

Forecasting the Outcome of Closed-Door Decisions: evidence from 500 years of betting on papal conclaves

October 2014

Leighton Vaughan Williams

Professor of Economics and Finance

Nottingham Business School

Nottingham Trent University

Burton Street

Nottingham NG1 4BU

United Kingdom

Tel: +44 115 848 6150

Email: Leighton.Vaughan-Williams@ntu.ac.uk

David Paton*

Professor of Industrial Economics

Nottingham University Business School

Wollaton Road

Nottingham NG8 1BB

United Kingdom

Tel: +44 115 846 6601

Fax: +44 115 846 6667

Email: David.Paton@Nottingham.ac.uk

* corresponding author

Acknowledgements

We would like to thank participants at a Nottingham Trent University Staff Seminar and at Georgia State University's 'Center for the Economic Analysis of Risk' Workshop, held in January, 2014 at the University of Cape Town, for several helpful comments and suggestions.

Forecasting the Outcome of Closed-Door Decisions: evidence from 500 years of betting on papal conclaves.

Abstract

Closed-door decisions may be defined as decisions in which the outcome is determined by a limited number of decision-makers and where the process is shrouded in at least some secrecy. In this paper, we examine the use of betting markets to forecast on particular closed-door decision, the election of the Pope. Within the context of 500 years of papal election betting, we employ a unique dataset of betting on the 2013 papal election to investigate how new public information is incorporated into the betting odds. Our results suggest that the market was generally unable to incorporate effectively such information. We venture some possible explanations for our findings and offer suggestions for further research into the prediction and predictability of other ‘closed-door’ decisions.

JEL codes: D72, G14, L83, Z12

Keywords: closed-door decisions; information; Papal conclave; market efficiency.

Forecasting the Outcome of Closed-Door Decisions: evidence from 500 years of betting on papal conclaves.

1. Introduction and Motivation

The forecasting of decision-making in the context of elections and public policy is the subject of an extensive literature. An increasingly significant part of this work considers the relevance and efficiency of betting or prediction markets in making these decisions (e.g. Page, 2011; Saville et al., 2011). Betting on election outcomes has a long history, and is particularly well-documented in the case of presidential elections in the US (Rhode and Strumpf, 2013). It has indeed been traced, according to contemporaries, back to the election of George Washington and has existed in organized markets since the administration of Abraham Lincoln.

Within the context of political betting and forecasting, we can distinguish between open- and closed-door decisions. The former are open to public scrutiny whilst the latter can be defined as decisions taken by an individual or group of individuals whose choices are shrouded, at least from outsiders, in a layer of secrecy. Examples of closed-door decisions include court rulings and verdicts.

The forecasting of closed-door decisions has attracted a less extensive literature, confined for the most part to the prediction of Supreme Court rulings (e.g. Epstein et al, 2010; Johnson et al., 2009; Ruger et al., 2004), and only a very small part of this literature considers the relevance or efficiency of betting or prediction markets in making these decisions (Blackman et al., 2012; Cherry and Rogers, 2006).¹

This paper seeks to address this gap in the literature in the context of a closed door decision for which betting markets have existed for at least 500 years, the outcome of papal conclaves. The decision over the choice of Pope is of considerable significance for a number of reasons. The Pope is perhaps the last absolute monarch (Padovano and Wintrobe, 2013) and, as such, exercises a very significant degree of influence over the shape and direction of the Catholic Church. In turn this has a direct impact on the lives of millions of Catholics world-wide. Further, due to the widespread activities of the Church in the fields of health,

¹ Nevertheless, a recent analysis of the 2012 Supreme Court decision about the ‘Affordable Care Act’ does support the value of expert forecasters in predicting Court decisions (Vaughan Williams, 2014). In particular, the eventual outcome was one which might have been read months before the decision through close scrutiny of the findings of the American Bar Association’s (2012) survey of a select group of Supreme Court experts. In this respect, the expert forecasters performed significantly better than the real-money exchange betting market.

education, economic development and diplomacy, the selection of a Pope has a global significance well beyond the immediate sphere of the Catholic Church.

This is also an interesting type of election in that the electorate might be regarded as relatively ‘non-complex’, defined by Jottier et al. (2012) in terms of size and heterogeneity. Jottier et al. argue that in such cases, prediction accuracy should be high relative to where there are larger, more heterogeneous electorates.

The specific focus of this paper is to examine how efficiently (or otherwise) the papal betting markets were able to incorporate both public and private information revealed during the course of the conclave. In this task, we are fortunate in having access to a unique dataset which provides detailed information on the betting market for every potential papal contender. In the next section of the paper, we draw together several sources to summarise what is known about gambling in historical papal elections. We then explain the background and context for betting on the 2013 conclave. In section 4, we introduce our data and the empirical methodology we employ. Finally we report and discuss our empirical results.

2. Betting on Papal Elections: the historical context

The first recorded example of betting on a papal election can be traced to the papal conclave of September, 1503, at which time it was considered already “an old practice” (Baumgartner, 2003 p.250; Villard, 2009). The brokers in the Roman banking houses who made books and offered odds on who would be elected,² made Cardinal Francesco Piccolomini the 100 to 30 (against) favourite, ahead of Cardinals Guiliano della Rovere (100 to 15) and Georges d’Amboise (the favourite if judged by the vocal support of the street crowds) at 100 to 13 (Baumgartner, 2003 p.88) Although Piccolomini is thought to have trailed in the first round of voting with 4 votes to 13 for d’Amboise and 15 for della Rovere, Piccolomini apparently benefited from a switch of votes from d’Amboise to himself in subsequent voting, and duly became Pope Pius III. The bookmakers were proved right.

The next conclave for which we have the betting odds is that of December, 1521, in which odds were offered on no fewer than twenty cardinals. Giulio de’Medici, the cousin of Leo X, was the betting favourite, at 100 to 25 (4 to 1), followed closely by Cardinal Alessandro Farnese at 100 to 20 (5 to 1), whose odds shortened to 100 to 40 (5 to 2) after a Roman mob plundered his house (Baumgartner, 2003 pp. 95-6). Though Farnese at one point came close to being elected Pope, he could not reach the required two-thirds of the vote

² See Hunt (2012 pp. 367-8) for details of how and where they conducted their operations.

(Brewer, 1920 p. 798), and ultimately the cardinals looked outside of the conclave, electing Adrian of Utrecht as Pope Adrian VI.

During the papal conclave of 1549-1550, Dandolo describes how Cardinal Gianmaria del Monte (who was eventually elected Julius III) had opened in the betting as the 5 to 1 (against) favourite, but within three days Cardinal Reginald Pole had been established at odds of 4 to 1 (CSP, 1970 p.274-6). On December 5, as balloting began, Pole was clear favourite at 100 to 95. On that day, he received 26 of the 28 votes that would have given him the two-thirds majority required to elect him Pontiff. Although on the point of being made Pope by acclamation, Pole insisted on waiting until he won the formal two-thirds majority (Baumgartner, 1985 p. 306). By the time that four additional French cardinals, opposed to Pole, arrived December 11, however, he was trading at 5 to 2, and a month later he was being offered at odds of 100 to 16 (Baumgartner, 2003 pp.108-9). Dandolo had reported to his superiors early in the conclave: "It is more than clear that the merchants are very well informed about the state of the poll, and that the cardinals' attendants in Conclave go partners with them in wagers, which this causes many tens of thousands of *scudi* (crowns) to change hands" (CSP, 1970 p.281).

