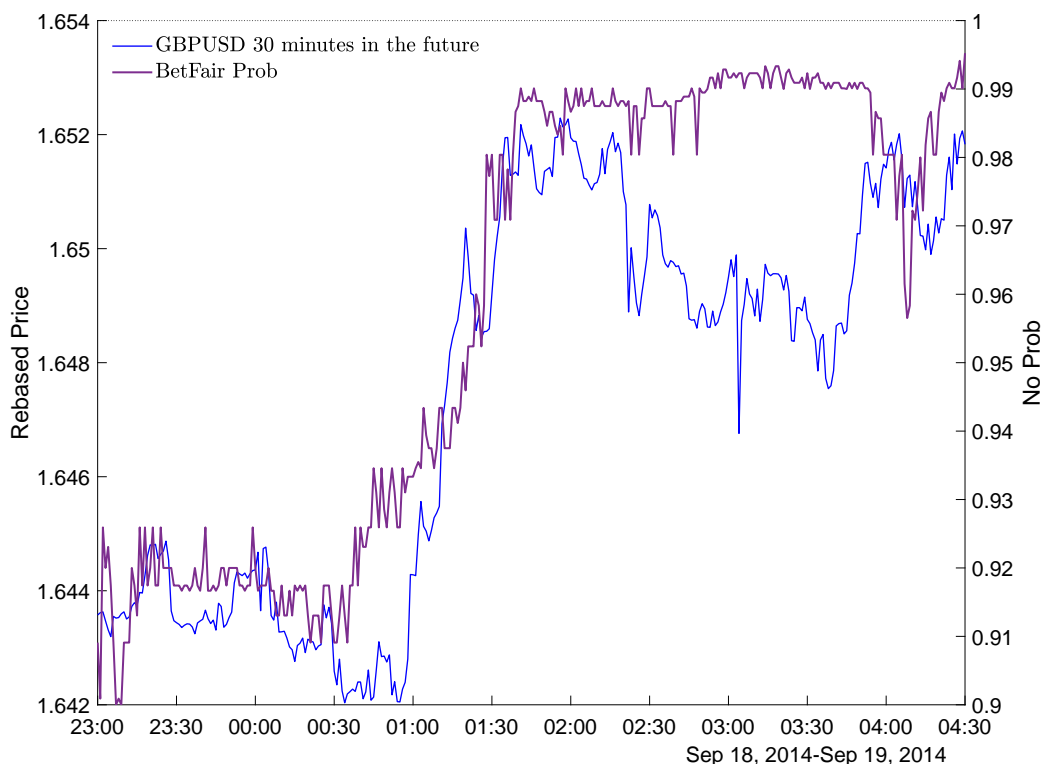


Figure 2.17: Betfair price for “No” bet and GBPUSD on the night of the independence referendum.

The truth is that on the night of the Scottish referendum the markets did not move a great deal, and what little movement there was for each market came relatively quickly and not quite in sync. The betting contract moved only 10p in the overnight session. This contrasts with the two events in 2016 where the results were a shock. The betting contracts moved close to their full range of £1. Our model fits well there as these political events dominated the markets. There was also sufficient coverage of price data there to generate statistical power, validating our model. In both cases in 2016 relevant information was released over a matter of hours. What is perhaps more interesting is the fact that the betting market appeared to lead the currency during the move up. This is similar to the political markets leading financial markets observed in the previous two events. There may in fact be a common trend in the markets relating to the probability of a ‘No’ vote;

it is just that the currency appears to lag that trend by around 30 minutes. To investigate this a little further we plot the betting price against the GBPUSD price bought forward 30 minutes from the future. This is shown in Figure 2.17. To the eye the market move is much more synchronized. We check the linear cointegration tests for this adjusted data. Results are shown in Table 2.23. We now see that there is evidence of a common trend between the two time series with the Phillips and Ouliaris tests both rejecting the null of no cointegration at the 95% level. The reader should note that this is not a rigorous statistical test of cointegration. The author has plotted more than one adjusted data series, choosing the one most pleasing to the eye.²⁵ A joint hypothesis test has effectively been performed whilst the pValues in Table 2.23 relate to only single hypothesis tests. Either way this exercise has demonstrated that the lack of statistical significance is likely due to the presence of a deviation of weak market efficiency of similar order in time to that of the move in the common trend. For the most part, and either side of the Clackmanshire result, the markets simply drifted.

Table 2.23: Results of linear cointegration tests for Scottish independence referendum where GBPUSD is shifted forward 30 minutes.

Symbol	\hat{p}_0	$\hat{\Delta}p$	$\hat{\lambda}$	$p_{Z_t}(a=0)$	$p_{Z_\alpha}(a=0)$	$p(r=0)^a$
GBPUSD ^a	1.5580	0.0934	—	0.041**	0.030**	0.601
The first stage regression is $p_t = \hat{p}_0 + \hat{\Delta}pPB_t + \hat{\epsilon}_t$						
The residual regression is $\hat{\epsilon}_t = a\epsilon_{t-1} + \eta_t$						
^a Result of the Johansen max eigenvalue test for zero cointegration rank						

2.5 Conclusion

In this chapter we attempted to describe the behaviour of financial and political betting markets in the hours after elections. The key assumption is that the only information that has a persistent affect on asset prices is that related to the *probability* of a political outcome (and nothing else). The further restrictions of equivalence of beliefs between the two markets (consistent with weak market efficiency), risk neutral pricing in betting markets and a ‘short’ overnight session lead to a non-linear cointegrating relationship between financial asset and binary options prices

²⁵Full disclosure: Graphs for GBPUSD bought forward by 20, 30, 35 and 40 minutes were eyeballed.

in the betting markets. When there are risk neutral investors in financial markets the relationship between the markets is one of standard linear cointegration. We test the model on data from three recent elections and find strong support for the theory for two of them.

For the 2016 Brexit referendum our base linear cointegrating model finds excellent support. The long term cointegration relationship is found to be so stable that a brave trader could estimate it in the first part of the night and profit from the convergence of deviations from that relationship in the second. It appears that market efficiency holds on a longer term basis that night (long term meaning hours) whereas there are some deviations from the model, and hence weak market efficiency on a timescale of tens of minutes. Our results are also in contrast to the model in Manasse et al. (2020) where a divergence from risk neutral preferences is required to explain the behaviour of the currency markets in the weeks and months preceding the vote. Of further interest is the fact that with respect to short term dynamics, the betting markets appear to lead and cause the financial markets but there is no apparent causation in the other direction. This suggests betting markets were more efficient at discounting the results of the vote in the hours after the referendum.

For the 2016 US presidential election linear cointegration, and hence risk neutrality, cannot explain the observed price action. However, we find strong evidence for our theory when deviations from risk neutrality are incorporated with a non-linear cointegrating relationship. The interpretation of the computed risk parameters are pleasing. Risk aversion is revealed for “risk-on” assets, whereas risk loving preferences are observed for the safe haven treasury future asset. Similar to Brexit we find deviations from the long term (albeit non-linear here) cointegrating relationship are predictable, indicating deviations from weak market efficiency on similar smaller timescales. The result that in the short run betting markets lead financial markets but that there is no causation in the other direction is also repeated. However, that brave trader who profited from the Brexit deviations would have had a terrible night repeating that strategy for this election. This is explained by the presence of non-trivial risk preferences. It is not possible to infer or estimate the non-linear cointegrating relationship in the first few hours. Foresight of the conditional expectations of the asset prices given election outcomes would be needed to make money.

Finally we do not find that our model holds on the night of the 2014 Scottish independence referendum. The markets did not move very much at all that night,

as the result was largely as expected. The move itself was also short lived. This is in contrast to the other two events where the results were both a huge shock, and large moves in all prices were observed that took several hours. The betting markets do appear to lead the currency market though by around 30 minutes. We explain the failure to demonstrate statistical significance as due to the similarity of timescales of deviations from market efficiency with those of the time taken for market to move upwards that night. For all three events it appears that betting markets reflected political information more quickly than financial markets, and on a timescale of minutes to tens of minutes.

Chapter 3

Political factor models

3.1 Introduction

In the previous chapter we investigated the price action of betting and financial markets in the hours after an election. But, the question remains, how do markets behave on a longer term, during the weeks and months preceding an election? This is the subject of this chapter.

The successful election night model used a key assumption. This was that asset prices are uniquely determined by the probability of a binary political outcome. Combined with other conditions this led to the presence of a single stochastic trend and the presence of cointegration. However, it is difficult to argue that, outside of the overnight session following the vote before the result is known, financial prices do not respond to the ebb and flow of other non-political information. One implication of the ‘No Shocks’ assumption in the previous model is that the conditional expectations of asset prices given binary political outcomes is fixed in time. This does not apply on a longer term basis. Conditional expectations of prices will change with the arrival of non-political information. Instead, the key assumption in this chapter is that the *difference* of the two conditional expectations is fixed. This is equivalent to assuming that the outcome of the elections has a fixed effect on the prospects of a company or financial asset. The residue of the price is allowed to vary. This is due to the existence of other economic and commercial information. The model leads to the existence of a political factor driving a part of the variance of asset price returns. The model is naturally extended for stocks using the ubiquitous Fama–French (Fama and French (2015)) factors to describe variance in returns not related to the political event.

This chapter proceeds as follows. In section 2 we briefly outline the failure of the cointegration model of the previous chapter on longer time-frames and explain why a new model is needed. Section 3 builds a pricing model using the assumption of a fixed difference in conditional expectations of stock prices given the outcome of a binary political event. The empirical and testing framework is outlined in section 4. Section 5 tests the model on several recent elections and section 6 concludes the chapter.




Figure 3.1: Difference in performance of the FTSE250 and FTSE100 indices versus the Betfair contract for Remain.

3.2 Motivation

We begin this section by applying the cointegration model of chapter 2 to the months preceding the Brexit referendum of 2016. The Conservative party manifesto for the 2015 general election contained a promise to hold an in/out referendum on whether or not the UK should leave the European Union. Following their victory, the European Union referendum Act 2015 was introduced to parliament at the end

of May 2015. It subsequently passed in early September 2015. The act gave a legal basis for a consultative referendum on the UK remaining within the EU by the end of 2017. On 20th February 2016 David Cameron, the Prime Minister, declared that the vote would take place on 23rd June 2016. The opposition Labour party was in favour of ‘Remain’ and the Conservative party was neutral. However, members of both main parties were free to campaign publicly on either side of the question, which began shortly after the announcement of the date.

Leaving the EU would no doubt lead to trade frictions between the UK and the bloc. It was argued by some that leaving the EU could result in a recession and a sharp depreciation of the currency. This would make the country poorer¹. It was envisaged that domestic focused companies would fare much worse than multi-nationals, much of whose revenues would be in appreciating foreign currency. Figure 3.1 shows the difference in the price series of the logarithm of the FTSE250 index and the logarithm of the FTSE100 with the price of the betting contract for Remain. The period of the figure is from the announcement of the date of the vote, to the day of the vote itself (the campaigning period). The FTSE100 is an index consisting of the 100 companies listed on the London Stock Exchange (LSE) with the highest market capitalisation. This index is dominated by large multinational companies with operations in multiple countries. The FTSE250 index consists of the next 250 largest companies on the LSE. These companies are typically smaller cap domestic-focused firms. Both the FTSE100 and FTSE250 are performance indices. Plotting the difference of the log of the indices in Figure 3.1 shows the relative performance of smaller domestic firms, potentially more affected by a vote for Brexit, versus larger companies that may even benefit from an appreciation of their earnings in British Pounds, in a Brexit scenario. We consider this difference series, as information unrelated to the upcoming referendum may effect both of the indices similarly. This is in contrast to changes in the odds of Brexit which is likely to will effect them differently. The series is an attempt to create a simple proxy for financial Brexit risk whose variance can mostly or entirely be explained by the changing odds of Brexit. This series does appear to the eye to be positively associated with the implied probability of Remain (or negatively related to Brexit). Increases in the probability of Remain do seem to occur with outperformance of the FTSE250 with the FTSE100. This is not consistent though. There was a period of underperformance in April which occurred without meaningful change

¹The UK had a current account deficit of around 5% of GDP at the time of the referendum.

Table 3.1: Results of the Phillips Ouliaris Z_t and Z_α tests for cointegration between $\log(\text{FTSE250}) - \log(\text{FTSE100})$ and the betting markets using daily data for the Brexit campaign period.

Symbol	\hat{p}_0	$\hat{\Delta}p$	$\hat{\lambda}$	$p_{Z_t}(a=0)$	$p_{Z_\alpha}(a=0)$	$p(r=0)^c$
$\log(\text{FTSE250}) - \log(\text{FTSE100})^a$	1.030	-0.095	-	0.178	0.117	0.394
$\log(\text{FTSE250}) - \log(\text{FTSE100})^a$	1.024	-0.131	0.080	0.338	0.289	0.174
^a The first stage regression is $\log(\text{FTSE250}) - \log(\text{FTSE100}) = \hat{p}_0 + \hat{\Delta}pPB_t + \hat{\epsilon}_t$						
^b The first stage regression is $\log(\text{FTSE250}) - \log(\text{FTSE100}) = \hat{p}_0 + \hat{\Delta}pPB_t + \hat{\lambda}PB_t \cdot (1 - PB_t) + \hat{\epsilon}_t$						
The residual regression is $\hat{\epsilon}_t = a\epsilon_{t-1} + \eta_t$						
^c Result of the Johansen max eigenvalue test for zero cointegration rank						

in the betting market. Throughout the campaigning period polls showed a small but consistent lead for Remain. The risk neutral probability of Remain implied by Betfair was between 60% and 82% during the campaign period but then increased on the day of the vote to up to 90%. As we now know, the UK did indeed vote to leave the EU. The confidence demonstrated on the day of the vote in a Remain outcome was apparently misplaced.

We now turn to the application of the overnight cointegration model. Table 3.1 shows the results tests for cointegration between the $\log(\text{FTSE250}) - \log(\text{FTSE100})$ price series and the betting contract that pays out £1 in the event of Brexit. All tests, both for the linear as well as non-linear cointegration, fail to reject the null of no cointegration. This no great surprise. The overnight model assumes that the only determinant of prices is related to the odds of Brexit. This may well occur on election night. However, over the period of around four months preceding the referendum there would have been numerous announcements made regarding both the domestic and international economy that had nothing to do with Brexit. Examining the ratio of the two indices was helpful but could not eliminate all information relating to non-political factors. The cointegration model, although hard to justify, and without econometric evidence, does not seem completely without use. The estimated conditional expectation of the market price series given Brexit is 0.935. The day of the vote the series closed at 1.006, thus the model predicted a fall of around 7% in the FTSE250 relative to the FTSE100. As it happened the price series fell to 0.964 on the day after the vote but continued its downward trajectory on the next trading day to 0.917. This price action and the conditional expectation are shown in Figure 3.2. The prediction of the model performed remarkably well.




Figure 3.2: Prediction based on cointegration model for $\log(\text{FTSE250}) - \log(\text{FTSE100})$ following the Brexit referendum.

Despite the apparent utility of building a model in levels, there is no theoretical justification for cointegration. The logical next step is to analyse first differences. The regression of the first difference of $\log(\text{FTSE250}) - \log(\text{FTSE100})$ on the first difference of the betting contract for Brexit over the campaigning period is significant. Results are shown in Table 3.2. The intercept is insignificant but the slope coefficient is significant at the 95% level². The slope parameter suggests a difference of around 4.5% in relative performance of the FTSE250 versus the FTSE100 given a counterfactual vote to remain versus the realised vote to leave. From the day of the vote (23rd June) to the day after (24th June) the betting contract for Brexit moved from 12% to 100%. This implies the next day value of $\log(\text{FTSE250}) - \log(\text{FTSE100})$ of 0.966 ± 0.016 ³. This is remarkably close to the actual value on 24th June of 0.964 and is closer to the forecast of the cointegration

²Leads and lags of the betting market $\Delta P(\text{Brexit})_{t-1}$ and $\Delta P(\text{Brexit})_{t+1}$ were insignificant.

³This is based on the change in the betting contract of 0.88 multiplied by the slope parameter. The MSE is ignored as it is of lower order

Table 3.2: A regression of the first differences of $\log(\text{FTSE250}) - \log(\text{FTSE100})$ and the betting probability of Brexit.

$\Delta [\log(\text{FTSE250}) - \log(\text{FTSE100})]_t$	Estimate	SE	tStat	pValue
intercept	4.6×10^{-5}	4.9×10^{-4}	0.09	0.0925
$\Delta P(\text{Brexit})_t$	-0.045**	0.019	-2.43	0.017

$N = 88, R^2 = 0.0643$

model. This is not surprising as the cointegration regression in levels will give equal weight to price values at the start of the training period as to those of the the end of the period. However, there will have been information not related to the Brexit vote between the start and end of the campaign. This information will have had a permanent effect of the level of prices and would not “unwind” given the realisation of the vote. Models based on differences or returns of prices do appear more appropriate than those in levels for longer-term time periods preceding an election. In the next section we will build such a theoretical model based on economic assumptions.

3.3 Pricing model

We present a model linking the prices of political and financial markets that applies on longer term periods of weeks or months in the run up to an election.

We begin by outlining the scenario of the model. Similar to the model in chapter 2 there exists a betting market which is liquid and trades multiple times a day in the run up to a scheduled event. Contracts are listed on the market relating to the outcome of a binary political event. Say $E = 1$ if this event occurs and $E = 0$ otherwise. The outcome of the election or event occurs at time $t = T$. For $t < T$ let PB_t be the price of the contract that pays out when $E = 1$. There are N financial assets indexed $i = 1, \dots, N$ whose prices are p_{it} at time t . For much of this section we consider only a single financial asset and then we label the price p_t . A difference between the scenario here and that of the overnight model is that a unit time period ($\Delta t=1$) is now equal to one day.

The key assumption in this model is that there is a constant effect of the election on the condition expectations of financial prices given the outcome of the event E .

CE – Constant effect of the political event

In the overnight model we assumed that the only determinant of the price of financial assets was the outcome of the political event. This led to conditional expectations given the outcome of the election as being fixed in time. We wish to relax this assumption for longer periods but still impose the fact that the election has some fixed political effect on prices. For instance, ex-ante, the market may believe Trump winning an election would have a 10% effect on the stock price of the Mexican bank Inbursa, versus a Clinton win⁴. This would be equivalent to

$$\frac{\mathbb{E}_t(\text{INBURSA}_T|E=1)}{\mathbb{E}_t(\text{INBURSA}_T|E=0)} = 0.9 \quad \forall t < T. \quad (3.3.1)$$

Similarly, for the duration of the election campaign, the market may have believed that the price would be 3 Pesos lower if Trump versus Clinton won⁵. This could be expressed as

$$\mathbb{E}_t(\text{INBURSA}_T|E=1) - \mathbb{E}_t(\text{INBURSA}_T|E=0) = 3 \quad \forall t < T. \quad (3.3.2)$$

In either case some function of the two conditional expectations is fixed for the time period before the event, and not allowed to vary. We will be applying our model typically to a relatively small period of time. As such, as long as the stock price does not vary too much the above two formulations of the CE assumptions will be approximately equivalent. Taking logarithms of the first formulation would, for a general price p_t , recover

$$\log(\mathbb{E}_t(p_T|E=1)) - \log(\mathbb{E}_t(p_T|E=0)) = \gamma \quad \forall t < T.$$

To make the assumption operational we make a further approximation. Stocks typically vary by much less than order 100% over a period of a few weeks. As such, $\log(p_t)$ will be approximately linear over the support of $p_t|E$. So

⁴Inbursa closed around 5.6% lower the day after Trump win, although it did initially trade significantly lower.