The first 1590 conclave, in September, is the earliest in which reports of insider trading emerged, when two of the key influencers of votes in the conclave, Cardinals Montalto and Sforza secretly agreed to join forces in support of Niccolo Sfondrato. It is reported that both made fortunes betting on him, at odds of 10 to 1 the day before he was elected as Pope Urban VII. As the conclave opened, he was trading at 100 to 11, compared to Giambattista Castagna, who was offered at 100 to 22 (Baumgartner, 2003). During the second conclave of 1590, Cardinal Gabriele Paleotti at one point increased to 70 per cent in the betting.³ The odds were not reflected in the outcome, when Giovanni Battista Castagna became Pope Urban VII.

Johnson (1974) reports bookmaker odds in Milan for the 1958 conclave which show Cardinal Angelo Roncalli the 2 to 1 favourite, followed by Cardinals Agagianian and Ottaviani at 3 to 1, then Stefan Wyszynski and Giuseppe Siri at 4 to 1. The odds were justified when Roncalli became Pope John XXIII.

For the first conclave of 1978, bookmakers in London were offering odds of 5 to 2 about Cardinal Sergio Pignedoli, 7 to 2 about Sebastian Baggio and Ugo Poletti and 4 to 1 about Carlo Benelli. The best odds about a non-Italian were 8 to 1 about Johannes

³ BAV, Ur. Lat 1058, October 13, 1590, fol. 525v, cited in Hunt (2012).

Willebrands. Of these only Pignedoli showed any strength in the voting, unconfirmed reports of the voting indicating that he obtained about 18 votes in the first ballot, compared to about 23 for Albino Luciani and 25 for Giuseppe Siri. Ultimately, Luciani became Pope John Paul I.

For the second conclave of 1978, following the death of Pope John Paul I, the Associated Press, on October 14, noted that:

“Once again, there is no odds-on favourite to be elected as the new pope of the Roman Catholic Church Those mentioned most often were Corradi Ursi, 70, of Naples; Salvatore Pappalardo, 60, of Plaermo, Sicily; Ugo Poletti, 64, of Rome; Giuseppe Siri, 72, of Genoa; Giovanni Colombo, 75, of Milan; Giovanni Benelli, 57, of Florence, and Antonio Poma, 68, of Bologna ... Non-Italian front-runners included Argentinian Eduardo Pironio, 57, and Dutchman Johannes Willebrands, 68.”⁴

In fact, Cardinal Carol Wojtyla, archbishop of Krakow, became Pope John Paul II, after the eighth ballot

In 2005, the man who became Pope Benedict XVI, Cardinal Joseph Ratzinger, opened in the betting, according to the odds offered by the bookmaker, Paddy Power, at 12 to 1. At that point, William Hill made Cardinal Arinze favourite, with Archbishop Tettamanzi, Cardinal Ratzinger and Cardinal Hummes as others to watch (Fleishman, 2005). At lunchtime Tuesday, April 19, after three ballots, Ratzinger was favourite on two out of the three online betting boards monitored by CNN,⁵ and by the last day of the conclave had shortened to a clear 3 to 1 favourite.

Toman (2004) analyses the dynamics of conclave voting using data collected from seven conclaves, beginning with the election of Benedict XV in 1914 to John Paul II in 1978. Modelling the election procedure using a linear feedback count panel data model, she found three significant patterns. First, that the number of votes obtained during the previous ballot is strongly and positively correlated with the votes obtained during the ongoing ballot. Secondly, she found a momentum effect, so that the growth in votes a cardinal obtains between the previous ballot (at time $t-1$) and the one before that (time $t-2$) is positively correlated with votes obtained during the current ballot (time t). In other words, candidates tend to give more votes to a cardinal whose votes are seen as growing, and vice-versa. Finally, the effect of “nocturnal conversations” (discussions made after the end of the day)

⁴ <http://fivethirtyeight.blogs.nytimes.com/2013/03/12/election-of-a-pope-tests-betting-markets/>

⁵ CNN.com, Bookmakers lay odds on new pope, Tammy Oaks, April 19, 2005.
<http://edition.cnn.com/2005/WORLD/europe/04/18/pope.betting/>

tends to sizeably reduce the number of votes obtained by the cardinal leading in the vote. A possible explanation for this finding is that these conversations allow cardinals to coordinate and hinder the election of the leading cardinal.⁶

In summary, papal conclaves do have some history, dating back to 1503, of electing one of the favourites in the betting, but this is by no means a general rule, and there is some evidence of predictable patterns of voting. The historical evidence is that the betting markets have a patchy record in assimilating information about the identity of the next Pope, but there is clear historic evidence, not least from tracking movements in the betting odds during the course of the conclave, that the markets do show evidence of having picked up in a number of documented cases genuine information of predictive value about the outcome of papal elections.

3. The 2013 Conclave

We now turn to the 2013 conclave which eventually led to the election of Cardinal Jorge Bergoglio as Pope Francis I. In the run-up to the conclave, a survey of Vatican watchers⁷ by YouTrend. It listed Cardinal Timothy Dolan of the United States as the second most likely Pope, after Cardinal Angelo Scola, followed in order by Cardinals Marc Ouellet, Odilo Scherer and Thomas O'Malley. Luis Tagle of the Philippines was ranked sixth.

An analysis was also reported in the National Catholic Register⁸ of the likely age of the incoming Pope based on three trends – age at vacancy, length of reign and age at election. In terms of age at leaving office, the last several Popes (except for John Paul 1) were over 80 years old, a result of a rising trend over the last 500 years. A similar rising trend in the length of reign indicates that a reign of 15 years is now an indicative guideline. An upward, though less steep, upward trend in the age at election, indicates that someone in their late 60s (around 68) is a good guideline. In conclusion, this analysis pinpointed the expected age of the new pope to be about 68, to be expected to reign for about 15 years. More generally, the analysis concluded that “the next pope is likely to be between 60 and 70.” Of the 115 cardinal electors, 47 were in this age range. Jorge Bergoglio (76) was not one of them.

⁶ See also: www.linkiesta.it/blogs/una-firma-di-tutto-riposo/how-do-cardinals-vote-statistical-analysis-papal-conclaves

⁷ www.youtrend.it/2013/03/12/totopapa-sondaggio-youtrend-vaticanisti-previsioni-papabili/

⁸ Akin, J, How Old Will the Next Pope Be? National Catholic Register, February 25, 2013. Available at: <http://www.ncregister.com/blog/jimmy-akin/how-old-will-the-next-pope-be>

Of perhaps more significance than age is country or region of origin. In particular, a Pew Research Report⁹ examining the regional distribution of the world's Catholics showed that while 65% of Catholics lived in Europe in 1910, by 2010 that had declined to 24%, while the share in Latin American –Caribbean countries rose from 24% to 39%, in Asia-Pacific countries from 5% to 12%, in Sub-Saharan Africa from less than one per cent to 16%, and in North America from 5% to 8%. Middle East-North African countries made up less than one per cent in both periods. Of arguably more importance than the regional distribution of the world's Catholics, however, may be the regional distribution of the 115 attending Cardinal electors. Of these, 28 were from Italy, 32 from the Rest of Europe, 20 from North America, 13 from South America, 11 from Africa and 11 from Asia (including the Middle East) and Oceania. For the 2005 conclave, the corresponding figures were 39 from Italy, 30 from the Rest of Europe, 17 from North America, 9 from South America, 10 from Africa and 10 from Asia (including the Middle East) and Oceania.

An assessment released on March 12, 2013, noted that “Almost like clockwork since 1878, every election has alternated between producing a favourite, and an almost complete surprise. In 2005, 1963, 1939, 1914 and 1878 the Cardinal Electors played it safe, while in 1978 (twice), 1958, 1922 and 1903 they were prepared to ‘roll the dice’, often in order to break a deadlocked conclave. If this pattern holds, 2013 will produce a (late) surprise ... Could the surprise be the first non-European pope in 1,282 years? After all, two-thirds of Catholics now reside outside of Europe, forecast to grow to three-quarters by 2050 ...”¹⁰

What is most notable about the pre-conclave speculations was how little attention was paid by observers to the chances of Cardinal Bergoglio. John Allen Jr. (the Vatican expert for the US publication, *The National Catholic Reporter*) was unusual in at least profiling Bergoglio as a possible contender, albeit in the context of profiles over twenty other ‘papabile’, although he was rather ambivalent about Bergoglio’s chances.¹¹

We now look at events during the process of the conclave itself. We are particularly interested in examining how much new information (either private or public) that became available during the conclave might have influenced perceptions about the outcome. We focus on four significant events as follows:

- A. 18:41 GMT on Weds 12th March. This is when the first black smoke appeared, indicating that the Cardinal electors had not agreed on a pope following the first

⁹ Pew Research Religion & Public Life Project, February 13, 2013, ‘The Global Catholic Population.’ Available at: <http://www.pewforum.org/2013/02/13/the-global-catholic-population/>

¹⁰ Rod Crosby, <http://ncronline.org/blogs/ncr-today/papabile-day-men-who-could-be-pope-13>

¹¹<http://ncronline.org/blogs/ncr-today/papabile-day-men-who-could-be-pope-13>

round of ballots. At this point in time, the Cardinals exit the Sistine chapel where the voting takes place, and the first opportunity arises for the leaking of private information on the first voting tallies.

- B. 08:10 GMT, publication of a ‘Vatican Insider’¹² report in La Stampa by Vatican reporter, Giacomo Galeazzi, identifying the group of candidates leading the ballot, namely Cardinals Dolan, Ouellet, Scherer, Scola and Bergoglio. This is the first and only report in the public domain claiming to know the state of the voting, and was offering what we now know to have been reliable information on the progress of the ballot. This was also the first credible report indicating that Cardinal Bergoglio was a serious contender. What is particularly important to note here is that ‘Vatican Insider’, and Giacomo Galeazzi in particular, are sources which were at the time already very well regarded by informed observers.¹³
- C. 10:38 GMT, second black smoke revealing that the second round of ballots had not led to a Pope being elected. This indicates a second point in time when new, but private information was available.
- D. 11:12 GMT, updated report by Galeazzi, suggesting that, in the most recent round of voting, the shortlist of candidates had reduced to just three: Scola, Ouellet and Bergoglio. At 11.57am, the Guardian, a UK national newspaper with a significant Internet presence, reported on the two Galeazzi articles on their Liveblog, thus bringing this information to the attention of a much wider, English-speaking audience.

A subsequent report, published in La Repubblica after the election of the Pope, claims that Scola received approximately 35 votes in the first vote, to 20 for Bergoglio and 15 for Ouellet,¹⁴ an account backed up in broad terms by informal post-conclave interviews with five cardinal electors by John Allen Jr. at the National Catholic Reporter - which additionally reported support for Scherer. Allen continues that “After two rounds of voting Wednesday morning, it had become clear that neither Scola nor Scherer were likely to cross the finish line and gain the 77 votes needed for election ... The fourth ballot, the first of Wednesday afternoon, saw Bergoglio separate himself from the pack.”¹⁵

¹² <http://vaticaninsider.lastampa.it/en/about-us/>

¹³ <http://ncronline.org/blogs/all-things-catholic/ferment-religious-life-new-american-leader-and-vatican-insider>

¹⁴ www.ilvelino.it/it/article/orenove5-papa-in-conclave-un-plebiscito-quasi-cento-voti/493c9365-6d5a-44a7-9763-11c9364f38d3

¹⁵ John L. Allen Jr., Path to the Papacy: ‘Not Him, Not Him, Therefore Him’, National Catholic Reporter, March 17, 2013, <http://ncronline.org/news/global/path-papacy-not-him-not-him-therefore-him>

So it is clear that both Galeazzi reports on the outcomes of each round of voting (the second one subsequently picked up and reported in the Guardian's 'Liveblog') are cases in which accurate information, of great significance to the eventual outcome, was made publicly available. It is interesting to observe that even the order of listing of the names coincided with the vote tally as reported on 19th March in La Repubblica. More importantly, by comparing the Galeazzi update with the earlier report, it is clear that Bergoglio's vote tally was rising, while the early front-runners had failed to reach the required two-thirds majority by the time of the 'nocturnal conversations'. For these reasons, we might have expected the reports to have led to a surge of betting interest on Cardinal Bergoglio, unless there was good reason to question the credibility of the Vatican Insider reports. We might also expect to have seen a flight of money away from previously favoured candidates not mentioned in the reports.

In the next section we examine how the extent to which prices for the key contenders actually responded to the release of both private and public information.

4. Methodology and Data

We focus on prices during the final twenty four hours of the conclave betting market. For our purposes, price is defined as the probability of a contender winning the election as implied by the Betfair odds. As described above, during this time, we identify four key time points (A to D) at which significant additional information was potentially available to the market.

Events A and C provided public information to the extent that they revealed that no Pope had been elected. More importantly, however, at these points, insiders who were aware of the breakdown of votes in the inconclusive ballots would have been in possession of information unavailable to the general public regarding the likely prospects of the key contenders. In principle, this information could have been exploited to engage in insider trading on the betting markets.

Events B and D indicate points at which that same information was made publicly available. At 11:57 on 13th March (a little under one hour after Event D, the publication of the second La Stampa article) the Guardian live blog reported on the article, naming the three remaining contenders and hence publicising this information more widely to an English-speaking audience. We consider the implications of the publication of the Guardian blog below.

The betting data we use comes from Betfair, the world's largest person-to-person betting exchange.¹⁶ Betfair supplied us with a unique dataset comprising a complete record of every bet (recorded and time-stamped in-running) for every candidate in the papal election. Each record includes the amount bet, the price achieved and the timestamp. For the majority of our analysis we aggregate the data up to the hourly level to ensure there is a reasonable level of liquidity. The mean size of a wager in the papal conclave was just over £10. In total, over the course of the market, over 17,000 bets were placed giving a total sum wagered of £180,312. Anyone based in Italy, however, at the point of trading, was not able to place a bet through Betfair on the conclave.