⁵Inbursa was trading in the range 27–30 Pesos in the weeks prior to the election.

$$\log(\mathbb{E}_t(p_T|E=i)) \approx \mathbb{E}_t(\log(p_T)|E=i) \quad \forall t < T.$$

Thus if either of the above forms of the CE assumption for Imbursa hold (equations 3.3.1 and 3.3.2), the quantity

$$\gamma_t = \mathbb{E}_t(\log(p_T)|E=1) - \mathbb{E}_t(\log(p_T)|E=0) = \gamma$$

will be approximately fixed in time for $t < T$. This is the mathematical form that our CE assumption takes.

The CE assumption will not hold for all assets all of the time. As an example, consider the case of a UK based bank unexpectedly announcing the opening of European operations during the Brexit referendum campaign. This may reduce the firm's reliance on cross-border regulations between the UK and the EU, reducing any negative effect of a vote for Brexit. As such it would be reasonable to expect γ to reduce for this bank over the announcement. Nonetheless this would be a fairly rare and specific event. Other scenarios could be systematic and affect large numbers of assets. For example, had the leave campaign in the Brexit campaign period unexpectedly announced that they had settled upon a policy of remaining in the EU single market should they win the vote ('softening' the Brexit), then one would expect the difference in conditional expectations to reduce. This would create a structural break in the political sensitivities γ . The strength of this assumption is a matter for debate. However, we believe for most assets and most elections it is fairly weak.

Derivation

We derive a model of financial asset and political betting prices using the key assumption of CE. In what follows below we assume a single asset price with price p_t at time t . $r_t = \log(p_t/p_{t-1}) = \log(p_t) - \log(p_{t-1}) = \Delta \log(p_t)$ the return of the asset on day t .

First we write down an equation linking the price of a financial asset today, and the expected value in the future. This is

$$\mathbb{E}_t(\log(p_T)) = \mu_{Tt} + \log(p_t)$$

where

$$\mu_{Tt} = \mathbb{E}_t(\log(p_T) - \log(p_t)).$$

Simply put, the expected log price of the asset price at the point the election result becomes known, $\mathbb{E}_t(\log(p_T))$, is the price today, p_t , plus the expected appreciation of the asset μ_{Tt} . This is of course equivalent to

$$\log(p_t) = -\mu_{Tt} + \mathbb{E}_t(\log(p_T)). \quad (3.3.3)$$

The purpose of this equation is to link future expected prices with the price today. Financial economics has many instances of such equations. These include the uncovered interest parity (UIP) relation applied to currencies. Another example is the cash and carry relationship linking the arbitrage relationship between forward or futures and spot prices. The common feature of these relationships is that any changes in future expectations of an asset's price are immediately reflected in changes to today's price.

The first difference of equation 3.3.3 recovers

$$r_t = \Delta \log(p_t) = -\Delta \mu_{Tt} + \Delta \mathbb{E}_t(\log(p_T)). \quad (3.3.4)$$

For notational convenience write

$$\begin{aligned} E_t^i &= \mathbb{E}_t(p_T | E = i) & i &= 1, 2 \\ P_t^i &= \mathbb{P}_t^f(E = i) & i &= 1, 2 \\ \gamma_t &= E_t^1 - E_t^0. \end{aligned}$$

where $\mathbb{P}_t^f(\cdot)$ represents the probability of an event as evaluated by a representative investor in the financial markets.

Next we expand the expectation of the log price using the total law of expectation.

$$\begin{aligned}
\mathbb{E}_t(\log(p_T)) &= P_t^1 \cdot E_t^1 + P_t^0 \cdot E_t^0 \\
&= P_t^1 \cdot E_t^1 + (1 - P_t^1) \cdot E_t^0 \\
&= E_t^0 + (E_t^1 - E_t^0) \cdot P_t^1 \\
&= E_t^0 + \gamma_t \cdot P_t^1
\end{aligned}$$

Taking first differences gives an expression for the second term of equation 3.3.4.

$$\begin{aligned}
\Delta \mathbb{E}_t(\log(p_T)) &= \Delta E_t^0 + \gamma_t \cdot P_t^1 - \gamma_{t-1} \cdot P_{t-1}^1 \\
&= \Delta E_t^0 + (\gamma_{t-1} + \Delta \gamma_t) \cdot P_t^1 - \gamma_{t-1} \cdot P_{t-1}^1 \\
&= \Delta E_t^0 + \gamma_{t-1} \cdot (P_t^1 - P_{t-1}^1) + \Delta \gamma_t \cdot P_t^1 \\
&= \Delta E_t^0 + \gamma_{t-1} \cdot \Delta P_t^1 + \Delta \gamma_t \cdot P_t^1
\end{aligned}$$

Note that $\Delta E_t^0 = \Delta E_t^1$ and that this quantity is the change in the expected future log price of the asset at time T , $\log(p_T)$, that is not related to the political event. Further, our CE assumption is precisely that $\Delta \gamma_t = 0$ and $\gamma_t = \gamma \forall t$. So CE \Rightarrow

$$\Delta \mathbb{E}_t(\log(p_T)) = \Delta E_t^0 + \gamma \cdot \Delta P_t^1. \quad (3.3.5)$$

Thus the CE assumption implies that the expectation of the log price from time $t - 1$ to t can be split up into a change not related to the upcoming election (ΔE_t^i $i = 1$ or 2) and the change due to the political event ($\gamma \cdot \Delta P_t^1$). The latter term is simply a constant multiplied by the financial market's evaluation of the probability of $E = 1$, P_t^1 . Assets that are expected to appreciate when $E = 1$ have $\gamma > 0$ whereas expected depreciation results in $\gamma < 0$. The greater the sensitivity to the political event the larger the magnitude of γ .

Cointegration model as a special case

Equation 3.3.5 decomposes asset returns into that due to changes in the probability of the outcome of the election ($\gamma \cdot \Delta P_t^1$), and that unrelated to the election (ΔE_t^0).

A key assumption of the overnight model, relaxed here, is that only changes in the odds of the event have an effect on asset prices. This is equivalent to setting $\Delta E_t^0 = 0$ in 3.3.5. This now becomes

$$\Delta \mathbb{E}_t(\log(p_T)) = \gamma \cdot \Delta P_t^1. \quad (3.3.6)$$

To first order, this is equivalent to

$$\Delta \mathbb{E}_t(p_T) = \gamma \cdot \Delta P_t^1.$$

Replacing γ with $\Delta p = \mathbb{E}(p_T|E = 1) - \mathbb{E}(p_T|E = 0)$, which is now fixed, recovers equation 3.3.5 of the preceding chapter. This equation, with the addition of market efficiency and risk neutrality, led to the derivation of the cointegrating relationship. Thus we see that the model of the preceding chapter is in fact a special case of that presented in this chapter, where the variance of factors affecting prices unrelated to the vote is zero.

Weak market efficiency and risk neutral pricing in betting markets.

As in the second chapter, we add a condition consistent with market efficiency equating the probability of $E = 1$ as assessed by a representative investor in the betting markets with that of $\mathbb{P}_t^f(E = 1)$,

$$\mathbb{P}_t^f(E = 1) = \mathbb{P}_t^B(E = 1)$$

where $\mathbb{P}_t^f(\cdot)$ is the probability of an event as evaluated by participants in the betting market. Next, the condition of risk neutral pricing in betting markets with an effective discount factor of unity⁶ implies that for a contract that pays out £1 when $E = 1$

$$PB_t = \mathbb{P}_t^B(E = 1)$$

⁶Our model applies only to the campaign period of an election of a few weeks or months. Thus any discount factor related to the payout of binary options that payout on the election result will be very close to one.

Thus the additional restrictions of equal beliefs in the two markets and risk neutral investors in betting markets \Rightarrow

$$\mathbb{P}_t^f(E = 1) = PB_t.$$

Taking first differences and substituting into equation 3.3.5 \Rightarrow

$$\Delta \mathbb{E}_t(\log(p_T)) = \Delta E_t^0 + \gamma \cdot \Delta PB_t. \quad (3.3.7)$$

To recover an expression for the daily return of an asset we still need to consider the change in the expected appreciation of the asset to the election result, $\Delta \mu_{Tt}$. This can be achieved using the fact that market efficiency implies that returns have a martingale property. First write

$$\begin{aligned} \mu_{Tt} &= \mathbb{E}_t(\log(p_T) - \log(p_t)) \\ &= \mathbb{E}_t(\log(\frac{p_T}{p_t})) \\ &= \mathbb{E}_t(\log(\frac{p_T}{p_{T-1}} \cdot \frac{p_{T-1}}{p_{T-2}} \cdot \dots \cdot \frac{p_{t+1}}{p_t})) \\ &= \mathbb{E}_t(\log(\frac{p_T}{p_{T-1}}) + \dots + \log(\frac{p_{t+1}}{p_t})) \\ &= \mathbb{E}_t(\log(\frac{p_T}{p_{T-1}})) + \dots + \mathbb{E}_t(\log(\frac{p_{t+1}}{p_t})). \end{aligned}$$

The final line uses the fact that the expectations are separable due to the martingale property of returns. If we assume an expected constant rate of daily return for the asset of μ then $\mathbb{E}_t(\log(\frac{p_{t+i}}{p_{t+i-1}})) = \mathbb{E}_{t-1}(\log(\frac{p_{t+i}}{p_{t+i-1}})) = \mu$ and

$$\mu_{Tt} = (T - t) \times \mu.$$

\Rightarrow

$$\begin{aligned}
\Delta\mu_{tT} &= \mu_{tT} - \mu_{t-1T} \\
&= (T-t)\mu - (T-t+1)\mu \\
&= -\mu.
\end{aligned} \tag{3.3.8}$$

Substituting equations 3.3.7 and 3.3.8 into equation 3.3.4 \Rightarrow

$$r_t = \mu + \Delta E_t^0 + \gamma \cdot \Delta PB_t. \tag{3.3.9}$$

This is linear equation linking changes in prices of betting market binary options and financial market returns. Changes in the odds of $E = 1$ affect returns directly through the factor loading γ . This is similar to the β s in the capital asset pricing model (CAPM). Whereas β in the CAPM relates to the sensitivity of the asset to the single factor market return, γ in our model represents the sensitivity to the political event.

Errors in market efficiency

We now consider what happens when beliefs in the betting and financial markets about the probability of $E = 1$ diverge. This would be consistent with deviations of weak market efficiency. To do this we add the stationary error ϵ_t to the condition linking beliefs in the two markets:

$$\mathbb{P}_t^f(E = 1) = \mathbb{P}_t^B(E = 1) + \epsilon_t = PB_t + \epsilon_t.$$

Equation 3.3.7 becomes

$$\Delta \mathbb{E}_t(\log(p_T)|\epsilon_t, \epsilon_{t-1}) = \Delta E_t^0 + \gamma \cdot \Delta PB_t + \gamma \cdot \Delta \epsilon_t$$

and equation 3.3.9 is varied to

$$r_t = \mu + \Delta E_t^0 + \gamma \cdot PB_t + \gamma \cdot \Delta \epsilon_t.$$

The factor model is now

$$r_t = \mu + \gamma \cdot \Delta P B_t + \zeta_t$$

$$\zeta_t = \Delta E_t^0 + \gamma \cdot \Delta \epsilon_t.$$

If market efficiency generally holds, but with some stationary fleeting errors between beliefs about the probability of the political event, the residual error of the factor model will be serially correlated.

3.3.1 Extension to multi-factor Fama–French model

We now consider a large portfolio of stocks, with returns r_{it} , $i = 1, \dots, N$. The single political factor model is

$$r_{it} = \mu_i + \gamma_i \cdot \Delta P B_t + \zeta_{it} \tag{3.3.10}$$

$$\zeta_{it} = \Delta E_{it}^0$$

$$\mathbb{E}(\zeta_{it}) = 0$$

The error ζ_{it} is allowed, indeed expected, to be correlated across different stocks. Under strict market efficiency it will be serially uncorrelated but fleeting deviations from efficiency will result in auto-correlation. Nonetheless the model above allows a parsimonious description of the covariance matrix of returns when a political event is upcoming. γ_i is the political factor loading and is a measure of stock i 's sensitivity to the upcoming election. Ceteris paribus, stocks with higher or lower γ s will have higher or lower correlation. $\mu_i + \xi_{it}$ is the part of the return not related to the political event. Indeed when there is no upcoming vote the returns model reduces to $r_{it} = \mu_i + \xi_{it}$. This is a poor description of the covariance structure of the portfolio of stocks as there are no common factors. We now extend the model by applying the ubiquitous Fama–French factor model to describe this part of stock returns.

The Fama–French factor model for equities is:

$$r_{it} - r_f = \alpha_i + \sum_{j=1}^K b_{ij} f_{jt} + \varepsilon_{it} \quad \mathbb{E}(\varepsilon_{it}) = 0 \quad i = 1, \dots, n \quad j = N, \dots, K. \tag{3.3.11}$$

There are N financial assets and K factors. r_{it} is the return of the i 'th stock at time t and r_f is the risk free return. The factors f_{1t}, \dots, f_{Kt} are K univariate random variables that vary with time. B is a $N \times K$ constant matrix describing the loading of the factors within the space of the N stocks (the elements of which are $\{b_{ij}\}$). ε is a vector of shocks, assumed to be serially uncorrelated and weakly correlated across different stocks as $N \rightarrow \infty$. The traditional market asset pricing model, the Capital Asset Pricing Model (CAPM), uses only one variable to describe the returns, being the market factor R_m ⁷. Fama and French's initial factor model (Fama and French, 1992) had three factors. These added factors related to the return of small versus large cap stocks and the returns of cheap versus expensive companies (measured by book-to-value ratios). The model was extended to five factors in Fama and French (2015). A factor relating to operating profitability as measured by profits to assets was added as was one based on a measure of investment by the company. The key result of the model is used to describe the cross-section of returns across stocks. It follows from arbitrage pricing theory and is

$$\mu_i = \mathbb{E}(r_i) = r_f + \sum_{j=1}^K b_{ij} \mathbb{E}(f_{jt}).$$

Equivalently $\alpha_i = 0 \ \forall i$ in equation 3.3.11. This implication has been studied in numerous papers and contexts, too many to discuss here. It is also not the topic of this chapter. We will not impose $\alpha_i = 0$. We simply wish to use the common factor approach to describe the non-political part of stock returns. Following the convention in the literature we write the excess return of stock i as $Z_{it} = r_{it} - r_f$. The most simple way to proceed is to replace the non-political part of the return, $\mu_i + \zeta_{it}$, in our one factor political model (equation 3.3.10) with the general factor expression above (equation 3.3.11). The result is

$$Z_{it} = \alpha_i + f_{1t} \cdot b_{i1} + \dots + f_{Kt} \cdot b_{iK} + \Delta P B_t \cdot \gamma_i + \eta_{it}. \quad (3.3.12)$$

The shock has been replaced by η in the full model above to distinguish it from the

⁷ R_m is the return of the market as a whole.

Fama–French residual, ε . We now have a $K + 1$ characteristic factor model which includes the additional political component.

ΔPB_t is unlikely to be uncorrelated with the factors $\{f_{jt}\}$. Indeed, the outcome of most political events will be expected to have an overall effect on the prices of stocks. Changes in the probability of $E = 1$ will be highly likely to affect the returns of the market R_m . When this is the case it is important to note that the factor loadings in the full political factor model (equation 3.3.12), will be different from those in the standard Fama–French model (equation 3.3.11). This is also the case for the political loadings $\{\gamma_i\}$. These loadings in the full model will differ from those in the single political factor model (equation 3.3.10). Only in the unlikely event of ΔPB_t being uncorrelated with $\{f_{jt}\}$ would loadings be equal between the full model and the sub-factor models.

In this chapter we are not interested in the accurate estimation of factor loadings in any one of these particular models. The purpose of the investigation is to examine whether political markets are explanatory of stock returns. As long as both the political loadings $\{\gamma_i\}$ are non-zero in either model, and ΔPB_t is not collinear with $\{f_{jt}\}$, this will be the case. It will also follow that ΔPB_t will be explanatory of the Fama–French residuals $\{\varepsilon_{it}\}$. When collinearity occurs, loadings in the full model are not identifiable, but, more importantly, adding betting market information to the model does not help describe stock returns. This would be because changes in the betting markets would already be fully explained by the Fama–French factors.

3.4 Empirical specification

Our theoretical model for stock returns is summarised as the full factor model

$$Z_{it} = \alpha_i + f_{1t}.b_{i1} + \cdots + f_{Kt}.b_{iK} + \Delta PB_t.\gamma_i + \eta_{it} \quad i = 1, \dots, N, \quad t \leq T.$$

The model shows that betting markets are explanatory of stock return i if and only if $\gamma_i \neq 0$. Testing the validity of our model is equivalent to testing that the set of parameters $\{\gamma_i\}$ are significant. Rejecting $H_0 : \gamma_i = 0 \forall i$ would be strong

evidence in favour for the model holding, stronger than assuming the model as the null hypothesis and failing to reject.

The model itself is a $K + 1$ characteristic factor model. The historical daily K Fama–French factor returns are readily available from Professor K.R. French’s website⁸. Histories from around 1990 are provided for download. The final political factor ΔPB_t is available from betting markets for periods preceding a political event. In this chapter we rely on data solely from the Betfair exchange platform. This acts like a limit order book. Contracts may be listed months and even years before an election. However, liquidity generally increases as the event approaches. Far out from an election, on most days there may be no trading at all. ΔPB_t appears in the model as a measurement of the changes in the probability of a political event $\mathbb{P}_t(E = 1)$. However, unless the betting market is trading and liquid then the measurement will not be valid. Care will be required to ensure there is sufficient liquidity when choosing which time period to apply the model to. In practice only a small number of months or a few weeks directly prior to the election will be considered.

One could proceed by estimating the full model above on this small period and attempting to generate significance for γ . However, this would be inefficient. The much longer history of the Fama–French factors would not be exploited for estimation.


First choose a period where we judge the betting market to be sufficiently liquid to be a valid measure of beliefs about the probability of the political event. Call this period the testing period. Say it starts at T_1 and ends at T_2 . (In practice T_2 will be the day before the election result is announced, $T - 1$). Define a longer period starting at T_0 and ending at $T_1 - 1$. Call this the training period. Fama–French factor data from both the testing and training periods will be used but only betting data from the testing set is considered. The approach is illustrated in Figure 3.3.

One way to proceed would be to estimate the following regression

$$Z_{it} = \hat{\alpha}_i + f_{1t}.\hat{b}_{i1} + \cdots + f_{Kt}.\hat{b}_{iK} + [\delta_{t,\{T_1, T_2\}} \times \Delta PB_t] . \hat{\gamma}_i + \hat{\eta}_{it} \quad t = T_0, \dots, T_2 \quad (3.4.1)$$

where

⁸https://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html



Fama–French factor loadings are estimated in the training window. Explanatory power of the betting market on the Fama–French residual is tested in the testing window.

Figure 3.3: Schematic of the empirical approach.

$$\delta_{t,\{T_1,T_2\}} = I(t \in \{T_1, T_2\}).$$

This replaces the political factor with $\delta_{t,\{T_1,T_2\}} \times \Delta PB_t$. This forces the factor to zero outside of the liquid testing period. However, this is a flawed approach. Political risk may still be present prior to the testing period, and changing. An election, say, three or four months in the future may be well known and the prospects for each of the candidates varying. However, it just may be too far in the future to be in the minds of those that choose to place bets in a political market. Lack of participation in a betting market may make an evaluation of the probability of an election outcome unmeasurable. This is not the same as it being constant and unchanging though. Conducting a regression where it is assumed to be zero in the training period will lead to invalid results.