Our approach is to track graphically prices at the end of each hour on the betting markets of the main contenders over the final twenty four hours of the conclave market and to examine the effect on prices at each of these points. We identify two groups of contenders. The first group comprises the five Cardinals identified by La Stampa (and subsequently verified by other sources) as having attracted significant numbers of ballots in the early rounds of voting. The second group comprises the six other Cardinals who were most favoured in the betting odds at the opening of the market, specifically Cardinals Bertone, Erdo, O'Malley, Schonborn, Tagle and Turkson.

Having examined the impact of our four events on the prices of each cardinal, we then go on to estimate a formal econometric model of market efficiency for each of the three main contenders (as revealed by the second La Stampa report and subsequently verified by other sources). We use a standard approach to testing for information inefficiencies by testing whether past movements in asset prices can be used to predict positive returns. On the assumption of an efficient market, current returns should follow a random walk process and lagged returns should have no explanatory power. When estimating such models, it is important to take account of the impact of time-varying volatility, or Autoregressive Conditional Heteroscedasticity (ARCH) (Engle, 1982). Not doing so is likely to lead to biased and inconsistent estimates. There exist a whole class of models to deal with ARCH effects. Most common in the analysis of asset prices is the use of Generalised ARCH (GARCH) models (Bollerslev, 1986). In these models, the time-dependent volatility is estimated as a function of observed prior volatility, measured as the lagged value(s) of the squared regression disturbances and, also, lagged value(s) of the conditional variance. The order of the GARCH model is given by the number of lags in each case.

¹⁶ A betting exchange is an on-line trading platform provides clients of the exchange with the facility to offer or accept odds about the outcome of a future event.

In the context of market efficiency for asset prices, the GARCH(p, q) model can be represented as follows:

$$R_t = \alpha_o + \sum_{i=1}^k \alpha_i R_{t-i} + \varepsilon_t \quad (1)$$

$$\sigma_t^2 = \gamma_0 + \sum_{i=1}^q \gamma_i \varepsilon_{t-i}^2 + \sum_{i=1}^p \delta_i \sigma_{t-i}^2 \quad (2)$$

where R_t represents returns in hour t , ε_t is assumed to follow a normal distribution with zero mean and variance σ_t^2 ; γ_i are the ARCH parameters; δ_i are the GARCH parameter(s). We define returns on hour t in the normal way as $R_t = \log(P_t/P_{t-1})$ where P_t is the betting price for the contender at the end of hour t . We use the Akaike Information Criterion (AIC) to determine the optimal lag length of the ARCH and GARCH parameters.

We supplement equation 1 by the inclusion of dummy variables for each of the four time points identified above. We include one lag of event D (publication of the second La Stampa article) to allow for the further dissemination of the information via the Guardian. Significant coefficients on dummy variables for events A and C will provide information on the extent to which private information held by insiders was incorporated into the market. The coefficients for events B and D indicate how the market incorporated public information. We would expect the effect of these dummy variables to vary with each contender. In particular, for Cardinal Bergoglio who was not originally considered by the markets to be a realistic contender, we would expect a significant and positive coefficient for the two public information dummies. A positive coefficient for the two private information dummies would indicate significant trading by insiders on the basis of such information.

We estimate this model using hourly data from the opening of the market at 21.00 GMT on the 28th February until the close of the market at 18.00 GMT on the 13th March. This gives us a total of 310 observations.

5. Results

In Figures 1a and 1b we track hourly mean price for individual Cardinals in the critical final 24 hours period of betting. We present plots for two groups of Cardinals. The first is those identified *ex post* as having attracted significant numbers of votes in the early rounds. The second are those cardinals who were expected to poll well prior to voting but who were identified *ex post* as receiving few or no votes.

Taking each of our events in turn, event A (the first appearance of black smoke) indicates a time at which insiders would have been aware that the first group of cardinals (reported in Figure 1) were all still in the running, whereas the chances of those in the second group (who had originally attracted a lot of betting but did not attract many or any votes) had clearly been overestimated. In fact, the time paths provide very little evidence that insiders exploited this information to any significant degree. In many cases, prices changed very little after that point or continued on a previously established trend. Indeed, in some cases the graphs suggest a perverse effect. For example, Cardinal Turkson's price increased markedly in the period soon after event A, contrary to what we would expect if insiders were exploiting their private information. Looking at the two candidates (Bergoglio and Dolan) who attracted votes despite previously not being thought of as serious contenders, if anything prices seem to drop after event A. A close look at the minute-by-minute data for Bergoglio confirms an increase in volatility around the time of event A, and the price eventually settles at a lower value than before the black smoke, the opposite to what we would expect had the private information made available at that point (namely that Bergoglio had attracted a significant number of votes) been exploited in the markets.

Event B indicates the point at which the private information becomes public. Again, there is little consistent pattern in the data to suggest that the markets successfully incorporated this information. Indeed, there is very little movement in prices at all at this point.

Event C (the second appearance of black smoke) marks the second period in which there was the potential for insider trading. This period lasts just one hour before the voting information was made public by La Stampa (event D). Again, there is no clear evidence that the prices of the candidates reacted in a way consistent with what we would expect if there was insider trading, although Bergoglio's price does appear to increase marginally after this point. At the time of Event D, Bergoglio's odds of winning appear to drop at first (again the opposite of what we would expect). His price eventually rises significantly, in line with expectations, and the timing of this lagged response matches the subsequent re-posting of the La Stampa report by The Guardian.. There are other apparent inconsistencies. For example, Schonborn's odds of winning increase markedly soon after Event D, despite there being no public indication (as far as we are aware) to indicate an increase in his chances.

We now go on to look at standard market efficiency regressions for some of the main candidates: Cardinals Bergoglio, Ouellet and Scherer. As explained above, we employ a simple GARCH specification using hourly data over the course of the whole market. We also

include dummy variables for the hour of occurrence of each of the four events discussed above. The key results are reported in Table 1.

Taking the market for Cardinal Bergoglio first, there is no overall evidence of market inefficiency in that mean returns (as revealed by the constant term) are not significantly different from zero and the lagged returns variable has no significant explanatory power. In contrast, for both Ouellet and Scola, the coefficient on lagged returns is negative and significant suggesting that lagged returns have significant explanatory power.

The results for the four timed events presents us with a mixed picture. The appearance of the first black smoke is associated with significantly negative returns for Bergoglio and Ouellet, the opposite to what we would expect had insiders used the private information that these candidates had attracted a significant number of votes in the early rounds. The second private information event (second black smoke) is more consistent with expectations: for both Bergoglio and Ouellet the coefficient on the dummy variable is positive and significant.

Looking at the revelation of public information regarding voting patterns, the publication of the first La Stampa article appears to have had little or no effect on returns for Bergoglio. This is striking given that it was the first appearance of any information indicating that he was a front-runner. Intriguingly, the second episode of additional public information (the second La Stampa article) is associated with negative returns at first and only when the Guardian blog further disseminated this information, do returns become strongly positive and significant. Taken together, these results are suggestive of the markets processing information in an inconsistent manner and also of significant lags in the processing of public information.

6. Discussion and Conclusions

In summary, there is evidence of some reaction in the betting market to the breaking information relating to the outcome of the conclave, but in retrospect a clear under-reaction in terms of both extent and speed of incorporation. We note in particular what we might term the ‘Galeazzi anomaly’.

In general, online person to person betting markets provide excellent potential for insiders to exploit any inefficiencies in the way the market adjusts to new information. The costs of entry into the Betfair market are low, simply a small up-front deposit into the trading account, and the market was reasonably liquid for clients of the exchange. However, it was not possible to trade from Italy with Betfair on the Conclave. This institutional feature may

explain, at least partly, why the market was slow to adjust to newly revealed information. That said, given that the relevant information was widely available and publicised in the UK where Betfair is much better known and very accessible, this is unlikely to be anywhere close to a complete solution.

What can we conclude from these findings? Either that the public release of accurate information from a generally well respected news source was not (for good or bad reason on this occasion) sufficiently believed by those able to place a trade on the exchange, or else it was significantly overlooked by traders. In other words, the betting market did not perform as well as might have been expected in terms of responding to new public information, which given its provenance and authoritative tone might have been expected to be accurate, and which turned out to be so. This cannot be explained in terms of the fog of conflicting signals as there were no other credible sources issuing conflicting information.

More generally, the main lesson we can perhaps take from this analysis is that decisions taken by individuals or groups of individuals whose choices are shrouded, at least from outsiders, in a layer of secrecy, may not always be as impenetrable as conventional wisdom about the secrecy of the deliberations might suggest. Rather, it is a question of knowing where to look for the information, and identifying which information to believe.

We venture that in the context of the most recent papal election, sufficient credible information was in the public domain to allow profitable exploitation of this information, yet markets failed to reflect this. To this extent, it reinforces the conclusion of Vaughan Williams (2014) about the relative predictive power of expert forecasters (American Bar Association) and a major real-money betting exchange market in identifying the high-profile outcome of the 2012 US Supreme Court decision on the ‘Affordable Care Act.’

This failure of the market to aggregate information in an effective way stands in clear conflict with the weight of evidence found in many other studies of prediction markets. We infer that the divergence in the efficiency of prediction markets in this context, as also found in relation to the high-profile 2012 US Supreme Court decision on the ‘Affordable Care Act’, can be attributed to the very different processes of evaluation used to discern and properly weight information flows in the context of open-door and closed-door decision-making.

We propose that further research might seek to examine whether the same clear divergence in market efficiency can be found in the context of other examples of open and closed-door decision-making.

References

- American Bar Association (2012), Preview of United States Supreme Court Cases, Special Edition, *Health Care and the High Court*, Preview, , 39 (6):, 1-40.
- Baumgartner, F.J. (2003), *A History of the Papal Elections*, New York: Palgrave Macmillan.
- Baumgartner, F. J. (1985), 'Henry II and the Papal Conclave of 1549', *Sixteenth Century Journal*, XVI (3): 301-314.
- Blackman, J., A. Aft and C.M. Carpenter, Jr. (2012), 'FantasySCOTUS: crowdsourcing a prediction market for the Supreme Court', *Northwestern Journal of Technology & Intellectual Property*, 10 (3): 125-167.
- Bollerslev, T (1986), 'Generalized autoregressive conditional heteroskedasticity', *Journal of Econometrics*, 31: 307-27.
- Brewer, J.S. (1920), *Letters and Papers, Foreign and Domestic, Henry VIII*, Volume 1, 1509-1514, London. Available at: www.british-history.ac.uk/report.aspx?compid=102701
- Calendar of State Papers and Manuscripts Relating to English Affairs Existing in Venice* [CSP], (1970), Liechtenstein: Kraus Reprints.
- Cherry, M.A. and R.L. Rogers, 2006. Tiresias and the Justices: using information markets to predict Supreme Court decisions, *Northwestern University Law Review*, 100 (3): 1141-96.
- Engle, R.F. (1982), 'Autoregressive conditional heteroskedasticity with estimates of the variance of U.K. inflation', *Econometrica*, 50: 987-1008.
- Epstein, L., W.M. Landes and R.A. Posner, (2010) 'Inferring the winning party in the Supreme Court from the pattern of questioning at oral argument', *Journal of Legal Studies*, 39 (2): 433-467.
- Fleishman, Jeffrey. Who the next pope will be is up for speculation. Los Angeles Times, April, 13, 2005.
- Fleishman, J. (2005), 'Who the next pope will be is up for speculation, Jeffrey Fleishman' *L.A. Times*, April 13
- Hunt, J.M. (2012), 'The Conclave from the "Outside In": rumor, speculation and disorder in Rome during early modern papal elections', *Journal of Early Modern History*, 16: 355-382.
- Johnson, P. (1974), *Pope John XXIII*, Boston: Little, Brown and Co.
- Johnson, T.R., R.C. Black, J. Goldman and S.A. Treul, (2009) 'Inquiring minds want to know: do Justices tip their hands with their questions at oral arguments in the U.S. Supreme Court', *Washington University Journal of Law and Policy*, 29: 241-261.

- Jottier, D., J. Ashworth and B. Hendels (2012), Understanding voters' preferences: how the electorate's complexity affects prediction accuracy and wishful thinking among politicians with respect to election Outcomes, *Kyklos*, 65 (3, August): 340-70.
- Padovano, F. and R. Wintrobe (2013), 'The dictatorship of the Popes', *Kyklos*, 66 (3, August): 365-77.
- Page, L. (2011), 'The ability of markets to predict conditional probabilities: evidence from the US Presidential campaign', in: *Prediction Markets: Theory and Applications*, ed. L. Vaughan Williams, London: Routledge, 123-136.
- Rhode, P.W. and K.S. Strumpf (2004), 'Historical Presidential betting markets', *Journal of Economics Perspectives*, 18 (2): 127-141.
- Rhode, P.W. and K.S. Strumpf (2013), 'The long history of political betting markets: an international perspective', *The Oxford Handbook of the Economics of Gambling*, ed. L. Vaughan Williams and D. Siegel, New York: Oxford University Press.
- Ruger, T.W., P.T. Kim, A.D. Martin, and K.M. Quinn (2004), 'The Supreme Court forecasting project: legal and political science approaches to predicting Supreme Court decision-making', *Columbia Law Review*, 104: 1150-1209.
- Sanuto, M., *I Diarii*, ed. F. Stefani (1903), 58 vols., Venice, 1879-1903, F. Visenti,
- Saville, B., H. Stekler and L. Vaughan Williams (2011), 'Do polls or markets forecast better? Evidence from the 2010 US Senate elections', *The Journal of Prediction Markets*, 5 (1), 64-74.
- Toman, J.T. (2004), *The Papal Conclave: how do Cardinals divine the will of God?*, University of Sydney. Available at:
www.international.ucla.edu/cms/files/PERG.Toman.pdf.
- Vaughan Williams, L. (2014), *Forecasting the 2012 US Supreme Court decision on the 'Affordable Care Act': a case study*, Mimeo.
- Villard, R. (2009), 'Le Conclave des parieurs: Paris, opinion publique et continuité du pouvoir à Rome au XVI^e Siècle', *Annales*, 64 (2, March-April): 375-403.