To overcome this issue, and use more of the history of the Fama–French factors, we proceed as follows:

1. Estimate the Fama–French loadings $\hat{\alpha}, \hat{B}$ during the training period $t = T_0, \dots, T_1 - 1$.

2. Evaluate the Fama–French estimated residuals $\hat{\varepsilon}_{it} = Z_{it} - \hat{\alpha}_i - \sum_{j=1}^K \hat{b}_{ij} f_{jt}$ on the testing period $t = T_1, \dots, T_2$.
3. Regress the estimated Fama–French residuals $\hat{\varepsilon}_{it}$ on changes in the political factor ΔPB_t on the testing period $t = T_1, \dots, T_2$ without intercept.

The above approach involves the same regression as that of equation 3.4.1 under the null hypothesis. For most event studies it would also produce identical statistics under alternatives of interest. However, if the political factor is in fact changing in the training set then the two step approach has greater power. This is demonstrated in Appendix C. As discussed in the previous section the Fama–French loadings estimated in step 1 above will not be the same as the loadings in the full factor model. This is due to the very likely correlation of ΔPB_t with the other factors $\{f_{1t}, \dots, f_{Kt}\}$. However, significance of γ_i in the regression

$$\hat{\varepsilon}_{it} = \alpha + \gamma_i \Delta PB_t + \xi_{it} \quad t = T_1, \dots, T_2$$

is equivalent to significance of γ_i in the full factor model. Note that we allow an intercept in the first stage regression as we do not impose the arbitrage pricing theory constraint. However, for the second stage we drop the intercept. This is because both $\mathbb{E}(\hat{\varepsilon}_{it}) = 0$ and $\mathbb{E}(\Delta PB_t) = 0$. The level of PB_t will change in the testing period but the unconditional expectation is zero⁹. Of course there will be estimation error in $\hat{\alpha}$, \hat{B} . This will lead to estimation error of $\hat{\varepsilon}_{it}$. However, the error will be in the dependent variable in the second stage regression. Further, as the training and testing periods are disjoint, the estimation error will be independent of the residual of the final equation ξ_{it} . No endogeneity will be present in the final step and estimates of γ_i will be unbiased, although there will be a reduction in precision.

3.4.1 Significance tests

In the next section we will test our model on a series of political events from recent years. For each event we will select an appropriate universe of stocks and a betting contract and follow the process described above. There are two tests we will perform which are outlined below.

⁹There is an exception to $\mathbb{E}(\Delta PB_t) = 0$. This is that the final Fama–French residual $\mathbb{E}(\varepsilon_{iT})$ has positive expectation, as political risk, not present in the general factor model, disappears. We do not however, in practice, include the final period in the testing set. PB_t is at all other times a martingale.

3.4.1.1 Portfolio test

The simplest test we apply is to consider the univariate regression for an equally weighted portfolio of the stocks of the chosen universe. If $\hat{\varepsilon}_{it}$ is the estimated Fama–French residual for stock i then define $\bar{\hat{\varepsilon}}_t = \frac{1}{N} \sum_{i=1}^N \hat{\varepsilon}_{it}$, the average estimated residual. We then conduct the following regression

$$\bar{\hat{\varepsilon}}_t = \bar{\gamma} \cdot \Delta PB_t + \xi_t \quad t = T_1, \dots, T_2.$$

If $\gamma_i = 0 \forall i$ then $\bar{\gamma} = 0$ in the above regression. Thus rejection of the null hypothesis $\bar{\gamma} = 0$ rejects the null hypothesis of insignificance of the betting market. This is a simple test to perform. It can also easily be made robust to serially correlated errors that may occur due to the presence in errors in market efficiency, as well as heteroskedasticity. However, the test is not robust to alternatives where the set of individual stock γ s can take different signs.

3.4.1.2 J_α tests of Pesaran and Yamagata

The above portfolio test may suffer from a lack of power. There is loss of information by grouping all securities together in a single portfolio. We wish to use a test which can exploit information from a large number of stocks. As we are only conducting tests over periods where the betting market will be liquid, T is likely to be of the order of 10s, and certainly no more than 100. This is small compared to the number of stocks available. For example, there are around 5000 stocks listed in the US alone. There is the possibility of using large number of assets to generate significance but this will require a test to apply when $N > T$.

Recall that the key pricing implication of arbitrage pricing theory is that the excess returns, $\{\alpha_i\}$ in the regression equation 3.3.11 are zero. There is a large literature in empirical finance on testing this implication. Whereas the empirical question there is whether or not the intercepts in the panel model of excess stock returns on factors is zero, we are concerned with whether the slope coefficients $\{\gamma_i\}$ are significant. However, many of the existing methods translate simply when applied to the slope coefficient rather than the intercept. Reviews of the literature can be found in Black et al. (1972) and Fama and French (2004). Jensen (1968) was the first to use t-statistics to test the significance of α for a given security in the CAPM equation. However, when there are large numbers of stocks, individual

tests lose their meaning. A joint test is sought. Further, the expected cross sectional dependence of the residuals leads to correlation of the individual t-statistics which means combining them into a single statistic is non-trivial. The standard test in the literature that addresses this problem is that of Gibbons et al. (1989). This is an exact test based on multivariate statistics. The test assumes normally distributed regression errors. It is also valid only when $N < T$. To overcome these limitations, typically monthly returns are used (to minimise deviations from non-normality), and securities are grouped into a small number of portfolios (to reduce N). Although superior to testing a single portfolio or stock, there is still loss of information, and potentially statistical power, due to grouping the securities into a smaller number of portfolios. There is also the possibility of introducing endogeneity when a large number of stocks are used as the portfolio return may become related to the market return factor R_m . Beaulieu et al. (2007) present a test of $\alpha_i = 0$. This is based on using simulation methods to calculate the test size for a wide class of non-normal distributions. Gungor and Luger (2009) and Gungor and Luger (2013) present distribution free non-parametric tests but are not robust to cross sectional dependence nor asymmetric distributions. Neither of these approaches is valid for $N > T$ though.

The problem of testing $\alpha_i = 0$ when the number of stocks is large relative to T was solved in Pesaran and Yamagata (2012). This uses a normalised Wald statistic, J_α , with a thresholding estimator on the cross-sectional error correlations which is robust to weak cross sectional dependence. Further, the test is demonstrated to be asymptotically valid in the case of non-normal errors. Monte Carlo evidence shows the test performs well in small samples. T is tested at $T = 60$ and $T = 100$ and N varied from 500 to as low as 50. Minimal size distortions are observed for the larger values of N considered. These values of N and T are in the range of those used in our application. The test is based on a Wald statistic of the individual t-statistics of univariate regressions, adjusted for (threshold applied) cross sectional correlations. It can be readily applied to testing the significance of the slope γ_i and it is this test that will be primarily relied upon in the next section. Call this the J_γ test. Improvements to the J_α (and other similar tests) have been proposed by Fan et al. (2015). They demonstrate that the J_α test lacks power against sparse alternatives where, for example, a finite number of stocks have significant α s. They add a power enhancement term to the test statistic which vanishes asymptotically under the null but has power against such sparse alternatives. This may well be a relevant alternative when testing arbitrage pricing theory, where a small number of stocks

may be responsible for failures in market efficiency. The empirical application in their paper, studying returns of the S&P500 index components, certainly suggests so. However, we do not believe this is relevant in our application. If betting markets are significant we expect explanatory effects beyond a small number of stocks, or particular industry sector. The number of stocks with significant α s is likely to grow as $n \rightarrow \infty$. The J_α (or J_γ) tests should have power against the alternatives of interest and we will rely solely on them without power adjustment.

The test is summarised as follows. Consider the following factor model in the form of a panel data regression and stacked by cross-sectional regressions

$$\mathbf{y}_t = \boldsymbol{\alpha} + \mathbf{B}\mathbf{f}_t + \mathbf{u}_t.$$

y_i are the individual stock returns, and \mathbf{f}_t are K known factors. \mathbf{B} is an $N \times K$ matrix of factor loadings. The Wald statistic for $\boldsymbol{\alpha} = 0$ can be estimated as

$$\hat{W}_\alpha = \sum_{i=1}^N t_{\alpha,i}^2$$

where $t_{\alpha,i}$ is the t-ratio of the intercept of the OLS regression of $(\mathbf{y}_t)_i$ on intercept and \mathbf{f}_t . In the case of cross sectionally independent errors it can be shown that under various regularity conditions

$$\begin{aligned} \mathbb{E}(W_\alpha) &\rightarrow \frac{\nu N}{\nu - 2} \\ \text{VAR}(W_\alpha) &\rightarrow \frac{2N(\nu - 1)}{(\nu - 4)} \left(\frac{\nu}{\nu - 2} \right)^2 \end{aligned}$$

as $N \rightarrow \infty$, $T \rightarrow \infty$ and $N/T \rightarrow \infty$ where $\nu = T - K - 1 > 4$. The following is an exactly standardised statistic

$$\hat{J}_{\alpha,1} = \frac{N^{-1/2} \sum_{i=1}^N \left(t_{\alpha,i}^2 - \frac{\nu}{\nu-2} \right)}{\left(\frac{\nu}{\nu-2} \right) \sqrt{\frac{2(\nu-1)}{(\nu-4)}}}$$

and thus is distributed as $N(0, 1)$ asymptotically under the null. The result holds in the case of non-normal errors. A second statistic $\hat{J}_{\alpha,2}$ is derived that is adjusted

for correlation of the individual t-statistics when the errors are cross sectionally correlated. The adjustment is based on a consistent estimate of the correlation matrix of the disturbances \mathbf{u}_t . Define

$$\hat{\rho}^2 = \frac{2}{N(N-1)} \sum_{i=2}^N \sum_{j=1}^{i-1} \hat{\rho}_{ij}^2 I(\nu \hat{\rho}_{ij}^2 \geq \theta_N)$$

where $\hat{\rho}_{ij}$ is the sample correlation of the regression residuals $\hat{\mathbf{u}}_t$ and θ_N is some threshold value. The latter is chosen so that the number of non-zero correlations decline steadily with N . We will follow the protocol set out in Pesaran and Yamagata (2012) with $\sqrt{\theta_N} = \Phi^{-1}(1 - \frac{p_N}{2})$ and $p_N = 0.1$. The adjusted statistic is now

$$\hat{J}_{\alpha,2} = \hat{J}_{\alpha,1} \times \frac{1}{\sqrt{[1 + (N-1)\hat{\rho}^2]}} = \frac{N^{-1/2} \sum_{i=1}^N (t_{\alpha,i}^2 - \frac{\nu}{\nu-2})}{(\frac{\nu}{\nu-2}) \sqrt{\frac{2(\nu-1)}{(\nu-4)} [1 + (N-1)\hat{\rho}^2]}}.$$

Pesaran and Yamagata (2012) show that $\hat{J}_{\alpha,2}$ is asymptotically distributed as $N(0,1)$ under various stricter regularity conditions and weak cross sectional correlation. They also show, via Monte Carlo simulation, small deviations from correct sizes for sample sizes similar to those used in this chapter. Of course $|\hat{J}_{\alpha,2}| < |\hat{J}_{\alpha,1}|$ since it has lower variance due to the correlation adjustment of the t-statistics. $\hat{J}_{\alpha,2}$ will never reject the null when $\hat{J}_{\alpha,1}$ does not. We will not consider the first statistic at all as we do expect that errors will indeed have some correlation.

The test generalises simply to the slope parameter. We thus consider

$$\hat{J}_{\gamma,2} = \frac{N^{-1/2} \sum_{i=1}^N (t_{\gamma,i}^2 - \frac{\nu}{\nu-2})}{(\frac{\nu}{\nu-2}) \sqrt{\frac{2(\nu-1)}{(\nu-4)} [1 + (N-1)\hat{\rho}^2]}}$$

where $t_{\gamma,i}$ is the t-statistic for slope in the regression of the Fama–French residual $\hat{\varepsilon}_{it}$, on intercept and changes in the betting market, ΔPB_t .

This test has been demonstrated to have excellent properties for our setting where there may be non-normal errors and some correlation amongst stock residuals. The test however is not demonstrated to be robust in the case of serial correlation (possibly present due to inefficiencies) or heteroskedasticity. In practice we will test the residuals of the second stage regression for the presence of these

effects. The lower power univariate portfolio test does have one advantage over this large N test which is that it is simple to use robust errors.

3.5 Results

This section presents the empirical approach applied to real world data from six+ elections from recent years. To apply the method several choices will need to be made for each event. These include which particular betting contract to use, what portfolio of stocks to consider and what training and testing periods to use. A plot of the logarithm of trailing 7-day average daily volume on the Betfair exchanges for our chosen event is shown in Figure 3.4. Any key announcement concerning the event is also shown on the figure. As can be seen, liquidity explodes exponentially as the day of voting approaches, but can be very low a few months out. For example, the daily volumes trading on exchange are under £10k 5 months from several of the elections. We will need to treat data from lower liquidity days cautiously. Below, we discuss each event in more detail. Table 3.3 lists each election along with their chosen specifications for testing. Note that we test two portfolios for the Brexit referendum as well as the 2016 US presidential election. Also, 1, 3 and 5 factor models are available for each Fama–French choice, the 1 factor equating to CAPM.

3.5.1 Events

2014 Scottish independence referendum

We have already analysed this event in the overnight section. The vote took place on 18th September 2014. The question on the ballot paper was ‘Should Scotland become an independent country’. To recap, polls showed a consistent lead for ‘No’ during July and August of that year. Polls tightened in September. There was even a poll that showed a small lead for ‘Yes’ published on 6th September. The pound and companies linked to Scotland depreciated on that day and the betting contract paying out £1 for ‘Yes’ rallied to 35p. The risk neutral Betfair implied probabilities can be seen from Figure 3.5. The figure shows the prices along with the chosen testing windows for this and the other political events. Ultimately though, polls reverted and the ‘No’ side prevailed. As this was a binary Yes/No referendum the choice of betting contract is simple. We chose the ‘No’ bet in our model. For stocks we chose all companies listed on the LSE, from Q2 2014, domiciled in the UK. Many small less liquid companies will be removed from the




Figure 3.4: Liquidity on Betfair in the months prior to our chosen political events.

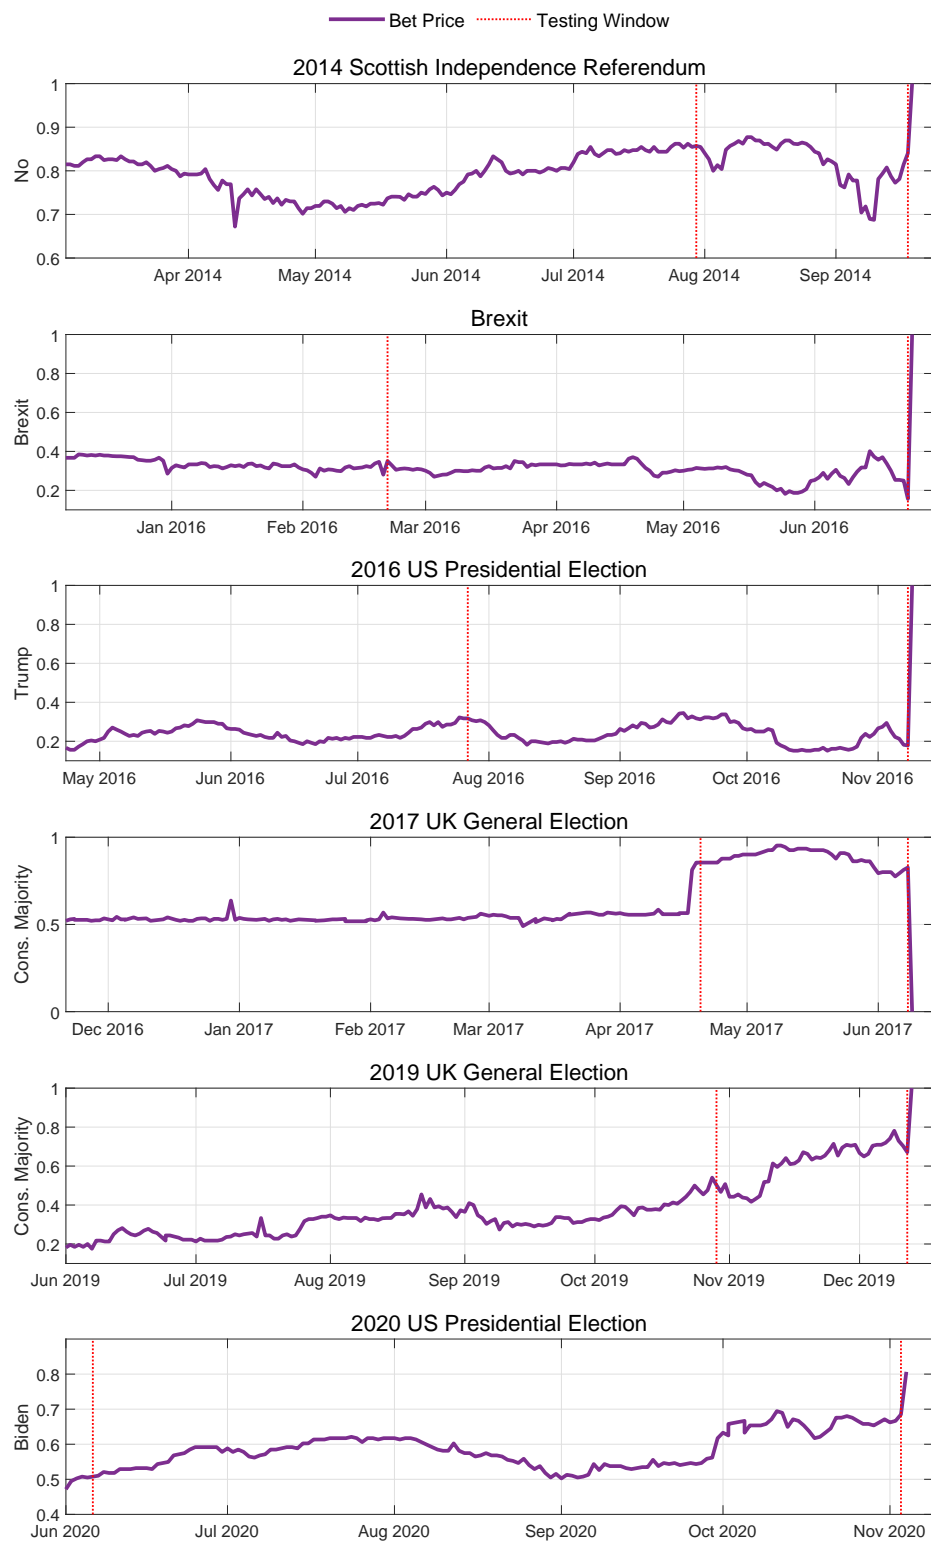


Figure 3.5: Betfair implied probabilities.

portfolio according to a further screening step. This is discussed at the end of this subsection and is applied to all the portfolios studied in this chapter. Fama–French daily factor returns are available for the European market and we utilise this as the most appropriate choice¹⁰. The choice of testing period is not straightforward. The Scottish Government announced the date of the referendum on 21 March 2013, around 18 months prior to the vote itself. Liquidity on betting markets was muted though. Little more than around a few thousand pounds exchanged hands per day on Betfair through April, May and June 2014. The volume increased to around £10k per day in July and peaked at well over £1m a day in the week before the referendum. We chose to start the testing window on July 30th to give us sufficient time samples in the testing window. The volume did start to significantly increase from around £25k per day from this point.