Figure 1a: Hourly prices for the 5 main Papal Contenders

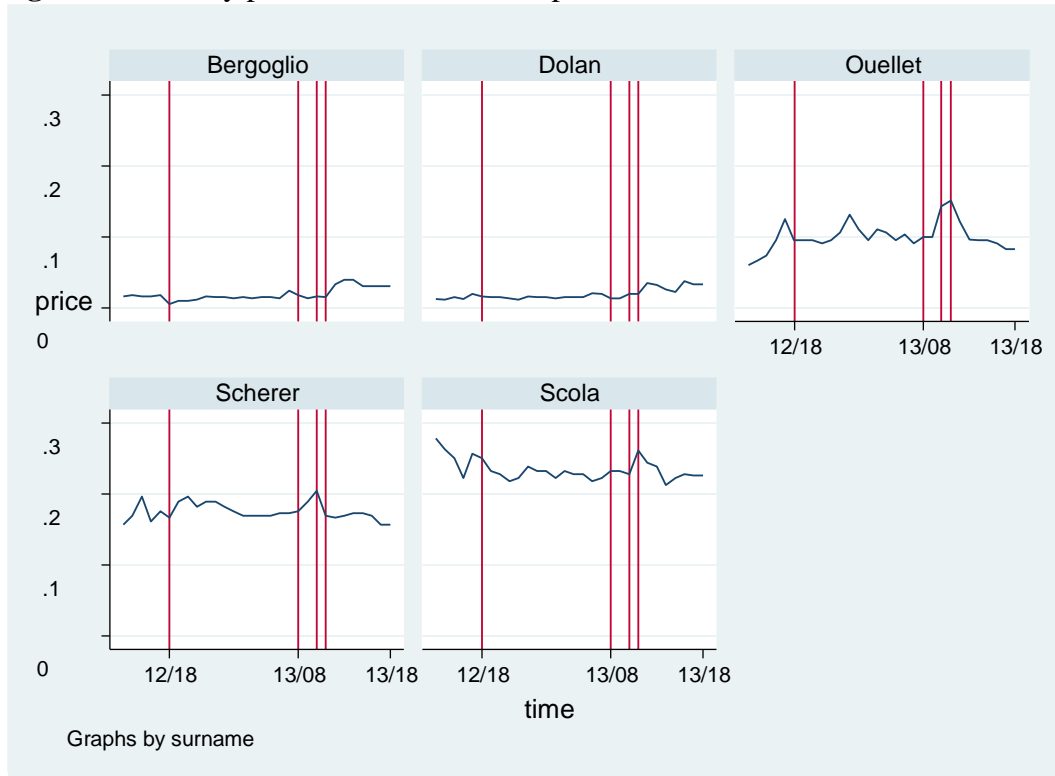
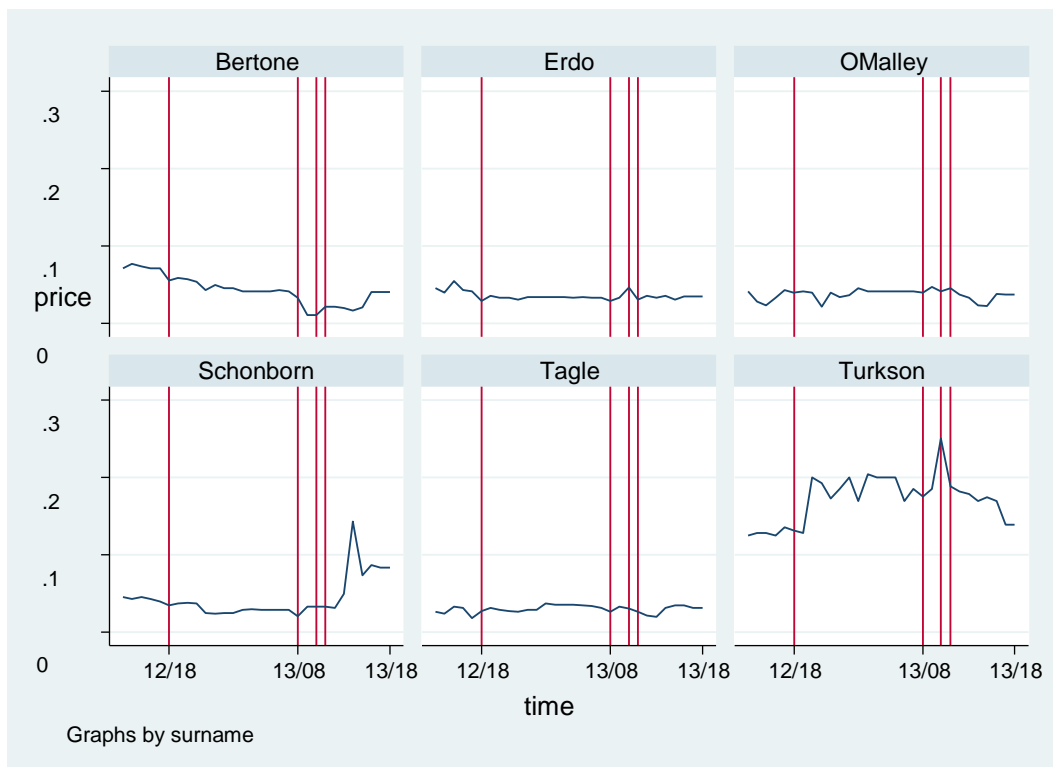


Figure 1b: Hourly prices for the selected other Papal contenders



Notes:

- (i) Graphs cover the period from 12:00 GMT on the 12th March 2013 until the closing of the market at 18:00 GMT on the 13th March 2013.
- (ii) Price is the probability of that candidate being elected as implied by the Betfair odds on matched bets.
- (iii) Vertical lines indicate the timing of each of the events A to D (as described in the text) in order.