2016 Brexit referendum

This event was discussed in the previous chapter too. Again this is a binary vote and so the choice of betting contract is straight forward. We use the contract for ‘Leave’. On 20th February 2016 David Cameron announced the date of the referendum to parliament as 23rd June 2016. Liquidity improved on Betfair from this point with around £100k per day trading the week after the announcement, with volume peaking at around £50m on the overnight session following the vote. This makes our choice of testing period easier. We chose the day after the announcement to begin the testing window. The betting percentage odds during this period put the probability of leaving the EU from between the high teens to the high 30s. However, by the day of the vote there was a widespread (misplaced) belief that the country would vote to remain. The implied probability of a Leave vote bottomed out around 10%, just after the vote closed. We make the similar choices for the equity universe and Fama–French factors as the 2014 Scottish independence referendum, using stocks listed on the UK based LSE from Q1 2016. We also separately test a portfolio of stocks that are made up of EU27 based companies that earned at least 25% of their revenues in the UK according to their 2015 full year accounts.

¹⁰Daily developed market factors are also available but the European centric-data set is preferred.

Table 3.3: Specifications for each political event.

Event	Bet	Portfolios	Factors	StartTest	EndTest ¹
2014 Scottish independence referendum	No	Stocks listed on LSE in Q2 2014	Fama/French European	30-Jul-2014	18-Sep-2014
Brexit	Yes	Stocks listed on LSE in Q1 2016 European Stocks with > 25% revenues from UK in 2015	Fama/French European Fama/French European	21-Feb-2016	23-Jun-2016
2016 US presidential Election	Trump	S&P500 in Q2 2016 Stocks Domiciled in Mexico in Q2 2016	Fama/French North American Fama/French North American	27-Jul-2016	8-Nov-2016
2017 UK General Election	Conservative Majority	Stocks listed on LSE in Q1 2017	Fama/French European	20-Apr-2017	8-Jun-2017
2019 UK General Election	Conservative Majority	Stocks listed on LSE in Q3 2019	Fama/French European	29-Oct-2019	12-Dec-2019
2020 US presidential Election	Biden	S&P500 in Q2 2020	Fama/French North American	6-Jun-2020	3-Nov-2020

¹The test period ends on the day of the vote. The result is generally known one day later.

2016 US presidential election

The election of Donald Trump as the 45th president of the United States was the second political surprise of 2016. As election day is set under statute as ‘the Tuesday next after the first Monday in the month of November’, the date is known well in advance. However, the two candidates are not nominated until the summer before the vote. Recall that our model assumes a fixed difference between an expectation of a particular outcome (say X) and the complement of that outcome (\bar{X}). Even if one candidate is known, the complement, that is the opponent, may not be. Thus the expectation of the complement can change without the betting odds for the known candidate changing. For example, Donald Trump’s nomination came before the Democratic candidate. The odds of the particular Democratic candidate can change, effecting asset prices, without a corresponding change to the probability of Donald Trump winning. So our model will not hold. Until both candidates are known the model will not apply. The vote is not binary.¹¹

¹¹The vote for US president is in fact never completely binary. There is always the chance that a candidate may become incapacitated during the campaign. Running mates often trade with

Trump was officially nominated at the Republican national convention on 19th July, Clinton, on July 26th at the Democrat event. Liquidity for the election was much higher than the preceding events, with hundreds of thousands of pounds exchanging hands on Betfair around the time of the nominations. We are comfortable with using the betting markets from the day after both nominations were known. We chose the contract for Donald Trump for the model. Trump was (incorrectly) never seen as the stronger candidate. The percentage odds of him winning never rose above the mid 30s during the testing window. For this event we will use two different stock universes and test both. The first will be the index stocks of the S&P500 in Q2 of 2016; the second, publicly traded stocks domiciled in Mexico. We do not limit ourselves to companies listed on the major Mexican exchanges as many top Mexican stocks trade offshore. As discussed in the previous chapter, it was expected that a Trump win would be very bad for trade with Mexico and hence its markets. The relevant Fama–French factors for both these portfolios are the North American ones.

2017 UK General Election

The result of the 2016 Brexit referendum was a surprise. The ruling conservative party had a working majority of 17 in the house of commons.¹² There was no election due under the Fixed-term Parliaments Act until May 2020. However, following the UK’s triggering of Article 50 of the Treaty on European Union, the Prime Minister Theresa May called a surprise election. This was to ‘strengthen her hand’ in negotiations with the EU. Ultimately though she failed in that aim. She lost her majority and following the result governed as the leader of a minority government with the support of the Democratic Unionist party.

Legislation to enact a general election was ratified by parliament on 19th April, 2017. The date was set for 8th June 2017. Until 18th April, the day before the announcement, liquidity for this event on Betfair was minimal, struggling to reach even £100 per day. However, trading jumped to around £100k a day once it became clear an election was imminent. It increased to over £1m per day in the days before the vote. The choice of testing set is clear with this event as there is simply no liquidity prior to the date of the election’s announcement.

very small but positive probabilities. The morbidly interested reader can look at the betting odds for Kamala Harris in the 2020 election to impute the market’s belief about the odds of the elderly president Biden passing during his campaign!

¹²There are 650 seats in the house of commons and 17 is considered a small majority which makes governing difficult.

Until the snap election became public, the odds for a conservative majority were barely trading. For the few trades that did occur, the implied probability was around 50%. The lack of liquidity was likely due to the fact that the election was not scheduled for several years. Gamblers were neither interested in, nor informed about, the details of the election. This observation is remarkably consistent with the findings of Page and Clemen (2013) that there is a ‘bias of the price towards 50%’ for events far into the future. Until the announcement of the election date, the vote was indeed expected to be ‘far into the future’. Going into the campaign, opinion polls consistently showed a lead for the Conservatives. They tightened as the vote approached, Wikipedia (2017). Implied odds for a majority rallied to over 80% when the election was announced. They stayed in the range 78–95 until the vote. The choice of which contract to use is not straightforward. If the conservatives did not achieve a majority, then there were potentially multiple alternatives. Bets for a ‘Labour majority’, a ‘hung parliament’ or ‘anything else’ were listed. However, the conservative’s lead over the Labour party was such that the implied probability of an actual Labour majority was very small. For the most part of the campaign, odds for this outcome hovered around 1–2.5%. It did increase a little as the polls tightened going into the election but was never implied to be higher than 4%, except for a few trades on election day. This means the event was close to binary with the majority of the probability distributed between a conservative majority and a hung parliament. However, we should note our model does not hold exactly. In terms of the universe of stocks, we will stick with all stocks headquartered in the UK listed on the LSE (in Q1 2017), with the European Fama–French factors.

2019 UK General Election

Following the UK government losing its majority in the house of commons in 2017 there was prolonged political deadlock. This led to Theresa May resigning as Prime Minister. Boris Johnson was elected as her replacement by the Conservative party in the summer of 2019. Johnson could not, though, convince the house of commons to pass a revised withdrawal agreement¹³. This caused him to call a snap election, the third general election in 4 years. Legislation was passed on 28th October 2019 and the date was set for 12th December. As was the case with the 2017 general election, trading on Betfair rose significantly over the announcement. For example,

¹³The withdrawal agreement was made between the UK and the EU and established the terms of the UK’s withdrawal from the EU.

the day before only £2,600 traded, whereas the day after volume jumped to over £100k. Volumes generally increased from this point with over £2m per day trading in the final week. Given the uptick in volume we will again start the testing period the day after the announcement.

Throughout the campaign period, the conservatives held a strong lead over the Labour party (Wikipedia, 2019). However, the implied odds of a Conservative majority was relatively low compared to opinion polls. It traded in the range 40–55% in the first fortnight of the campaign. This was likely due to the underperformance of the conservatives versus initial polling and expectations in the 2017 election. The odds did steadily increase to around 70–80% as Labour failed to significantly tighten the polls. The conservatives ultimately gained a landslide win with a huge majority of 80 seats in the house of commons. We again use the bet for ‘Conservative majority’. The prices for a labour majority did initially imply odds as high as 6% but this soon dropped to the 2–3% range as they failed to make ground. The alternative to a conservative majority was again dominated by a hung parliament. We again use the same choices for the equity universe as the preceding UK events, with UK stocks listed on the LSE in Q3 2019.

2020 US presidential election

In 2020 the presidential election was scheduled for 3rd November. There was never any doubt that Trump, as sitting president, would run for a second term. However, there was a political battle for the heart of the Democratic party between establishment Joe Biden and the left winger Bernie Sanders. On June 5th Joe Biden gained the required number of delegates from the Democratic national convention to secure the nomination. Volume on Betfair for this event was in the hundreds of thousands of pounds a day at the time of the nomination. This increased to around £1m a day in the days immediately prior to the vote. As we now know, president Biden did not secure the Presidency the day after the election. In the days and weeks after election day there were numerous recounts. There were also several (seemingly baseless) legal challenges by the Trump campaign looking to overturn what was apparently a clear win for the Democratic ticket. Contracts continued to trade on Betfair until 14th December 2020. The electoral college confirmed Joe Biden’s victory in the election on that date. There is no doubt that there are numerous opportunities to study financial and political markets in the period 9th November to 14th December 2020. However, we do not believe the model presented

in this chapter will apply. This is because the prospects of a Joe Biden win and its alternative will have fundamentally changed. Instead of political risk relating to this election disappearing overnight on 8th November, it substantially increased. The prospect of Donald Trump retaining the presidency via a typical presidential election and one where the apparent rightful winner is deposed by the courts (or worse) is very different. The risk becomes less about what the electorate is voting for and more about the strength of US institutions. Political implications aside, the conditional expectations of asset prices given a Biden or Trump win certainly changed on 4th November 2020. Given this, we have no confidence the model will hold. We end the testing set on 3rd November. We use two equity universes to test this event: the index stocks of the S&P500 in Q2 of 2020 and the North American Fama–French factors.

3.5.2 Data sources and handling

Equity data

For each event we have selected a universe of stocks and Fama–French factors. Our aim is to ascertain if there is evidence to reject the null hypothesis $\gamma = 0$. The input price series we use to generate stock returns is the CRSP adjusted close data. This adjusts the price to account for dividends, stock splits, other distributions and rights issues¹⁴. We use a variety of sources to conduct stock screens and source price data, including S&P Capital IQ, Thomson Reuters Datastream and Yahoo Finance. We do not exclude data points for days with specific company results or other announcements. This will no doubt introduce variance unrelated to either the political event or the Fama–French factors. Removing the corresponding return for such announcements would ‘clean’ the data set. Doing so would likely reduce the variance of the unexplained part of our regressions and potentially improve significance. It would also involve a considerable computational effort. We do not take this step, choosing to focus on testing a greater number of events, rather than a smaller number with cleaner data. This means there is an implicit assumption in our method. This is that the frequency of such stocks specific events is similar in training and testing sets. There is no reason to believe that this assumption is not valid.

¹⁴Adjustment methodology can be found at <https://www.crsp.org/products/documentation/crsp-calculations>.

Many of the stock universes we have chosen include all equities listed on a particular exchange, or domiciled in a particular country. This will include many illiquid and infrequently traded stocks. They may not respond to changes in the odds of a political event as they simply may not be traded. There may also be stocks included that do not exist for all of the training and testing sets. The estimates of the factor loadings may not be accurate for such stocks. To remove such securities we apply a filter to the portfolio. Firstly we remove any stocks that were not trading at the start of the training set. Secondly we only consider stocks that have trades, and non-zero returns, for 90% of days in both the training and testing sets. In practice this reduces the size of our stocks universes significantly.

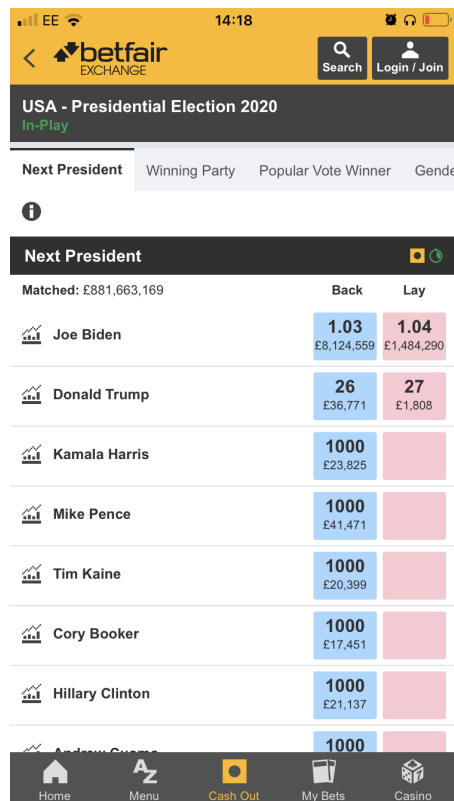


Figure 3.6: Betfair Exchange market for the ‘Next President’ on 21st Nov 2020, 18 days after voting closed.

Betting data

The betting company Betfair runs a platform called Betfair Exchange. This functions like a limit order book. For sporting, political and other events they list

markets where orders for bets can be placed and executed. For each market, various selections are listed. Only one of the selections on each market will ultimately win. Participants can either ‘back’ or ‘lay’ each selection. The prices quoted are in terms of odds. Backers of the winners receive the odds multiplied by their stake, less a small percentage commission levied by Betfair. The payout, gross of commission, is paid by the participant that laid the relevant matched bet. The risk neutral probability, ignoring commission, for a selection with odds o is $1/o$. In this thesis we only quote probabilities and not odds and think of each bet as a binary contract. Figure 3.6 shows a screen shot of the Betfair market for the 2020 US presidential election on 21st November 2020. This was in the height of the political mayhem that followed the election when Donald Trump was contesting results with various legal challenges. It can be seen that the risk neutral probability of him winning is seen at almost 4%, although we believe the discount to the payout from a Biden win is largely due to risk and cost of capital matters. Note that all contenders apart from Trump and Biden have orders to lay at 1000 to 1 (1000 is the maximum odds allowed on the exchange). The market has reflected the removal of the other candidates from the presidential race.

The raw data we use from Betfair are matched trades on the exchange for the relevant market. For non-binary markets, which have multiple possible outcomes, such as the US presidential election shown in Figure 3.6, we only consider trades in a single selection (corresponding to $E = 1$). We handle binary events (Brexit and the Scottish independence referendum) differently. Here we use both selections, converting the alternative implied probability to the chosen selection via $\mathbb{P}(E = 1) = 1 - \mathbb{P}(E = 0)$. The market is sufficiently efficient that arbitrage opportunities do not exist by holding every selection on a market (doing so guarantees a payout). Trades also happen at multiple times throughout the day and at various prices. We need to convert the trades into a daily difference that aligns with the stock returns. We do this by choosing the trade that happens closest to the time of the relevant equity market close on each day.

Table 3.4: Regression of Fama–French factor returns on ΔPB_t for political events in the testing window.

	European Factors					American Factors			
	Slope	Error	tStat	pValue		Slope	Error	tStat	pValue
	2014 Scottish independence ref.					2016 US election			
MktRF	1.15	3.24	0.36	0.722		-5.24	4.76	-1.10	0.272
SMB	-0.21	1.06	-0.20	0.840		-2.71*	1.51	-1.80	0.072
HML	-0.67	1.31	-0.51	0.608		2.07	2.41	0.86	0.390
RMW	0.89	1.17	0.76	0.445		1.37	1.36	1.01	0.313
CMA	-0.70	0.89	-0.79	0.431		2.98**	1.39	2.14	0.033
				$\overline{T} = 36$					$\overline{T} = 74$
	2016 Brexit referendum					2020 US election			
MktRF	-16.28***	5.00	-3.26	0.001		10.69	7.51	1.42	0.155
SMB	3.16*	1.90	1.66	0.097		-0.69	4.93	-0.14	0.889
HML	-0.83	1.81	-0.46	0.646		-2.11	7.51	-0.28	0.779
RMW	0.57	1.47	0.39	0.698		-0.78	2.89	-0.27	0.788
CMA	0.34	0.62	0.54	0.588		1.76	3.92	0.45	0.654
				$\overline{T} = 88$					$\overline{T} = 106$
	2017 UK GE								
MktRF	0.23	4.92	0.05	0.963					
SMB	0.35	1.88	0.19	0.851					
HML	5.34**	2.15	2.48	0.013					
RMW	-3.41*	1.77	-1.93	0.053					
CMA	2.70*	1.39	1.94	0.053					
				$\overline{T} = 35$					
	2019 UK GE								
MktRF	-0.30	1.66	-0.18	0.856					
SMB	0.91	0.98	0.93	0.353					
HML	-1.30	1.25	-1.04	0.297					
RMW	0.22	0.60	0.36	0.716					
CMA	0.64	0.74	0.86	0.390					
				$\overline{T} = 32$					

Heteroskedastic robust errors are used.

Fama–French factor data

Fama–French daily factor returns are available for download from Professor French’s website¹⁵. We source data for the five factor model for both the training and testing periods. The factors are labelled MktRF, the market return, SMB, the small minus big size companies, HML, high minus low value stocks, RMW, robust minus weak operating profitability and CMA, conservative minus aggressive investment.

In terms of the size of the training set, we settle on using two years worth of data. We use the data immediately prior to the testing set to estimate the Fama–French loadings. We judge this period sufficient to estimate the factor loadings. We avoid choosing a longer training set to avoid the possibility of structural breaks in stock loadings. Over long periods of time the factor loadings may change for a stock. For example, small companies can become large ones.

3.5.3 Results

Firstly we examine relationships between the factor returns and the betting markets. We regress each factor on ΔPB_t and an intercept in the event window for each political event. Note we exclude the day after the election. This is because both the changes in betting odds, and market moves, can be very large on that day. Including this day will dominate the OLS regression. Further, returns, and Fama–French factors, are calculated based on market close. If we were to include the day after the election, it would be better to look at the market open as this will be the first snapshot of prices post-result. There can be large intra-day volatility the day after the outcome becomes known. For example, after the 2016 US presidential election the market rout that followed the surprise win was totally reversed within the trading day. This followed a reassuring morning address from president Trump. If there is a relationship between the betting and stock markets we want to test that it is persistent over a reasonable length of time and not just when the election result becomes known. Results for the slope for these regressions are shown in Table 3.4. Errors robust to heteroskedasticity are used. The betting markets do not explain the variation in the European factors for the Scottish independence referendum. This is not surprising as this was a UK based event and the factors are based on the whole of the European market. Correspondingly though for the Brexit referendum, there is a significant relationship at the 99% level for the market return, and at

¹⁵https://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html

the 90% level for the SMB factor. Brexit was an event that had potential effects beyond the UK and could affect many large European economies and exporters. (The Fama–French European factor covers both the UK and mainland Europe.) The estimated slopes in the regressions suggests a difference in price of the average European company of 17% ($\pm 5\%$) between the UK remaining in the EU and leaving. Smaller companies would be expected to outperform by 3% ($\pm 1.9\%$) under Brexit, perhaps as smaller EU companies are less likely to export to the UK than large ones. The 2017 UK general election has surprising results. MktRF and SMB do not appear related to the betting market, but the other three factors do. This election was called by the Prime Minister Theresa May to ‘strengthen her hand’ in any Brexit negotiations. The outcome of this election could have been expected to have an effect on the type of deal the UK and the EU would ultimately strike. It then follows that the probability of Theresa May gaining a majority could have an effect on stock prices throughout Europe. However, it is puzzling that effects were seen on stock prices according to their relative value, operating profitability and investment strategy, but not on the overall level of the market. Further, no significant effects were seen on the Fama–French factors in the general election two years later, where the new Prime Minister Boris Johnson was vowing to ‘Get Brexit Done’. Turning now to the 2016 US presidential elections, we see significant effects in the SMB and CMA factors. The relationship between MktRF and ΔPB_t is not seen as significant though. This is a surprise given the huge sell off in asset prices observed on the night of the election itself when Donald Trump unexpectedly won. We do note that the estimated slope coefficient does suggest the market would be 5.2% ($\pm 4.8\%$) lower if Trump were to win¹⁶. For the 2020 election, no significant effects are seen on the Fama–French factors from the betting market. To experiment to see if the betting markets become more informative as the election approaches we re-run the regressions, but begin the event window only a calendar month before the dates of the election. The new results are shown in Table 3.5 and are striking. All factors are now significant for the 2016 election bar RMW and MktRF and CMA are significant for the 2020 election. Biden is now seen as having a rather unbelievable positive effect of around 25% ($\pm 8\%$) on the stock market if he were to replace Trump as president. This suggests that the betting markets become more informative as the election nears. Note that liquidity also exponentially increases during this time too.

¹⁶The S&P500 future initially sold off by 4.8% overnight. The odds for Trump increased from around 20% to 99% during this time.

Table 3.5: Regression of Fama–French factor returns on ΔPB_t for US presidential elections from the month prior to election day.

	2016 US election					2020 US election			
	Slope	Error	tStat	pValue		Slope	Error	tStat	pValue
MktRF	-12.03**	(5.92)	-2.03	0.042		23.42**	(10.54)	2.22	0.026
SMB	-4.89**	(2.22)	-2.20	0.028		2.18	(7.34)	0.30	0.766
HML	4.45	(2.96)	1.50	0.133		18.27*	(11.02)	1.66	0.097
RMW	1.82	(1.60)	1.14	0.254		3.12	(4.83)	0.65	0.517
CMA	5.09***	(1.61)	3.16	0.002		13.62**	(5.60)	2.43	0.015
				$\overline{T} = 22$					$\overline{T} = 22$

Heteroskedastic robust errors are used.

3.5.3.1 Significance tests

We now turn to results of significance tests for the model. As with the Fama–French factor regressions, we exclude the day after the election from the testing period. Including the day after in the regressions would almost certainly generate significance for any events with surprise results. Returns will likely have very large values on that day. However this does not test that the model holds over a general period and would relate to significance ‘after the fact’. We apply the steps in the empirical specification for the residuals of the K=0, 1 and 5 factor Fama–French models. For K=0 we simply regress stock returns, less their mean from the training set, on ΔPB_t . K=1 corresponds to residuals from the CAPM model. Both the mean weighted portfolio and individual stock Pesaran and Yamagata tests are run. Results are shown in Table 3.6. We perform various diagnostics to check for serial correlation and heteroskedasticity.

From the theoretical pricing model recall that the presence of serial correlation in the regression errors is consistent with errors in market efficiency. It would also invalidate the results of the J_γ^2 test. Ljung Box statistics are calculated for all regressions, including the multivariate form of the statistic. There is little evidence of serial correlation from the mean weighted portfolio regressions. Of the 24 regressions conducted a single one has a significant Ljung-Box statistic at the 95% level. This is for the one factor model for the 2016 US presidential election. This is not repeated for the other factor models, nor for the individual

Table 3.6: Results of the significance tests for γ .

Univariate regressions ^a									Individual stock regressions						
K	$\hat{\gamma}$	$\hat{\sigma}_{\gamma}$	t	p	R^2	p_Q ^b	p_{LM}^{BP} ^c	p_{LM}^E ^d	$\bar{\gamma}$	\bar{t}	$\sqrt{\bar{t}^2}$	$\hat{J}_{\gamma,2}$	$P_{J_{\gamma,2}}$	p_Q ^e	p_{LLn} ^f
2014 Scottish independence referendum															
0	0.025	(0.036)	0.70	0.483	0.016	0.234	0.575	0.138	0.025	0.29	1.17	2.01***	0.022	1.000	0.898
1	0.017	(0.016)	1.06	0.290	0.028	0.676	0.543	0.179	0.017	0.22	1.22	3.33***	<0.001	1.000	0.193
5	0.018	(0.016)	1.08	0.282	0.032	0.798	0.276	0.354	0.018	0.21	1.22	3.37***	<0.001	1.000	0.426
$N = 188, T = 36$															
2016 Brexit referendum - LSE stocks															
0	-0.147***	(0.042)	-3.51	<0.001	0.171	0.848	0.946	0.006***	-0.147	-1.92	2.28	15.66***	<0.001	1.000	0.014**
1	-0.028*	(0.016)	-1.70	0.089	0.025	0.775	0.600	0.192	-0.028	-0.30	1.39	6.02***	<0.001	1.000	0.970
5	-0.026	(0.016)	-1.63	0.103	0.021	0.795	0.588	0.328	-0.026	-0.27	1.34	5.88***	<0.001	1.000	0.705
$N = 219, T = 85$															
2016 Brexit referendum - European exporters to UK															
0	-0.133***	(0.041)	-3.26	0.001	0.121	0.755	0.773	0.071*	-0.133	-1.64	1.89	9.06***	<0.001	1.000	0.723
1	0.010	(0.021)	0.48	0.631	0.002	0.186	0.251	0.487	0.010	0.15	0.91	-1.16	0.878	1.000	0.909
5	0.014	(0.020)	0.71	0.479	0.004	0.112	0.341	0.593	0.014	0.18	0.94	-0.85	0.802	1.000	0.843
$N = 102, T = 87$															
2016 US presidential election - S&P500 constituents															
0	-0.041	(0.049)	-0.82	0.413	0.016	0.031**	0.575	0.136	-0.040	-0.52	1.07	0.52	0.300	1.000	0.536
1	0.009	(0.007)	1.51	0.132	0.021	0.414	0.092*	0.772	0.011	0.17	1.06	0.49	0.313	1.000	0.531
5	0.005	(0.006)	0.92	0.356	0.006	0.558	0.575	0.410	0.006	0.04	1.03	0.32	0.376	1.000	0.557
$N = 488, T = 74$															
2016 US presidential election - Mexican stocks															
0	-0.180***	(0.047)	-3.86	<0.001	0.229	0.569	0.311	0.969	-0.180	-2.05	2.35	14.91***	<0.001	0.520	0.644
1	-0.156***	(0.029)	-5.32	<0.001	0.270	0.970	0.749	0.895	-0.155	-1.85	2.20	15.66***	<0.001	0.518	0.795
5	-0.158***	(0.030)	-5.24	<0.001	0.281	0.906	0.603	0.985	-0.158	-1.88	2.22	16.33***	<0.001	0.421	0.675
$N = 51, T = 75$															
2017 UK general election - LSE stocks															
0	-0.008	(0.030)	-0.28	0.781	0.001	0.606	0.910	0.743	-0.008	-0.07	0.99	-0.66	0.745	1.000	0.661
1	-0.012	(0.039)	-0.31	0.754	0.004	0.151	0.194	0.595	-0.012	-0.13	1.00	-0.46	0.677	1.000	0.924
5	-0.011	(0.040)	-0.27	0.789	0.003	0.196	0.211	0.679	-0.011	-0.07	0.99	-0.63	0.735	1.000	0.438
$N = 215, T = 34$															
2019 UK general election - LSE stocks															
0	0.014	(0.019)	0.72	0.472	0.007	0.648	0.527	0.821	0.014	0.344	1.28	3.72***	<0.001	1.000	0.945
1	0.014	(0.010)	1.30	0.193	0.019	0.772	0.338	0.043**	0.014	0.36	1.37	6.86***	<0.001	1.000	0.947
5	0.014	(0.010)	1.35	0.178	0.018	0.984	0.414	0.047**	0.014	0.38	1.47	9.18***	<0.001	1.000	0.871
$N = 243, T = 33$															
2020 US presidential election - S&P500 constituents															
0	0.083	(0.092)	0.79	0.432	0.004	0.598	0.509	0.210	0.074	0.49	0.86	-0.43	0.672	1.000	0.860
1	0.004	(0.049)	-0.11	0.911	0.000	0.194	0.429	0.368	-0.005	0.01	0.84	-0.81	0.417	1.000	0.396
5	0.006	(0.022)	-0.15	0.884	0.000	0.868	0.521	0.114	-0.003	-0.01	0.88	-1.13	0.260	1.000	0.330
$N = 487, T = 105$															

^a Heteroskedastic robust errors are used for the univariate portfolio regression^b Ljung-Box serial correlation test p-value^c Breusch-Pagan heteroskedasticity test p-value^d Engle ARCH test p-value^e Multivariate Ljung-Box serial correlation test p-value^f Ling-Li multivariate heteroskedasticity test p-value

stock regressions for that election. We ascribe little relevance to this observation. There is also no evidence for serial correlations from the portfolio regressions with no significant Ljung-Box statistics. The conclusion is that the data demonstrate no evidence of market failure. Chapters 1 and 2 found deviations of weak market efficiency of the order of minutes to tens of minutes. It is not surprising that no inefficiency was found here. We consider only less frequent daily data. Discovering an inefficiency on that time-scale would be a much more surprising result.¹⁷

Turning now to heteroskedasticity we conduct both the Breusch-Pagan and Engle ARCH tests for the univariate portfolio regressions. The Engle test provides some evidence of deviations from homoskedasticity in these univariate regressions. We note though that only one event, the 2019 UK general election, is significant, at the 95% level, once the first market wide factor has been controlled for. However, results of these regressions will be valid as heteroskedasticity robust errors are employed. For the $J_{\gamma,2}$ tests we conduct the Ling and Li (1997) test. We observe a single significant value at the 95% level for the zero factor model in the Brexit test. Again we do not ascribe much relevance to this finding as it is not repeated for the higher factor models. In general we are confident that conclusions drawn from the modified Pesaran and Yamagata $J_{\gamma,2}$ tests are valid.

In short, we observe highly significant values of γ for four of the six events studied. We also note that, as expected, the $J_{\gamma,2}$ tests have higher power than the univariate mean-weighted portfolio test. Significance of the portfolio test is only ever found when the $J_{\gamma,2}$ is significant, but the opposite is not true. We discuss the results below in more detail.

For the 2014 Scottish independence referendum, $J_{\gamma,2}$ is significant at the 95% level for $K=0$ and at the 99.9% level for $K=1$ and 5. This is consistent with the conclusions of Darby et al. (2019) and Acker and Duck (2015) that showed that the CAPM residuals are related to betting odds. However, we go further and demonstrate that there is still explanatory power in the betting markets when the additional Fama–French factors are controlled for. We note that the average value of γ implies an average decrease in the price of UK stocks of 2.5% between Scotland voting for independence and voting to remain in the UK. The regressions on the residuals of the higher factor models estimates the decline relative to the European wide market. This is because the factors used are Europe wide. The

¹⁷Given the number of regressions performed, we would expect a small number of seemingly significant Ljung Box statistics at the 95% level under the null hypothesis of no serial correlation.

underperformance of UK stocks relative to their European peers is estimated at 1.7–1.8% given a vote for independence versus Scotland remaining in the UK.

The 2016 Brexit referendum also produced highly significant findings for UK based equities listed on the LSE. All the models produced $J_{\gamma,2}$ significant at the 99.9% level. Two of the three univariate regressions are also significant at the 99.9% level for $K=0$, and at the 90% level for $K=1$. Estimated values of γ imply very steep falls in stocks given a vote for Brexit. The $K=0$ model suggests an overall decline of 15% (individual regressions) relative to a vote to Remain. The $K=1$ and 5 models estimate declines for UK stocks relative to the overall European market of around 2.6–2.8%. Results for the Brexit referendum for the European exporters portfolio are interesting. Recall that the betting markets were highly explanatory for this European wide market factor. γ is significant for $K=0$ but once we control for the first market factor, MktRF , they become insignificant. The significance seen for European exporters in the raw returns is due to ΔPB_t driving the wider European market. However, ΔPB_t does not explain the CAPM residual, even for European companies that generated over 25% of their revenues in the UK in 2015. These companies were not expected to perform significantly differently than the wider European stock market given a vote for Brexit. We note that the estimated fall for these European exporters of 13% is consistent with both the estimated fall of UK stocks (around 15%) and the estimated underperformance of UK stocks (3%) to within standard errors. Recall also that a regression of the European wide MktRF factor on the betting market for this event was significant. The estimated effect of a 1% change in the odds of Brexit led to a fall of $0.16\% \pm 0.05\%$. Again this is broadly consistent with the fall of European exporters estimated in the portfolio regressions.

The 2016 US presidential election was a major event for Mexico. Changes in the odds for a Trump victory were highly significant for Mexican stocks, both outright and relative to the North American market. All univariate and multivariate tests were significant at the 99.9% level. The difference between a Clinton and a Trump win was estimated to have an average effect on Mexican stock prices of 18%. The decline relative to the North American market was estimated at around 16%. We note that for the $J_{\gamma,2}$ test $N < T$. Although the test is not asymptotically valid in such cases, Monte Carlo simulation of test statistics presented in Pesaran and Yamagata (2012) suggest good performance for similar values of N and T when $N < T$. They also demonstrate for these values of N and T superior performance to the standard multivariate test of Gibbons et al. (1989). Furthermore we

do not need to rely on the modified test of Pesaran and Yamagata here as the weaker mean-weighted portfolio test generated significance anyway. For the historical S&P500 constituents our tests do not generate significance, even for $K=0$. This is particularly surprising given the sell off on the night of the election itself when Trump's surprise win became apparent. This is a surprising result, but not dissimilar to the finding that ΔPB_t did not explain MktRF when regressing over the full event window. This will be explored further later.

Our final significant event was the 2019 UK general election. Boris Johnson was the firm favourite to win this election promising to 'Get Brexit Done'. His opponent Jeremy Corbyn had doubled down on his left wing policies after failing to win the 2017 election. He launched what was seen as the most left wing set of policy ideas seen for a generation, Pickard (2019). Policies included raising taxes sharply for companies and higher earners as well as nationalising key UK companies. A failure for Boris Johnson to gain a majority may have opened up the possibility of a reversal of Brexit as opposition parties were suggesting a second referendum. However, for asset markets, Jeremy Corbyn was now apparently seen as the bigger threat. Also, failure to gain a majority could likely continue the intense political uncertainty that had dogged the UK for the preceding few years. A conservative majority was seen as significantly positive for UK stocks, both outright and relative to the European market. $J_{\gamma,2}$ is significant for all three models (at the 99.9% level). Estimated values of λ imply an average premium for stocks of 1.4% under a conservative majority.

The two events that did not generate significant γ are the 2017 UK general election and the 2020 US presidential election. It may be the case for these events that either, one, $\gamma = 0$, and the political event is not informative for asset prices, or two, that the elections are informative but that noise in the data fail to generate statistically significant regression coefficients due to γ being small relative to the variance of the full model residual η_t . We cannot be certain of the answer. However, it is plausible that neither event is in fact, for the average stock, informative for prices. $\gamma = 0$ appears consistent with the very muted reaction seen in the markets for the 2017 General election on what was in effect a shock result. For example the FTSE100 blue chip index opened the morning after the election at 7,450 which was exactly where it closed the night before on election day! Prior to this election there was political gridlock due to disagreements about Brexit. Although the opposition Labour party was seen as less friendly to business, power for them would have made the probability of remaining in the EU more likely. Results for the market were not

clear. However, by the time of the 2019 election, opposition policies had become either more extreme or more progressive, depending upon one's political viewpoint. Either way, in 2019 the implications for the UK stock prices from a Corbyn win were generally seen as being more negative than from the certainty of Brexit under the more generally perceived business friendly incumbents. The possibility that the outcome of the 2017 election was not informative for stock prices at all, and that $\lambda = 0$, is also suggested by the very small values of the t-statistics in the stock by stock regressions. This is -0.07 (or -0.13 for $K=1$) and $\sqrt{t^2} = 0.99 - 1.00$. Similarly for the 2020 presidential election, as the US stock market had been on a roar during Trump's presidency (contrary to 2016 fears), implications for stocks were not at all clear either in that election. Either result is less surprising than the failure of the 2016 presidential election to generate significance for the S&P500 constituents.

3.5.3.2 Weighted regression tests

The failure to demonstrate that betting markets significantly explained the returns of the S&P500 stock returns (and their Fama–French residuals) was a puzzling result for the 2016 US presidential election. This is a similar result though to the failure of the betting markets to explain the returns of the first market factor, MktRF, over the same period. However, significance was found when regressing MktRF over the shorter period of the month preceding the election, both for this election and the 2020 presidential election. Could it be that the betting market's explanatory power increases closer to an election? This would be consistent with the findings of Page and Clemen (2013). They demonstrated the performance of prediction markets is negatively correlated with time to expiry of the market. As the election nears it tends to dominate the news cycle and will more likely be on the minds of potential bettors. Betting volumes also increase hugely.

To explore this idea we apply a weighted regression scheme to our significance tests, with the weights increasing as the election nears. The J_γ tests of Pesaran and Yamagata allow this as long as the same weighting is used for each individual stock regression. The weighted estimates of correlation coefficients are also required when adjusting for cross sectional dependence of returns. This is trivial to apply. A weighting scheme will need to be chosen. It would be natural to weight by the volume traded in the betting markets. However, as Table 3.4 shows, the liquidity can increase by several orders of magnitude during the testing period, and peaks

Table 3.7: Results for weighted regressions, with the weight increasing by 5% each week.