Figure 2a: Bergoglio prices minute by minute

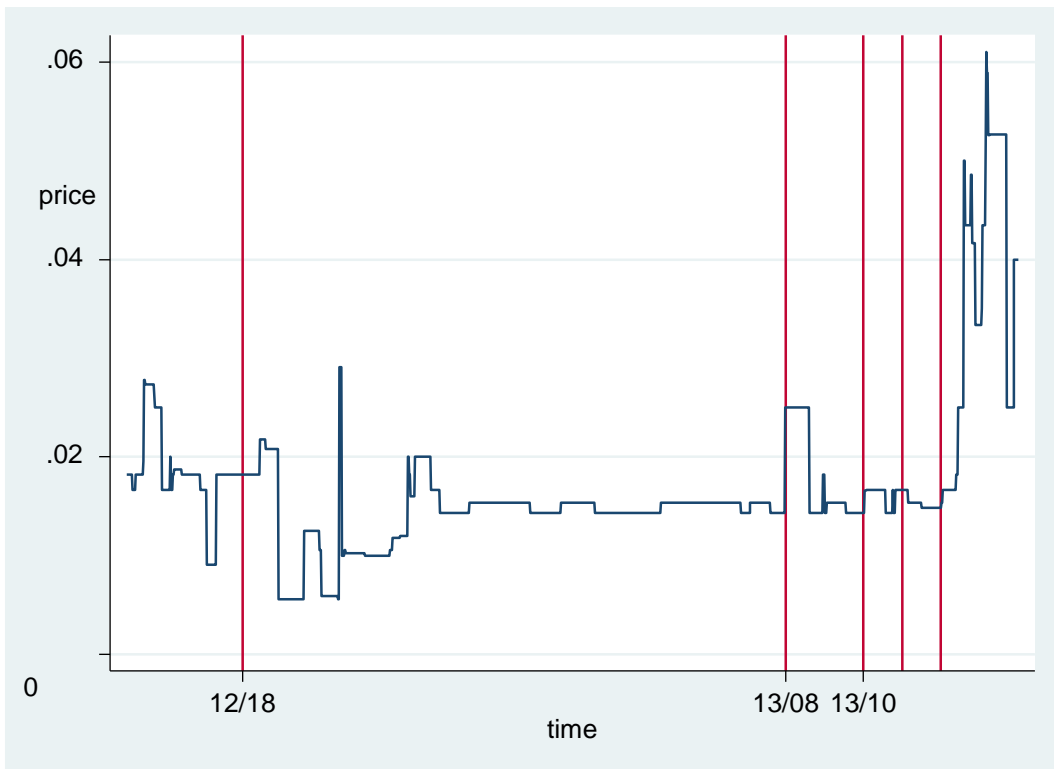


Figure 2b: Scola prices minute by minute

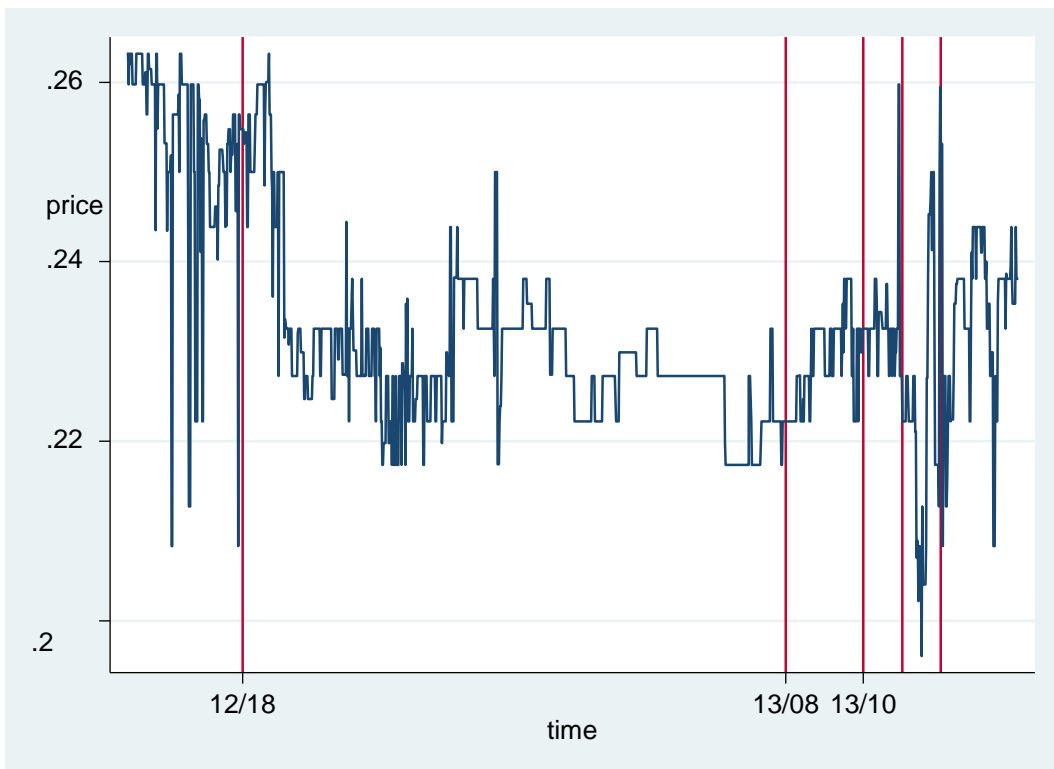


Figure 2c: Ouellet prices minute by minute

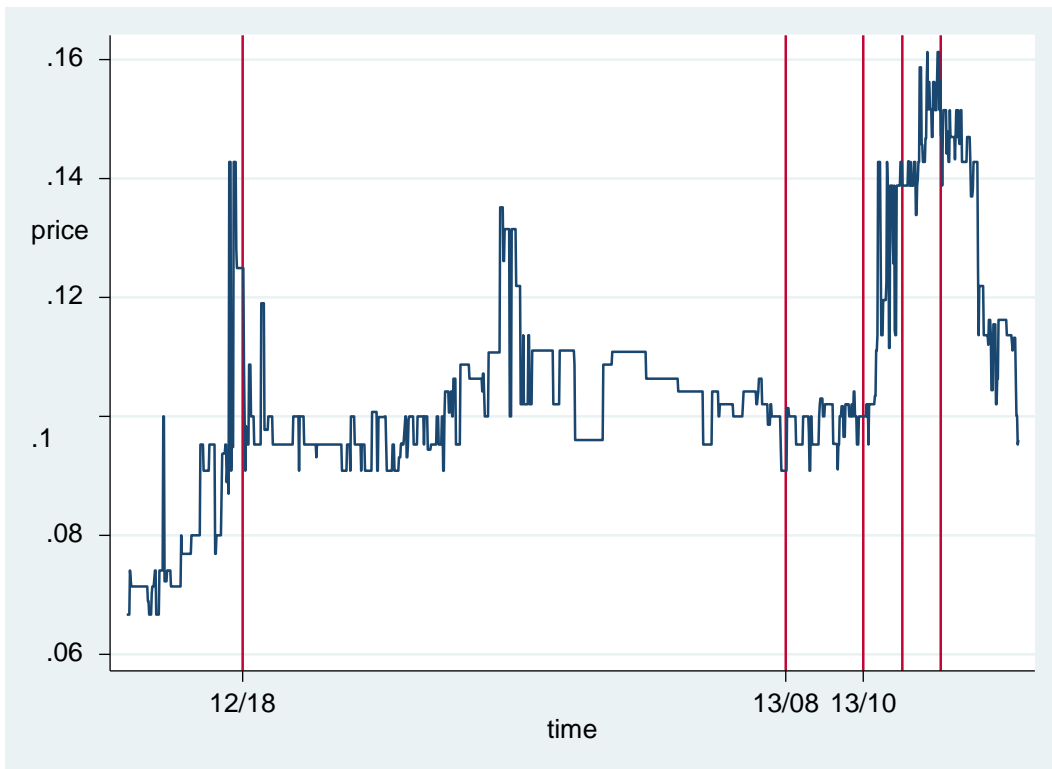


Figure 2d: Bertone prices minute by minute

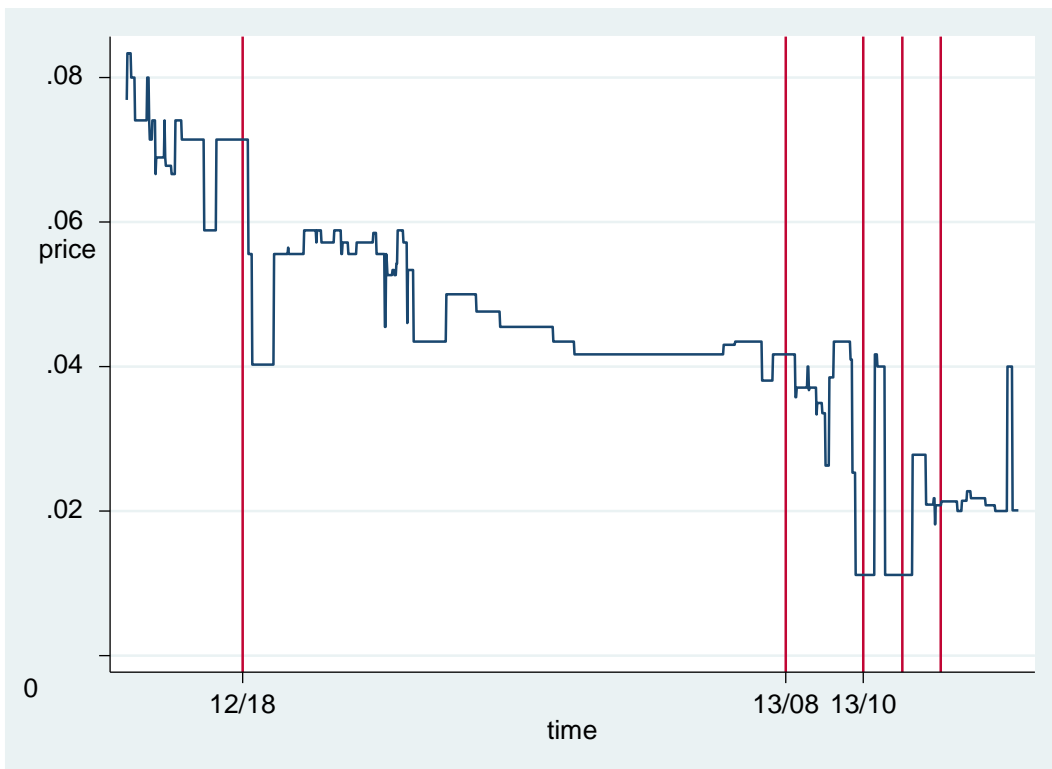
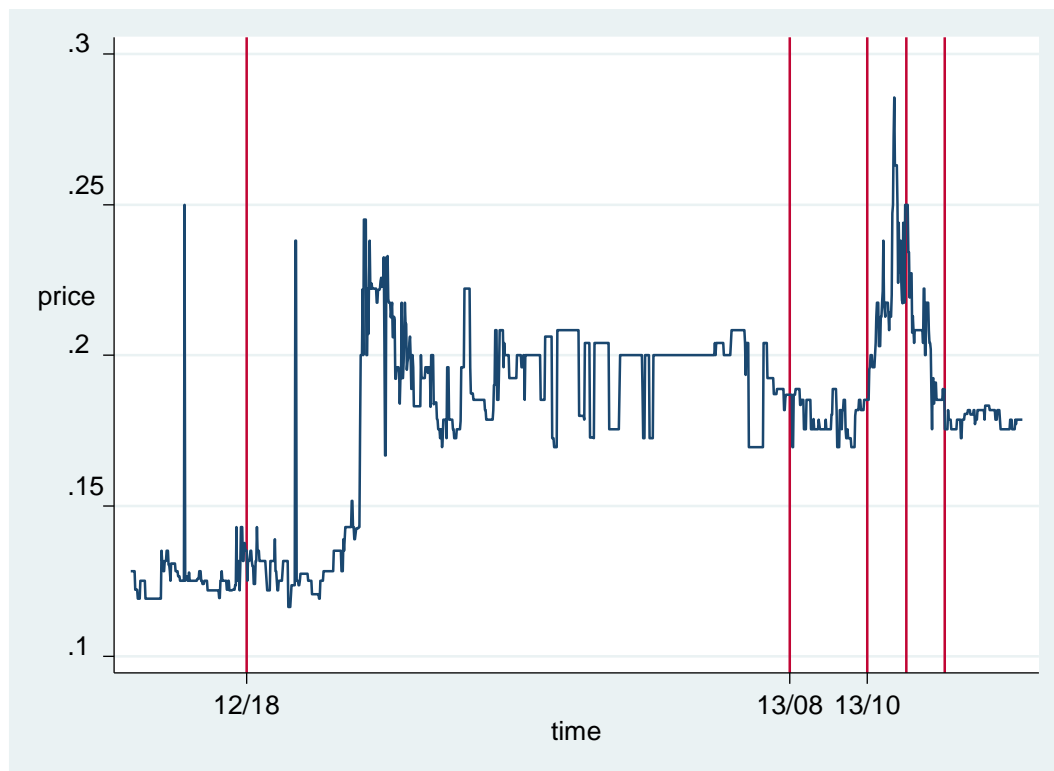


Figure 2e: Turkson minute by minute



Notes:

- (i) Graphs cover the period from 12:00 GMT on the 12th March 2013 until 14:00 on the 13th March 2013.
- (ii) Price is the probability of that candidate being elected as implied by the Betfair odds.
- (iii) Vertical lines indicate the timing of each of the events A to D (as described in the text) in order. The 5th line indicates the Guardian report which further disseminated the public information revealed at point D.

Table 1: ARCH Estimates of Returns in Papal Contender Markets

	1	2	3
	Bergoglio	Ouellet	Scola
Return (t-1)	-0.116 (0.130)	-0.283*** (0.087)	-0.334*** (0.111)