Univariate regressions ^a									Individual stock regressions						
K	$\hat{\gamma}$	$\hat{\sigma}\gamma$	t	p	R^2	p_Q ^b	p_{LM}^{BP} ^c	p_{LM}^E ^d	$\bar{\gamma}$	\bar{t}	$\sqrt{\bar{t}^2}$	$\hat{J}_{\gamma,2}$	$P_{J_{\gamma,2}}$	p_Q ^e	$p_{LL\alpha}$ ^f
2014 Scottish independence referendum															
0	0.023	(0.030)	0.78	0.438	0.017	0.234	0.575	0.142	0.023	0.30	1.27	3.60***	<0.001	1.000	0.768
1	0.014	(0.017)	0.82	0.412	0.019	0.693	0.543	0.143	0.014	0.19	1.31	5.23***	<0.001	1.000	0.580
5	0.014	(0.016)	0.90	0.366	0.023	0.824	0.276	0.347	0.014	0.19	1.32	5.34***	<0.001	1.000	0.512
$N = 187, T = 37$															
2016 Brexit referendum - LSE stocks															
0	-0.170***	(0.029)	-5.93	<0.001	0.289	0.666	0.946	0.017	-0.170	-2.88	3.34	32.62***	<0.001	1.000	0.286
1	-0.030**	(0.015)	-1.98	0.048	0.044	0.805	0.600	0.199	-0.030	-0.45	1.85	15.49***	<0.001	1.000	0.026
5	-0.028*	(0.015)	-1.84	0.065	0.037	0.823	0.588	0.338	-0.028	-0.41	1.79	15.73***	<0.001	1.000	0.554
$N = 218, T = 85$															
2016 Brexit referendum - European exporters to UK															
0	-0.156***	(0.030)	-5.15	<0.001	0.232	0.906	0.773	0.119	-0.156	-2.53	2.89	24.00***	<0.001	1.000	0.139
1	0.014	(0.018)	0.79	0.430	0.005	0.199	0.251	0.457	0.014	0.27	1.32	4.30***	<0.001	1.000	0.106
5	0.019	(0.018)	1.02	0.309	0.010	0.119	0.341	0.550	0.019	0.32	1.35	4.64***	<0.001	1.000	0.403
$N = 101, T = 87$															
2016 US presidential election - S&P500 constituents															
0	-0.080**	(0.036)	-2.23	0.026	0.068	0.032	0.587	0.061	-0.080	-1.09	1.61	5.39***	<0.001	1.000	0.954
1	0.008	(0.007)	1.10	0.273	0.013	0.338	0.398	0.764	0.008	0.16	1.30	4.22***	<0.001	1.000	0.372
5	0.004	(0.006)	0.62	0.537	0.003	0.569	0.720	0.822	0.004	0.04	1.29	4.31***	<0.001	1.000	0.223
$N = 488, T = 74$															
2016 US presidential election - Mexican stocks															
0	-0.215***	(0.036)	-5.91	<0.001	0.320	0.569	0.311	0.821	-0.215	-2.74	3.06	24.71***	<0.001	0.501	0.736
1	-0.172***	(0.028)	-6.11	<0.001	0.333	0.966	0.751	0.813	-0.172	-2.30	2.69	23.65***	<0.001	0.499	0.843
5	-0.175***	(0.028)	-6.27	<0.001	0.345	0.897	0.601	0.893	-0.175	-2.34	2.73	24.82***	<0.001	0.441	0.736
$N = 50, T = 75$															
2017 UK general election - LSE stocks															
0	-0.012	(0.037)	-0.32	0.747	0.002	0.586	0.910	0.749	-0.012	-0.11	1.06	0.48	0.316	1.000	0.194
1	-0.009	(0.044)	-0.20	0.842	0.004	0.157	0.194	0.601	-0.009	-0.11	1.06	0.46	0.325	1.000	0.944
5	-0.007	(0.044)	-0.15	0.882	0.003	0.205	0.211	0.689	-0.007	-0.04	1.06	0.45	0.328	1.000	0.893
$N = 214, T = 34$															
2019 UK general election - LSE stocks															
0	0.016	(0.025)	0.65	0.518	0.011	0.679	0.527	0.874	0.016	0.41	1.28	3.67***	<0.001	1.000	0.748
1	0.014	(0.017)	0.81	0.418	0.019	0.784	0.338	0.042	0.014	0.39	1.37	6.70***	<0.001	1.000	0.336
5	0.015	(0.019)	0.78	0.433	0.019	0.970	0.414	0.045	0.015	0.40	1.47	9.17***	<0.001	1.000	0.979
$N = 242, T = 33$															
2020 US presidential election - S&P500 constituents															
0	0.200***	(0.090)	2.24	0.025	0.045	0.595	0.493	0.228	0.200	1.41	1.71	3.41***	<0.001	1.000	0.390
1	0.026	(0.046)	0.58	0.561	0.003	0.170	0.658	0.488	0.026	0.21	1.12	0.84	0.201	1.000	0.840
5	0.003	(0.019)	0.15	0.881	0.000	0.947	0.202	0.108	0.003	0.01	1.19	2.93***	0.002	1.000	0.321
$N = 487, T = 105$															

^a Equally weight portfolio regression

^b Ljung-Box serial correlation test p-value

^c Breusch-Pagan heteroskedasticity test p-value

^d Engle ARCH test p-value

^e Multivariate Ljung-Box serial correlation test p-value

^f Ling-Li multivariate heteroskedasticity test p-value

just before the election. Using raw volume will simply put all the weight on the final few days. Using the logarithm of volume could be a choice but probably does not increase sharply enough given the large increases in volume. Rather than come up with a scheme whereby we make some choice of function of volume to use, we will simply weight in the time dimension. We chose to increase the weighting exponentially by a relatively modest 5% per week. Regressions and significance tests are run for both individual stocks as well as the equally weighted portfolio. Results are shown in Table 3.7.

For the previously significant stock universe and event pairs, results are repeated but with increased significance. Absolute values of t-statistics and estimates of $J_{\gamma,2}$ are higher. Values of λ are also similar. We will not in general discuss these results individually. One meaningful change is with the European exporters portfolio during Brexit. Not only has the significance of the betting markets improved for the K=0 model (and become significant in the univariate test) but the K=1 and 5 models now produce significant results (at the 99.9% level). A small outperformance of these stocks versus the European average given Brexit over Remain of 1.4–1.9% is estimated. This is surprising as this author expected European exporters to the UK to underperform given a Brexit scenario. There may be other factors at play here, such as these companies exporting more outside Europe, potentially mitigating a downturn in EU revenues, than the average European company.

With regards to the S&P500 universe in the 2016 presidential election, results are now highly significant using the weighted regression. All $J_{\gamma,2}$ tests are significant at the 99.9% level and the K=0 univariate equally weighted portfolio test is significant at the 95% level. It appears that the information content of the betting markets may indeed improve as the election is approached. The K=0 models imply a fall in price of the average S&P500 stock of 8% given a win for Trump versus Biden. The K=1 and 5 models actually predict a small outperformance of 0.4–0.8% of the S&P500 versus the North American market as a whole. This may be highly significant but it is an economically small effect.

The 2017 UK general election remains insignificant using the weighted regressions. Again very small values of t-statistics are found in the stock by stock regressions. \bar{t} is in the range -0.04 to -0.11 and $\sqrt{\bar{t}^2} = 1.06$. We conclude that our model is not relevant for this event and indeed $\lambda = 0$.

Finally we see that using a weighted regression generates some significance for S&P500 stocks in the 2020 US presidential election. For K=0 the results are significant at the 99.9% level for the stock by stock test and at the 95% level for

the equally weighted portfolio regression. The typical index stock is expected to outperform by 20% given a Biden win over a Trump win. The $K=5$ model also generates a significant $J_{\gamma,2}$ at the 99% level. However, the predicted outperformance of the index stocks of 0.3% versus the wider North American market is economically insignificant. We note that using the weighted scheme has generated significance for the S&P500 portfolios for both US elections studied in this chapter. Betting markets do explain the moves in outright stock returns. However, despite the betting markets having significant explanatory effects on the residuals of the Fama–French model for these index stocks the predicted effects are very small and economically insignificant.

In general, using significance tests based on the weighted regressions has sharpened our results. Events that were deemed to be significant using the standard regression over the whole testing period remain significant but more so. We have also demonstrated that the betting markets explain the returns of S&P500 stocks using the weighted scheme. This suggests that the information content increases as the election nears. This is a not unexpected result given the explosion in trading volumes on betting exchanges as the election is approached. However, we do not generate results of economically meaningful magnitude for the Fama–French residuals of the S&P500 index stocks. Finally our model does not appear to hold in that $\gamma = 0$ for the 2017 UK general election. The results of that election did not seem clear for Brexit and hence stock prices given the particular political situation in the UK at that time.

3.5.3.3 Political factor loading characteristics

Next we turn to an investigation of the political factor loadings γ and how they vary with any common observable characteristics of the stocks. We will examine how γ varies for the 5 factor model. This will identify the political sensitivity of individual stocks, controlled for common characteristics related to the Fama–French factors. Before we begin we need to caveat our results. We regress the loadings against some observable characteristics. However, not all stocks had the necessary data we sought. As such we cannot discount the possibility of some sample selection bias in what we report. However, most events had reasonably good data coverage and the relationships uncovered were generally strong. We are confident of the direction of the relationships revealed if not their exact magnitude.

Table 3.8: Regressions of factor loadings γ against stock characteristics for different events.

Dependent Variable	Estimate	Error	tStat	pValue
γ				
2014 Scottish independence referendum				
Intercept	-0.030**	0.012	-2.56	0.011
% 2013 UK revenue	0.080***	0.018	4.54	<0.001
$I(\text{HQ in Scotland})$	0.073**	0.036	2.03	0.044
$\overline{N = 158, R^2 = 0.136}$				
2016 Brexit referendum - UK Stocks				
Intercept	0.036**	0.014	2.50	0.013
% 2015 developed Europe revenue	-0.098***	0.021	-4.78	<0.001
$\overline{N = 183, R^2 = 0.112}$				
2016 Brexit referendum - EU Exporters				
Intercept	-0.000	0.023	0.00	0.999
% 2015 UK revenue	0.044	0.050	0.87	0.385
$\overline{N = 101, R^2 = 0.008}$				
2016 US presidential election - S&P500				
Intercept	0.058***	0.013	4.50	<0.001
% 2015 US revenue	-0.075***	0.017	-4.41	<0.001
$\overline{N = 461, R^2 = 0.041}$				
2016 US presidential election - Mexican stocks				
Intercept	-0.094	0.063	-1.50	0.156
% 2015 US revenue	-0.199	0.174	-1.15	0.271
$\overline{N = 16, R^2 = 0.086}$				
2019 UK general election				
Intercept	-0.041*	0.021	-1.97	0.051
% 2016 UK revenue	0.111***	0.033	3.41	<0.001
$I(\text{Corbyn to nationalise})$	0.004	0.050	0.07	0.944
$\overline{N = 202, R^2 = 0.071}$				
2020 US presidential election - S&P500				
Intercept	-0.030	0.019	-1.56	0.120
% 2015 US revenue	0.033	0.025	1.32	0.187
$\overline{N = 404, R^2 = 0.002}$				

For the 188 stocks listed on the LSE that survived our selection process for the 2014 Scottish independence referendum, 158 of them published easy to access percentages of revenues for the UK in their 2013 full year accounts. A greater sensitivity to independence is likely for stocks that generate more of their revenues on shore as the UK economy would likely suffer more than international peers. thirteen of the companies were also headquartered in Scotland. These latter stocks would be much more likely to affected by independence than companies based in the rest of the UK.

Table 3.8 shows the results of a regression of γ against these two stock characteristics. Results for similar regressions from the other events are also shown in the table. Both characteristics coefficients are significant. For every additional 1% revenue earned onshore in 2013, γ would be 0.0008 higher. This means the sensitivity of the stock price to the election increases 8 basis points (0.08%) from that 1% additional UK based revenue. By sensitivity here we mean the expected difference of prices between Scotland leaving or remaining as part of the UK. The coefficient of the ‘HQ in Scotland’ indicator implies an increased sensitivity due to being based in Scotland of 7.3%. Darby et al. (2019) demonstrated that sensitivity to betting odds helped predict cross sectional returns of Scottish based companies. The work in this chapter goes further and shows that Scottish headquartered companies had significantly more risk to the 2014 Scottish independence referendum than other UK firms.

For stocks listed on the LSE in the Brexit referendum, we use percentage of revenue generated in developed European markets published in the 2015 full year accounts. This includes the revenue earned in the UK. This developed EU markets figure is found to be very similar to the revenue figure for the UK only. Most of the revenue earned by UK companies in developed European markets is earned in the UK. Joint regression of γ on both these figures is problematic due to co-linearity. We chose the European number to include any effects on exporters to the EU (who are likely to be also affected by Brexit). These figures are available for 183 of the 219 stocks used in the regressions. European revenue is significantly explanatory for γ at the 99.9% level. The political sensitivity increases by about 10 basis points for every additional percent of developed European revenue earned. Put another way, for every additional percent of revenue earned outside of developed European markets, stocks are expected to have prices 8 basis points higher in price under Brexit (versus Remain). We also regress γ on revenue earned in the UK in 2015 for the EU exporters portfolio (this was available for all the exporters). No relationship

was found. The UK sales figure was not informative with a very low R^2 of 0.008 reported for the regression.

For the 2016 election where Trump won the presidency, we study both the S&P500 index stocks as well as Mexican companies. For the index stocks we source percentage of sales earned locally in the US in 2015 (we hypothesise that exporters would have more sensitivity to a Trump win given the potential for increased trade disruption). The characteristic was available for 461 of the 488 index stocks that were used in the significance tests. The local percentage of revenue figure was highly significant with γ falling 7.5 basis points for every additional percent of revenue earned locally. This is puzzling, as this means γ was increasing with increasing exports and that exporters would be likely to perform better under a Trump versus a Clinton presidency. We have no explanation for this result but do note that the estimated values of γ for this model had on average an expectantly positive sign, although the effect was very small ($\bar{\gamma} = 0.004$).

For the Mexican portfolio we source the percentage of revenue earned in the US in 2015. Unfortunately this was only reported for 16 of the 51 stocks used. The regression does not generate significant results, not unexpectedly since N is so low. The coefficient is estimated at -0.199 ± 0.174 ($p = 0.271$). This suggests a Mexican stock that earns all revenue from the US is estimated to have prices 20% lower under Trump versus Clinton than a Mexican company that earns no revenue in the US.

For the 2017 UK General election we do not proceed as we have concluded our model does not hold ($\gamma = 0$). For the 2019 UK General election portfolio we use the 2015 reported percentage revenue earned in the UK. This is available for 202 of the 243 stocks considered. Jeremy Corbyn, the opposition Labour party leader, also promised to nationalise 17 UK based companies, if he won, Rees (2019). We regress γ on the UK revenue percentage and an indicator based on whether or not the company would be nationalised under Labour. UK Revenue is significant at the 99.9% level. Political sensitivity increases with UK revenue, 11 basis points for every additional percentage point revenue earned onshore in 2018. This means UK focused businesses were expected to perform relatively poorly had the Conservatives failed to get a majority, relative to exporters. The coefficient on the indicator of nationalisation risk under Corbyn is not significant, but is estimated at 3.9%. We note that the average γ for stocks identified as being nationalised under Corbyn is 0.055 ± 0.12 versus 0.012 ± 0.17 for the other stocks. We thus test equivalence of the distributions of γ amongst stocks conditional on nationalisation

risk versus conditional on no nationalisation risk. We apply a one-sided two-sample Kolmogorov–Smirnov test. The null is that the two samples come from the same distribution versus an alternative that the CDF for the sample given nationalisation risk is higher. The Kolmogorov–Smirnov test statistic is 0.325 with a p-value of 0.028. The 17 stocks Corbyn named as nationalisation targets are significantly more sensitive to the result of the election than others at the 95% level.

Finally we seek a relationship for the S&P500 index stocks for the 2020 US presidential election. 2019 local percentage revenue is available for 404 of the 487 stocks used in the model estimation stage. This is not explanatory for γ with a very low R^2 reported of 0.002 for the regression.

Elections often involve domestic or geographical risks. It is likely that companies that earn more revenue offshore or internationally are more able to diversify away from election risks. Some evidence for this hypothesis was present in the literature for Brexit, Hill et al. (2019). Internalisation was found to decrease the sensitivity of UK stock prices to the betting odds for Brexit as well as the day-after-result return. In this chapter we have extended these findings, testing this idea for all six events. We included a proxy for national revenue (or developed EU revenue for Brexit), finding that internalisation is indeed significantly explanatory for four of the six events. These results both complement Hill et al. (2019) and improve them. Not only do we consider more events, we check the explanatory power after the five Fama–French factors have been controlled for.

Similarly we have repeated an observation made in one of the oldest studies of elections risk and financial markets (Gemmell, 1992). For the 1987 UK general election this paper showed there was a much larger increase as the election was approached in option volatility of stocks that were at risk of nationalisation under the Labour opposition. History appears to have repeated itself in 2019. The then Labour leader Jeremy Corbyn threatened to nationalise various utility and other companies if he won. Here we measure political risk as the sensitivity to betting odds of the Conservatives failing to get a majority, finding the risk significantly larger for companies at risk of nationalisation.

3.6 Conclusion

The information content of political prediction markets is well documented. As is the effect on asset prices of elections and other political risks. Given this, it is natural to ask is there a way to formally link prices between political markets

and financial markets? The previous chapter sought to answer this question in the very particular circumstances of the overnight session following a vote. In this chapter we seek to describe the relationship between these two types of markets in the weeks and months leading up to a political event. There are a small number of examples in the literature that study this question (see Manasse et al. (2020), Hanna et al. (2021), Acker and Duck (2015) and Darby et al. (2019)). However, they consider only a single event and typically study only empirical relationships. We build a pricing model from the ‘ground-up’ using economic assumptions and common pricing restrictions taken from asset pricing literature. That we recover a relationship between asset price returns and the first difference of betting markets is perhaps no surprise. Indeed other papers have studied this relationship empirically. However, this work differs in that we have a theoretical basis for the model, and also test it on multiple events. Indeed the main contribution of this chapter is to present a model that is grounded in an economic assumption with common pricing restrictions and can be applied to *any* political event¹⁸.

To build the model we make a key modification to the assumptions of the cointegrating model of the preceding section. This is that the *difference* of the conditional expectations of asset prices (given the result of the election) is fixed, as opposed to the conditional expectations themselves being fixed. This leads to the relationship in first differences. The variance of the returns of asset prices are separated into a political part, explained by political markets, and a residue, related to commercial, economic and other non-political factors. In fact the cointegration model of the preceding chapter is a special case of this model, where this latter residue has zero variance. The model is naturally extended using the Fama–French factors to describe the variance of the non-political residue. The resulting model is an extended characteristic factor model, where all factors, both Fama–French and political, via betting markets, are observed.

We test the model on six recent elections. We find strong evidence in favour of our model for four of the six events. One election has mixed results (the 2020 presidential election). A weighting scheme was required there to generate significance. Data closer to the election was weighted more highly than data far from the vote. Although a modification of the original model, this can be explained by political markets having greater information content the closer we are to an election. Some justification for the approach is from Page and Clemen (2013) which

¹⁸The model applies in more general settings. It can be used for events with a prediction market and whose outcome has an effect on asset prices.

showed that the forecasting ability of prediction markets is negatively correlated with time to expiry. Further, betting volumes also increase exponentially as the election approaches. We find the weighting scheme sharpens the results of the four other events found to be significant with the unweighted scheme. Significance of the political factor loading, γ , increased under the weighted approach. This is consistent with the huge increases in volumes observed on betting platforms as the election drew near. We find no support for our model for a single event, the 2017 UK general election. We conclude that this election is not informative for asset prices. In the model we believe the weighting on the political factor, γ , is in fact zero.

An exploration of the factor loadings reveals some pleasing relationships. Consistent with Hill et al. (2019) we find that domestic (or EU based) revenue is a strong explanatory factor of political risk. Indeed, this characteristic was significant for four of the six elections studied. This provides evidence for the hypothesis that companies with a greater reliance on international sales are more able to diversify domestic political risk. We find that the location of company headquarters, and nationalisation risk under a given outcome, are also found to increase γ markedly.

Concluding remarks

This thesis studied the behaviour of financial and prediction markets around political events. We believe there are two main contributions of the work. Firstly, we demonstrated without a shadow of a doubt that there was irrational behaviour by market participants on the night of the EU referendum. This extends the hypothesis that there was a ‘bubble’ in public opinion for Remain before the vote. Secondly, we presented pricing models that link the prices of financial and betting assets that apply in a general setting and also have strong support in the data. Although there are some studies in the literature studying this problem empirically for single events, we believe the work presented in the final two chapters is without peer. A further finding, that prediction markets lead financial on small timescales as election results unfold is also a valuable contribution. This adds to the weight of evidence in the literature showing that prediction markets have excellent forecasting ability.

It is appropriate to highlight potential areas of further research. Four are identified.

Firstly, the application of the model in chapter 3, that demonstrates a relationship between financial returns and changes in prediction market binary option prices, could be applied to asset classes beyond equities. We suggest investigating the relationship for leading indices for other asset classes (commodities, fixed income, currencies and cryptocurrencies) would be valuable.

Chapters 1 and 2 identify the fact that betting markets lead financial markets on the order of minutes to tens of minutes during the overnight session following a vote. This deviation from efficiency is not found in chapter 3. This is not surprising as we only consider much longer daily returns. Our second area of future research is to apply the model of chapter 3 to higher frequencies. This could answer the question of whether or not this lead–lag relationship occurs in the period before the vote, and not just on the night after an event.

In this thesis we have only considered models of prediction market binary options. The third area of suggested further research is to extend the analysis to betting market contracts-for-differences. Such contracts may settle against, for example, vote share or numbers of seats in a legislature, for a given party. Extensions of the model should be achievable by considering the conditional expectation of asset prices given the settlement of the underlying variable.