A. Black Smoke 1	-0.937*** (0.200)	-0.187*** (0.024)	0.022 (0.050)
B. La Stampa 1	-0.003 (0.171)	0.065*** (0.013)	0.053*** (0.005)
C. Black Smoke 2	0.121*** (0.034)	0.364*** (0.005)	-0.023*** (0.004)
D. La Stampa 2	-0.067*** (0.020)	0.127 (0.123)	0.131*** (0.003)
La Stampa 2 (t-1)	0.759*** (0.015)	-0.328*** (0.093)	-0.022 (0.017)
Constant	0.005 (0.008)	-0.008* (0.005)	1.233 e-4 (3.73 e-3)
Log-Likelihood	177.39	343.63	443.09
AIC	-336.78	-669.26	-868.18
N	310	310	310

Notes

(i) Sample period is 21.00 GMT on 28th Feb until 18.00 GMT on 13th March.

(ii) Dependent variable is the Betfair return for hour t, defined as $\log(P_t/P_{t-1})$ where P_t is the mean Betfair price at the end of hour t.

(iii) Figures in brackets are standard errors.

(iv) *** indicates significance at the 1% level; ** at the 5% level; * at the 10% level.

(v) AIC is the Akaike Information Criterion for model selection and is calculated as $AIC = -2(L - k)$ where k is the number of parameters being estimated. The AIC suggests a maximum lag order of one for the ARCH and GARCH terms.