Our models can be applied to any prediction market binary option, not just political ones. The condition is that the underlying event must have an effect on financial prices. Prediction markets exist not just for elections. Public markets exist for other events, for instance economic data releases and geopolitical risks. Our final area of future research is to apply our models to other types of events. A particularly poignant current example would be to study the relationship between betting markets related to the outcome of the Ukraine war and financial prices.

Bibliography

- Acker, D., Duck, N. W. 2015. Political risk, investor attention and the scottish independence referendum. *Finance Research Letters*. 13, 163–171.
- Alkhatib, A., Harasheh, M. 2018. Performance of Exchange Traded Funds during the Brexit Referendum: An Event Study. *IJFS*. 6 (3), 1–12.
- Anderson, C. W., Garcia-Feijóo, L. 2006. Empirical evidence on capital investment, growth options, and security returns. *The Journal of Finance*. 61 (1), 171–194.
- Aristeidis, S., Elias, K. 2018. Empirical analysis of market reactions to the UK's referendum results - How strong will Brexit be? *Journal of International Financial Markets, Institutions and Money*. 53, 263–286.
- Auld, T., Linton, O. 2019. The behaviour of betting and currency markets on the night of the EU referendum. *International Journal of Forecasting*. 35 (1), 371–389.
- Banz, R. W. 1981. The relationship between return and market value of common stocks. *Journal of Financial Economics*. 9 (1), 3–18.
- Baxter, M. 2017. Electoral Calculus; Historical Data and Plots. www.electoralcalculus.co.uk.
- Beare, B. K. 2018. Unit Root Testing with Unstable Volatility. *Economics Papers 2008-W06*, Economics Group, Nuffield College, University of Oxford.

- Beaulieu, M.-C., Dufour, J.-M., Khalaf, L. 2007. Multivariate Tests of MeanVariance Efficiency With Possibly Non-Gaussian Errors: An Exact Simulation-Based Approach. *Journal of Business & Economic Statistics*. 25, 398–410.
- Berg, J., Rietz, T. 2006. The Iowa electronic markets: Stylized facts and open issues, in *information markets: A new way of making decisions*. AEI-Brookings Joint Center for Regulatory Studies. 142–169.
- Bernhard, W., Leblang, D. 2006. *Democratic processes and financial markets: Pricing politics*. Cambridge University Press.
- Białkowski, J., Gottschalk, K., Wisniewski, T. P. 2008. Stock market volatility around national elections. *Journal of Banking & Finance*. 32 (9), 1941–1953.
- Black, F., Jensen, M. C., Scholes, M. 1972. The capital asset pricing model: Some empirical tests. In *Studies in the Theory of Capital Markets*, New York: Praeger.
- Brander, J. A. 1991. Election polls, free trade, and the stock market: Evidence from the 1988 Canadian general election. *The Canadian Journal of Economics / Revue canadienne d'Economique*. 24 (4), 827–843.
- Brodkin, J. 2008. Google bets on value of prediction markets. Internet, <https://www.networkworld.com/article/2284098/google-bets-on-value-of-prediction-markets.html>, 5 Mar 2008.
- Caballero, R. J., Simsek, A. 2016. A model of fickle capital flows and retrenchment. Working Paper 22751, National Bureau of Economic Research.
- Caporale, G. M., Gil-Alana, L., Trani, T. 2018. Brexit and Uncertainty in Financial Markets. *IJFS*. 6 (1), 1–9.
- Carr, P. 2017. Bounded brownian motion. *Risks*. 5 (4), 61.
- Cascione, S. 2016. Mexico peso seen plunging 10 percent if Trump wins U.S. election-Reuters poll. Reuters, 1 Nov 2016.

- Cavaliere, G., Taylor, A. M. R. 2008. Bootstrap unit root tests for time series with nonstationary volatility. *Econometric Theory*. 24 (1), 43–71.
- Cavaliere, G., Taylor, A. M. R. 2009. Heteroskedastic time series with a unit root. *Econometric Theory*. 25 (5), 1228–1276.
- Cavaliere, G., Taylor, R. 2007. Testing for unit roots in time series models with non-stationary volatility. *Journal of Econometrics*. 140 (2), 919–947.
- Cavaliere, G., Taylor, R. 2008. Time-transformed unit root tests for models with non-stationary volatility. *Journal of Time Series Analysis*. 29 (2), 300–330.
- Choi, I., Saikkonen, P. 2010. Tests for nonlinear cointegration. *Econometric Theory*. 26 (3), 682–709.
- Clark, I. J., Amen, S. 2017. Implied distributions from GBPUSD risk-reversals and implication for brexit scenarios. *Risks*. 5 (3), 1–17.
- Clark, I. J., Amen, S. 2019. Implied distributions from risk-reversals and Brexit/Trump predictions. In *Innovations in Insurance, Risk-and Asset Management*, 117–134, World Scientific.
- Clinet, S., Potiron, Y. 2019. Cointegration in high frequency data. *Electronic Journal of Statistics*.
- Curtice, J. 2014. Scottish independence: “Invaluable second chance” for Alex Salmond. BBC News Website, <https://www.bbc.co.uk/news/uk-scotland-scotland-politics-28921754>, 25 Aug 2014.
- Curtice, J., Fisher, S. D., Kuha, J. 2011. Confounding the commentators: how the 2010 exit poll got it (more or less) right. *Journal of Elections, Public Opinion and Parties*. 21 (2), 211–235.
- Dahlgreen, W. 2014. ‘Yes’ campaign lead at 2 in Scottish Referendum. YouGov, <https://yougov.co.uk/topics/politics/articles-reports/2014/09/06/latest-scottish-referendum-poll-yes-lead>, 6 Sep 2014.

- Darby, J., Gao, J., Lucey, S., Zhu, S. 2019. Is heightened political uncertainty priced in stock returns? Evidence from the 2014 Scottish independence referendum. University of Strathclyde.
- DellaVigna, S., Pollet, J. M. 2009. Investor inattention and Friday earnings announcements. *The Journal of Finance*. 64 (2), 709–749.
- Diepenbrock, G. 2016. Like polls, prediction markets failed to see Trump’s victory coming, economist says. University of Kansas, <https://phys.org/news/2016-11-polls-trump-victory-economist.html>, 11 Nov 2016.
- Ding, P. 2016. On the conditional distribution of the multivariate t distribution. *The American Statistician*. 70 (3), 293–295.
- Econoday. 2016. 2016 Economic Calendar. Econoday economic calendar, <http://www.econoday.com>, 23 Jun 2016.
- Economist. 2016. Who said Brexit was a surprise?. Internet, <https://www.economist.com/graphic-detail/2016/06/24/who-said-brexit-was-a-surprise>, 24 Jun 2016.
- Engle, R. F., Granger, C. W. J. 1987. Co-integration and error correction: Representation, estimation, and testing. *Econometrica*. 55 (2), 251–276.
- Escanciano, J.-C., Escribano, A. 2009. *Econometrics: Non-linear Cointegration*. 2757–2769, New York, NY, Springer New York.
- Fama, E. F. 1965. The behavior of stock-market prices. *The Journal of Business*. 38 (1), 34–105.
- Fama, E. F., French, K. R. 1992. The cross-section of expected stock returns. *The Journal of Finance*. 47 (2), 427–465.
- Fama, E. F., French, K. R. 1993. Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics*. 33 (1), 3–56.
- Fama, E. F., French, K. R. 2004. The capital asset pricing model: Theory and evidence. *Journal of Economic Perspectives*. 18 (3), 25–46.

- Fama, E. F., French, K. R. 2015. A five-factor asset pricing model. *Journal of Financial Economics*. 116 (1), 1–22.
- Fan, J., Liao, Y., Yao, J. 2015. Power enhancement in high-dimensional cross-sectional tests. *Econometrica*. 83 (4), 1497–1541.
- Forsythe, R., Frank, M., Krishnamurthy, V., Ross, T. W. 1995. Using market prices to predict election results: The 1993 UBC election stock market. *The Canadian Journal of Economics / Revue canadienne d'Economie*. 28 (4a), 770–793.
- Forsythe, R., Nelson, F., Neumann, G. R., Wright, J. 1992. Anatomy of an Experimental Political Stock Market. *American Economic Review*. 82 (5), 1142–1161.
- Foster, A. 2016. Brexit: How important will turnout be to the EU referendum result?. *The Express*, 23 Jun 2016.
- Fry, J., Brint, A. 2017. Bubbles, blind-spots and brexit. *Risks*. 5 (3), 37.
- GabrielaTitan, A. 2015. The efficient market hypothesis: Review of specialized literature and empirical research. *Procedia Economics and Finance*. 32, 442–499.
- Gao, J., King, M., Lu, Z., Tjøstheim, D. 2009. Nonparametric specification testing for nonlinear time series with nonstationarity. *Econometric Theory*. 25 (6), 1869–1892.
- Gelman, A., Rothschild, D. 2016. Something's Odd About the Political Betting Markets. *PredictWise blog*, <https://blog.predictwise.com/somethings-odd-about-the-political-betting-markets>, 12 Jul 2016.
- Gemmell, G. 1992. Political risk and market efficiency: Tests based in British stock and options markets in the 1987 election. *Journal of Banking & Finance*. 16 (1), 211–231, Special Issue on European Capital Markets.
- Gibbons, M. R., Ross, S., Shanken, J. 1989. A test of the efficiency of a given portfolio. *Econometrica*. 57 (5), 1121–52.

- Gjerstad, S. 2004. Risk Aversion, Beliefs, and Prediction Market Equilibrium. *Microeconomics* 0411002, University Library of Munich, Germany.
- Goodell, J. W., Bodey, R. A. 2012. Price-earnings changes during US presidential election cycles: voter uncertainty and other determinants. *Public Choice*. 150 (3/4), 633–650.
- Goodell, J. W., McGroarty, F., Urquhart, A. 2015. Political uncertainty and the 2012 US presidential election: A cointegration study of prediction markets, polls and a stand-out expert. *International Review of Financial Analysis*. 42, 162–171.
- Goodell, J. W., Vähämaa, S. 2013. US presidential elections and implied volatility: The role of political uncertainty. *Journal of Banking & Finance*. 37 (3), 1108–1117.
- Graefe, A. 2016. Political markets. In *SAGE Handbook of Electoral Behavior*, SAGE.
- Granger, C. 1983. Co-integrated variables and error correction models. UCSD, Discussion Paper.
- Gregory, A., Hansen, B. 1996. Residual-based tests for cointegration in models with regime shifts. *Journal of Econometrics*. 70 (1), 99–126.
- Gungor, S., Luger, R. 2009. Exact distribution-free tests of mean-variance efficiency. *Journal of Empirical Finance*. 16 (5), 816–829.
- Gungor, S., Luger, R. 2013. Testing linear factor pricing models with large cross sections: A distribution-free approach. *Journal of Business & Economic Statistics*. 31 (1), 66–77.
- Gutteridge, N. 2016. Working class revolution? Reports of huge EU referendum turnout which “would favour Leave”. The Express website, <https://www.express.co.uk/news/politics/682890/EU-referendum-Reports-huge-EU-referendum-turnout-favour-Leave-Brexit>, 24 Jun 2016.

- Hanna, J., Hansen, M. N.-J. H., MacDonald, M. M. 2021. The Impact of Political Uncertainty on Asset Prices: The Case of the United Kingdom's EU Membership Referendum. International Monetary Fund.
- Hanretty, C. 2016. The EU referendum: what to expect on the night. Blog, <https://medium.com/@chrishanretty/the-eu-referendum-what-to-expect-on-the-night-521792dd3eef>, blog.
- Hanretty, C., Lauderdale, B., Vivyan, N. 2016. Comparing strategies for estimating constituency opinion from national survey samples. *Political Science Research and Methods*. 1–21.
- Hansen, B., Phillips, P. 1990. Statistical inference in instrumental variable regression with $I(1)$ processes. *Review of Economic Studies*. 57, 99–125.
- Hanson, R. D., Caplan, B., du Pont Jarmas, R., Lott, J. R., Masse, C., Shaw, U. 2005. The Informed Press Favored the Policy Analysis Market. <http://mason.gmu.edu/rhanson/PAMpress.pdf>.
- Hanson, R., Oprea, R., Porter, D. 2006. Information aggregation and manipulation in an experimental market. *Journal of Economic Behavior & Organization*. 60 (4), 449–459.
- Hatemi-J, A. 2008. Tests for cointegration with two unknown regime shifts with an application to financial market integration. *Empirical Economics*. 35 (3), 497–505.
- Hill, P., Korczak, A., Korczak, P. 2019. Political uncertainty exposure of individual companies: The case of the brexit referendum. *Journal of Banking & Finance*. 100, 58–76.
- Hirshleifer, D., Hsu, P.-H., Li, D. 2013. Innovative efficiency and stock returns. *Journal of Financial Economics*. 107 (3), 632–654.
- Hirshleifer, D., Lim, S. S., Teoh, S. H. 2009. Driven to distraction: Extraneous events and underreaction to earnings news. *The Journal of Finance*. 64 (5), 2289–2325.

- Horn, C. F., Ivens, B. S., Ohneberg, M., Brem, A. 2014. Prediction markets: A literature review 2014. *The Journal of Prediction Markets*. 8 (2), 89–126.
- Hou, K., Xiong, W., Peng, L. 2009. A tale of two anomalies: The implications of investor attention for price and earnings momentum.
- Huang, J. Y., Shieh, J. C., Kao, Y.-C. 2016. Starting points for a new researcher in behavioral finance. *International Journal of Managerial Finance*. 12 (1), 92–103.
- Jensen, M. C. 1968. The performance of mutual funds in the period 1945–1964. *The Journal of Finance*. 23 (2), 389–416.
- Johansen, S. 1995. *Likelihood-Based Inference in Cointegrated Vector Autoregressive Models*. Oxford University Press.
- Johansen, S., Juselius, K. 1990. Maximum likelihood estimation and inference on cointegration – with applications to the demand for money. *Oxford Bulletin of Economics and Statistics*. 52 (2), 169–210.
- Joyce E. Berg, F. D. N., Rietz, T. A. 2008. Prediction market accuracy in the long run. *International Journal of Forecasting*. 24 (2), 285–300, US Presidential Election Forecasting.
- Karlsen, H. A., Myklebust, T., Tjøstheim, D. 2007. Nonparametric estimation in a nonlinear cointegration type model. *The Annals of Statistics*. 35 (1), 252–299.
- Keim, D. B. 1983. Size-related anomalies and stock return seasonality: Further empirical evidence. *Journal of Financial Economics*. 12 (1), 13–32.
- Kelly, B., Pastor, L., Veronesi, P. 2016. The price of political uncertainty: Theory and evidence from the option market. *The Journal of Finance*. 71 (5), 2417–2480.
- Krauss, C., Herrmann, K. 2017. On the power and size properties of cointegration tests in the light of high-frequency stylized facts. *Journal of Risk and Financial Management*. 10 (1).

- Ling, S., Li, W. K. 1997. Diagnostic checking of nonlinear multivariate time series with multivariate arch errors. *Journal of Time Series Analysis*. 18 (5), 447–464.
- Maki, D. 2012. Tests for cointegration allowing for an unknown number of breaks. *Economic Modelling*. 29 (5), 2011–2015.
- Malkiel, B. G. 2003. The efficient market hypothesis and its critics. *The Journal of Economic Perspectives*. 17 (1), 59–82.
- Manasse, P., Moramarco, G., Trigilia, G. 2020. Exchange rates and political uncertainty: The brexit case. *SSRN Electronic Journal*.
- Mann, A. 2016. The power of prediction markets. *Nature*. 538 (7625), 308–310.
- Newey, W. K., West, K. D. 1994. Automatic Lag Selection in Covariance Matrix Estimation. *The Review of Economic Studies*. 61 (4), 631–653.
- Noh, J., Kim, T.-H. 2003. Behaviour of cointegration tests in the presence of structural breaks in variance. *Applied Economics Letters*. 10 (15), 999–1002.
- Novy-Marx, R. 2013. The other side of value: The gross profitability premium. *Journal of Financial Economics*. 108 (1), 1–28.
- Nyblom, J., Harvey, A. 2000. Tests of common stochastic trends. *Econometric Theory*. 16 (2), 176–199.
- Oliven, T. A., K. & Rietz. 2004. Suckers are born but markets are made: Individual rationality, arbitrage, and market efficiency on an electronic futures market. *Management Science*. 50 (3), 336–351.
- Page, L., Clemen, R. T. 2013. Do prediction markets produce well-calibrated probability forecasts? *The Economic Journal*. 123 (568), 491–513.
- Pantzalis, C., Stangeland, D. A., Turtle, H. J. 2000. Political elections and the resolution of uncertainty: The international evidence. *Journal of Banking & Finance*. 24 (10), 1575–1604.

- Pesaran, M., Yamagata, T. 2012. Testing CAPM with a large number of assets. IZA Discussion Papers 6469, Institute of Labor Economics (IZA).
- Peters, W. 2016. Forecasts 2016: JPMorgan See GBP/USD Bottoming at 1.32. Internet, <https://www.poundsterlinglive.com/exchange-rate-forecasts/3404-jpmorgan-chase-exchange-rate-forecast-summary>, 25 Apr 2016.
- Phillips, P. 1991. Optimal inference in cointegrated systems. *Econometrica*. 59 (2), 283–306.
- Phillips, P., Ouliaris, S. 1990. Asymptotic properties of residual based tests for cointegration. *Econometrica*. 58 (1), 165–93.
- Pickard, J. 2019. Corbyn launches Labour’s most leftwing manifesto in generation. *The Irish Times*, 21 Nov 2019.
- Polgreen, P. M., Nelson, F. D., Neumann, G. R., Weinstein, R. A. 2007. Use of Prediction Markets to Forecast Infectious Disease Activity. *Clinical Infectious Diseases*. 44 (2), 272–279.
- Rallings, C., Johnston, R., Thrasher, M. 2008. Changing the boundaries but keeping the disproportionality: The electoral impact of the fifth periodical reviews by the Parliamentary Boundary Commissions for England and Wales. *The Political Quarterly*. 79 (1), 80–90.
- Rees, T. 2019. The 17 stocks trading at a 10pc discount because of their exposure to a Jeremy Corbyn-led government. *The Telegraph*, 29 Jun 2019.
- Reyes, B. 2016a. Gibraltar will vote to remain in EU – poll. *Gibraltar Chronicle*, 15 Apr 2016.
- Reyes, B. 2016b. Support for Remain strengthens in Gibraltar – poll. *Gibraltar Chronicle*, 16 Jun 2016.
- Rhode, P. W., Strumpf, K. S. 2004. Historical presidential betting markets. *Journal of Economic Perspectives*. 18 (2), 127–141.

- Rosenberg, B., Reid, K., Lanstein, R. 1985. Persuasive evidence of market inefficiency. *The Journal of Portfolio Management*. 11 (3), 9–16.
- Rothschild, D., Sethi, R. 2016. Trading strategies and market microstructure: Evidence from a prediction market. *The Journal of Prediction Markets*. 10 (1).
- Schwert, G. W. 1989. Tests for unit roots: A Monte Carlo investigation. *Journal of Business & Economic Statistics*. 7 (2), 147–159.
- Shapiro, Y. 2016. EU Referendum: Growing support in Northern Ireland for Brexit but balance of power lies with the undecided. *Belfast Telegraph*, 20 Jun 2016.
- Shin, Y. 1994. A residual-based test of the null of cointegration against the alternative of no cointegration. *Econometric Theory*. 10 (1), 91–115.
- Silvapulle, P. S., Podivinsky, J. M. 2000. The effect of non-normal disturbances and conditional heteroskedasticity on multiple cointegration tests. *Journal of Statistical Computation and Simulation*. 65 (1-4), 173–189.
- Smales, L. A. 2016. The role of political uncertainty in Australian financial markets. *Accounting & Finance*. 56 (2), 545–575.
- Stock, J. 1987. Asymptotic properties of least squares estimators of cointegrating vectors. *Econometrica*. 55 (5), 1035–56.
- Sultonov, M., Jehan, S. N. 2018. Dynamic linkages between Japan’s foreign exchange and stock markets: Response to the Brexit referendum and the 2016 U.S. presidential election. *Journal of Risk and Financial Management*. 11 (3), 34.
- Taleb, N. N. 2018. Election predictions as martingales: an arbitrage approach. *Quantitative Finance*. 18 (1), 1–5.
- Titman, S., Wei, K. C. J., Xie, F. 2004. Capital investments and stock returns. *The Journal of Financial and Quantitative Analysis*. 39 (4), 677–700.

- Twyman, J. 2016. Analysis: the impact of turnout on the EU Referendum. YouGov, <https://yougov.co.uk/topics/politics/articles-reports/2016/05/21/assessing-impact-turnout-eu-referendum>, 21 May 2021.
- Wagner, M., Hong, S. H. 2016. Cointegrating polynomial regressions: Fully modified OLS estimation and inference. *Econometric Theory*. 32 (5), 1289–1315.
- Wall, M., Costello, R., Lindsay, S. 2017. The miracle of the markets: Identifying key campaign events in the Scottish independence referendum using betting odds. *Electoral Studies*. 46, 39–47.
- Wang, Q. 2015. Limit theorems for nonlinear cointegrating regression. World Scientific.
- Wang, Q., Wang, Y. X. R. 2013. Nonparametric Cointegrating Regression With Nnh Errors. *Econometric Theory*. 29 (1), 1–27.
- Wearden, G. 2015. Scottish independence poll hits shares and sends pound sliding. The Guardian Website, 8 Sep 2014.
- Wikipedia. 2016. Opinion Polling for the United Kingdom European Union membership referendum. www.wikipedia.org.
- Wikipedia. 2017. Opinion polling for the 2017 United Kingdom general election. www.wikipedia.org.
- Wikipedia. 2019. Opinion polling for the 2019 United Kingdom general election. www.wikipedia.org.
- Wolfers, J., Zitzewitz, E. 2006. Interpreting prediction market prices as probabilities. Working Paper 12200, National Bureau of Economic Research.
- Wu, K., Wheatley, S., Sornette, D. 2017. The British Pound on Brexit Night: A Natural Experiment of Market Efficiency and Real-Time Predictability. Swiss Finance Institute Paper no. 17-12.
- YouGov. 2016. YouGov EU Referendum On the Day Poll. <https://yougov.co.uk/topics/politics/articles-reports/2016/06/23/yougov-day-poll>, 23 Jun 2016.

Appendix A: Election forecast model summary

A.1 Variables

$p_i, q_i, s_i :$	Voting area vote percentage, turnout and size, order by time of announcement
$p_n, \mu_N :$	National vote share and prior mean for Brexit
$\sigma_P^2, \hat{\sigma}_P^2 :$	National vote share prior variance and estimate, for Brexit
$\mu_{p_i} :$	Expectation of p_i
$\mu_H :$	Vector of expectations provided by Hanretty study
$\sigma_{p_i}^2 :$	Marginal variance of p_i
$v_i :$	Expectation of q_i
$\sigma_v^2 :$	Marginal variance of q_i , independent of i
$r :$	Vector of variables = $(p_1, \dots, p_n, q_1, \dots, q_n)^\top$
$\tilde{r} :$	Re-ordered vector by time of announcement = $(p_1, q_1, \dots, p_n, q_n)^\top$
$F_i, \tilde{F}_i :$	Marginal CDFs of i 'th components of r and \tilde{r}
$\rho_\theta, \rho_\phi :$	Inter-area prior vote share and turnout correlation
$\rho_{\theta\phi} :$	Intra-area prior vote and turnout correlation
$\Sigma_0 :$	Covariance matrix of copula prior

$\tilde{\Sigma}_0 :$	Covariance matrix of copula prior with sequentially reordered rows and columns
$\tilde{\Pi} :$	Mean of prior after the announcement of m results
$\tilde{\Sigma} :$	Covariance matrix of prior after the announcement of m results

A.2 Prior probability of Brexit

$$P(\text{BREXIT})_0 = P\left(\frac{\sum_i p_i q_i s_i}{\sum_i q_i s_i} > \frac{1}{2}\right)$$

$$\Phi^{-1}(F_1(r_1)), \dots, \Phi^{-1}(F_{2n}(r_{2n})) \sim N(0, \Sigma_0)$$

$$\Sigma_0 = \begin{pmatrix} \Sigma_p & \Sigma_{pq} \\ \Sigma_{pq}^T & \Sigma_q \end{pmatrix}$$

$$\Sigma_p = (1 - \rho_p) I_n + (\rho_p) i_n i_n'$$

$$\Sigma_q = (1 - \rho_q) I_n + (\rho_q) i_n i_n'$$

$$\Sigma_{pq} = \rho_{pq} \times [(1 - (\rho_q \rho_p)) I_n + (\rho_q \rho_p) i_n i_n']$$

A.3 Prior marginal calibration

$$\mu_p = \mu_H + \alpha_N \times i$$

$$\sigma_{p_i}^2 = (\sigma_H^2)_i + \sigma_N^2$$

$$\sigma_N^2, \alpha_N : E(p_N) = \mu_N, \sigma_P = \hat{\sigma}_P^2$$

A.4 Update

$$\tilde{\Sigma}_0 = \begin{pmatrix} \tilde{\Sigma}_{m,m} & \tilde{\Sigma}_{m,\eta} \\ \tilde{\Sigma}_{\eta,m} & \tilde{\Sigma}_{\eta,\eta} \end{pmatrix}$$

$$\tilde{x}_m = \Phi^{-1}(F_{p_1}(p_1), F_{q_1}(q_1), \dots, F_{p_m}(p_m), F_{q_m}(q_m))'$$

$$\tilde{\Pi}_{\eta} = \tilde{\Sigma}_{\eta,m} \tilde{\Sigma}_{m,m}^{-1} \tilde{x}_m$$

$$\tilde{\Sigma}_{\eta} = \tilde{\Sigma}_{\eta,\eta} - \tilde{\Sigma}_{\eta,m} \tilde{\Sigma}_{m,m}^{-1} \tilde{\Sigma}_{m,\eta}$$

$$P(\text{BREXIT})_m = P\left(\frac{\sum_{i \geq m} p_i q_i s_i}{\sum_i q_i s_i} > \frac{1}{2} - \frac{\sum_{i \leq m} p_i q_i s_i}{\sum_i q_i s_i}\right)$$

$$\Phi^{-1}\left(F_{p_{m+1}}(p_{m+1}), F_{q_{m+1}}(q_{m+1}), \dots, F_{p_n}(p_n), F_{q_n}(q_n)\right)' | \tilde{x}_m \sim N(\tilde{\Pi}_{\eta h}, \tilde{\Sigma}_{\eta h})$$

Appendix B: Review of probability model in Wu et al. (2017)

The model under consideration in Wu et al. (2017) performs (in the one factor case) the following Weighted Least Squares regression following the announcement of k results

$$p_i = \alpha\mu_i + \beta + \epsilon_i, \quad \epsilon_i \sim N(0, \sigma_\epsilon^2) \quad i = 1, \dots, k.$$

The national vote share and thus the probability of Brexit is then simulated by generating M realisations and evaluating the relevant sum. A correct application of this method would involve sampling unknowns $(\alpha, \beta, \sigma_\epsilon^2)$ from the joint distribution

$$N \left(\begin{pmatrix} \hat{\alpha} \\ \hat{\beta} \end{pmatrix}, \begin{pmatrix} \hat{\sigma}_\alpha^2 & \rho_{\alpha\beta}\hat{\sigma}_\alpha\hat{\sigma}_\beta \\ \rho_{\alpha\beta}\hat{\sigma}_\alpha\hat{\sigma}_\beta & \hat{\sigma}_\beta^2 \end{pmatrix} \right), \chi^2(k-2),$$

where the slope and the intercept from linear regression are mutually correlated. Then the correct covariance and variance of unannounced results would be:

$$\begin{aligned} \text{cov}(p_i, p_j) &= E_{(\alpha, \beta, \sigma_\epsilon^2)} [\text{cov}(p_i, p_j | \alpha, \beta, \sigma_\epsilon^2)] \\ &\quad + \text{cov}_{(\alpha, \beta, \sigma_\epsilon^2)} [E(p_i | \alpha, \beta, \sigma_\epsilon^2), E(p_j | \alpha, \beta, \sigma_\epsilon^2)] \\ &= \text{cov}_{(\alpha, \beta, \sigma_\epsilon^2)} [\alpha\mu_i + \beta, \alpha\mu_j + \beta] \\ &= \mu_i\mu_j\hat{\sigma}_\alpha^2 + \hat{\sigma}_\beta^2 + [(\mu_i + \mu_j)\rho_{\alpha\beta}\hat{\sigma}_\alpha\hat{\sigma}_\beta] \\ \text{var}(p_i) &= \hat{\sigma}_\epsilon^2 + \mu_i\mu_j\hat{\sigma}_\alpha^2 + \hat{\sigma}_\beta^2 + [(\mu_i + \mu_j)\rho_{\alpha\beta}\hat{\sigma}_\alpha\hat{\sigma}_\beta]. \end{aligned} \tag{B.0.1}$$

However, an assumption that $\hat{\alpha}$ and $\hat{\beta}$ are uncorrelated appears to be used when in fact they are close to being perfectly anti-correlated with¹

$$\rho_{\alpha\beta} = \frac{-\sum_1^k \mu_i}{\sqrt{n \sum_1^k \mu_i^2}} = -\frac{\bar{\mu}}{\sqrt{\mu^2}} \approx -1.$$

The calculation in Wu et al. (2017) would thus be calculating the variance structure of the unknown referendum results as:

$$\begin{aligned} \text{cov}(p_i, p_j) &= \mu_i \mu_j \hat{\sigma}_\alpha^2 + \hat{\sigma}_\beta^2 \\ \text{var}(p_i) &= \sigma_\epsilon^2 + \mu_i \mu_j \hat{\sigma}_\alpha^2 + \hat{\sigma}_\beta^2, \end{aligned}$$

which is different to the values in equation B.0.1.

Another issue with the method is that the assumption of a normal distribution for $(\hat{\alpha}, \hat{\beta})$ with correctly specified errors is highly questionable, particularly in small samples. $(\hat{\alpha}, \hat{\beta})$ only follows such a distribution in finite samples if the errors are normal, and if not, it would be biased but consistent. The normal distribution is an asymptotic result for non-normal errors, and even then a correct evaluation assumes no heteroskedasticity; otherwise error estimates are likely to be too low and implied probabilities of Brexit to be too confident. This could be overcome using robust errors, but only for large data sets. Using robust errors in small samples can produce severely biased estimators.

The model in Wu et al. (2017) and that presented in this paper use different approaches to estimate the covariance structure of the conditional distributions used to form predictions. That of Wu et al. (2017) requires no prior (beyond expectations) and attempts to infer the covariance structure from an OLS regression of results announced so far. Our model, by contrast, starts with a prior for the covariance structure and updates that prior as results come in. Both methods will produce the same covariance and result in larger samples but will be different for small samples. The different approach is illustrative of the differences between a Frequentist and Bayesian approach to inference. However, we suggest that the Frequentist approach presented in Wu et al. (2017) is not appropriate for the small

¹A simulation of their results yielded $\rho_{\alpha\beta} \in (-0.97, -1)$.

numbers of results available at the times of predictions (<20 results). More sophisticated corrections for small sampling estimation would be desirable. We believe our Bayesian approach is a more suitable way to proceed in the case of this application.

Appendix C: Power of two-step empirical specification versus a single regression approach

We demonstrate that using our two step approach to testing political betting significance (outlined in section 3.4) has superior power than using a single larger factor model regression (equation 4.1). The latter (falsely) assumes an unchanging political probability on the training set. The key assumption is that the political factor, although not necessarily observable, has non-zero variance in some part of the training period. To show this consider the simple univariate regression

$$r_t = \alpha + \beta \Delta P_t + \epsilon_t \quad t = 1, \dots, T_1, T_1 + 1, \dots, T_2 \quad \mathbb{E}(\epsilon_t) = \sigma_\epsilon^2 \quad (\text{C.0.1})$$

where r_t is the return of a financial asset at time t and ΔP_t is the change in the probability of the upcoming event. Assume that this is the correct model, that $\beta \neq 0$ and note that $\mathbb{E}(\Delta P_t) = 0$. Data is available for r_t for entire period $t = 1, \dots, T_2$. ΔP_t is in general not observed for the first period $t = 1, \dots, T_1$ but can be measured by the betting market prices ΔPB_t for $t = T_1 + 1, \dots, T_2$. Consider the regression on the period where we have data for both r_t and ΔP_t , ($t = T_1 + 1, \dots, T_2$). This is analogous to regressing the estimated Fama-French residual on changes in the political markets on the testing set only. The t-statistic for β in this regression is t_β where

$$t_\beta \longrightarrow \sqrt{T_2 - T_1} \times \frac{\text{COV}(r_t, \Delta P_t)}{\sigma_\epsilon \cdot \sigma_{\Delta P_t}}.$$

Next define

$$\tilde{\Delta P}_t = I(t > T_1) \cdot \Delta P B_t$$

which is ΔP_t forced erroneously to zero on the first set $T = 1, \dots, T_1$ where data is lacking. Now consider a regression of r_t on \tilde{P}_t on the whole period for $t = 1, \dots, T_2$

$$r_t = \alpha + \beta \tilde{\Delta P}_t + \tilde{\epsilon}_t \quad t = T_1 + 1, \dots, T_2 \quad \mathbb{E}(\tilde{\epsilon}_t) = \tilde{\sigma}_\epsilon^2$$

This is analogous to the single regression empirical approach where we jointly estimate the Fama-French loadings with a false assumption that the political probability is unchanging in the first period. The relevant t-statistic for β is now

$$\begin{aligned} \tilde{t}_\beta &= \sqrt{T_2} \times \frac{\frac{1}{T_2} \sum_{t=1}^{T_2} r_t (\tilde{\Delta P}_t - \frac{1}{T_2} \sum_{s=1}^{T_2} \tilde{\Delta P}_s)}{\hat{\sigma}_\epsilon \times \sqrt{\frac{1}{T_2} \sum_{t=1}^{T_2} (\tilde{\Delta P}_t - \frac{1}{T_2} \sum_{s=1}^{T_2} \tilde{\Delta P}_s)^2}} \\ &= \sqrt{T_2} \times \frac{\frac{1}{T_2} \sum_{t=1}^{T_2} r_t (\tilde{\Delta P}_t - \frac{1}{T_2} \sum_{s=T_1+1}^{T_2} \tilde{\Delta P}_s)}{\hat{\sigma}_\epsilon \times \sqrt{\frac{1}{T_2} \sum_{t=1}^{T_2} (\tilde{\Delta P}_t - \frac{1}{T_2} \sum_{s=T_1+1}^{T_2} \tilde{\Delta P}_s)^2}} \\ &= \sqrt{T_2} \times \frac{\frac{1}{T_2} \sum_{t=1}^{T_2} r_t (\tilde{\Delta P}_t - \frac{(T_2-T_1)}{T_2} \hat{\Delta P})}{\hat{\sigma}_\epsilon \times \sqrt{\frac{1}{T_2} \sum_{t=1}^{T_2} (\tilde{\Delta P}_t - \frac{(T_2-T_1)}{T_2} \hat{\Delta P})^2}} \\ &= \sqrt{T_2} \times \frac{\frac{1}{T_2} \sum_{t=T_1+1}^{T_2} r_t \Delta P_t - \frac{T_1(T_2-T_1)}{T_2^2} \hat{\Delta P}}{\hat{\sigma}_\epsilon \times \sqrt{\frac{1}{T_2} \sum_{t=T_1+1}^{T_2} \Delta P_t^2 - \frac{2(T_2-T_1)}{T_2^2} \hat{\Delta P}}} \\ &= \sqrt{T_2} \times \frac{\frac{(T_2-T_1)}{T_2} \hat{COV}(r_t, \Delta P_t) - \frac{T_1(T_2-T_1)}{T_2^2} \hat{\Delta P}}{\hat{\sigma}_\epsilon \times \sqrt{\frac{(T_2-T_1)}{T_2} \hat{VAR}(\Delta P_t) - \frac{(T_2-T_1)^2}{T_2^2} \hat{\Delta P}^2}} \end{aligned}$$

where we have used the fact that $\tilde{\Delta P}_t = 0$ for $t \leq T_1$ and

$$\begin{aligned}
\frac{\hat{\Delta P}}{\Delta P} &= \frac{1}{(T_2 - T_1)} \sum_{s=T_1+1}^{T_2} \Delta P_s \longrightarrow 0 \\
\hat{\sigma}_\epsilon &= \sqrt{\frac{1}{T_2} \sum_{t=1}^{T_2} \tilde{\epsilon}_t^2} \longrightarrow \sigma_{\tilde{\epsilon}} \\
C\hat{O}V(r_t, \Delta P_t) &= \frac{1}{(T_2 - T_1)} \sum_{t=T_1+1}^{T_2} r_t \Delta P_t \longrightarrow COV(r_t, \Delta P_t) \\
V\hat{A}R(\Delta P_t) &= \frac{1}{(T_2 - T_1)} \sum_{t=T_1+1}^{T_2} \Delta P_t^2 \longrightarrow \sigma_{\Delta P_t}^2.
\end{aligned}$$

The first result comes from the fact that $\mathbb{E}(\Delta P_t) = 0$. Thus

$$\tilde{t}_\beta \longrightarrow \sqrt{T_2} \frac{\frac{(T_2-T_1)}{T_2} COV(r_t, \Delta P_t)}{\sqrt{\frac{(T_2-T_1)}{T_2}} \sigma_{\tilde{\epsilon}} \cdot \sigma_{\Delta P_t}} = \sqrt{T_2 - T_1} \times \frac{COV(r_t, \Delta P_t)}{\sigma_{\tilde{\epsilon}} \cdot \sigma_{\Delta P_t}}$$

Note that $t_{\tilde{\beta}}$ is very similar to t_β and differs only in the denominators σ_ϵ and $\sigma_{\tilde{\epsilon}}$. However, we have assumed that

1. $\beta \neq 0$
2. $\{x_1, \dots, x_{T_1}\}$ is not everywhere zero
3. Equation (C.0.1) represents the *true* model.

However, if the model is indeed true it cannot provide a worse fit than the model

$$y_t = \alpha + \beta \tilde{x}_t + \tilde{\epsilon}_t \quad t = 1, \dots, T_2.$$

When assumptions 1 and 2 above hold, the fit must be strictly worse. Thus $\sigma_{\tilde{\epsilon}} > \sigma_\epsilon$. The limit of the second t-statistic has smaller absolute size than that of the first t-statistic. Thus under an alternative hypothesis $\beta \neq 0$ the second t-statistic will never reject the null hypothesis $\beta = 0$ unless the first t-statistic does. t_β has greater power than \tilde{t}_β .

If we are testing the hypothesis $\beta \neq 0$ in a univariate regression and we lack data for the explanatory variable for some part of a sample set, then a higher power test is provided by regressing on only the sub-sample where data is known for both

the explanatory and dependent variables rather than regressing on the whole set and forcing the dependent variable to zero where there are no observations. This is exactly analogous to showing that jointly estimating the Fama-French residuals along with the political component, whilst forcing the political factor to be zero, has lower power than the two step approach we employ.