ISM UNIVERSITY OF MANAGEMENT AND ECONOMICS

MASTER OF SCIENCE IN FINANCIAL ECONOMICS PROGRAMME

Vaidotas Statulevičius

MASTER'S THESIS

INVESTIGATION OF MARKET EXPECTED RETURNS AND VOLATILITY OF DJIA THROUGH ANALYSIS OF OPTIONS AND PREDICTIONS MARKETS

Supervisor

Dr. Asta Klimavičienė

VILNIUS, 2011

ABSTRACT

Statulevičius, V. Investigation of market expected returns and volatility of DJIA through analysis of options and predictions [Manuscript]: Master Thesis: economics.

Vilnius, ISM University of Management and Economics, 2011.

This paper aims to provide an innovative nonparametric approach to extract information about the equities market out of prediction markets. The procedure is carefully laid out and empirically tested on DJIA contracts in prediction markets. Results show that market expected DJIA values are not distributed normally or lognormally. A Jiang and Tian (2005) nonparametric implied volatility extraction method is applied to DJIA options. This paper also overviews prediction markets and their theoretical and empirical support.

Keywords: Prediction markets, Implied volatility, Options, Nonparametric modelling

SANTRAUKA

Statulevičius, V. Tikėtinų DJIA verčių ir grąžos svyravimų tyrimas analizuojant pasirinkimo sandorių ir prognozių rinkas [Rankraštis]: magistro baigiamasis darbas: ekonomika.

Vilnius, ISM Vadybos ir ekonomikos universitetas, 2011.

Šis darbas siekia pademonstruoti inovatyvų neparametrinį informacijos apie akcijų rinką išgavimą iš prognozių rinkos. Metodas kruopščiai aprašomas ir empiriškai patikrinamas naudojantis DJIA sandoriais prognozių rinkoje. Rezultatai rodo kad rinkos tikimasis DJIA verčių skirstinys near normalinis ar lognormalinis. Taip pat, naudojantis Jiang ir Tian (2005) neparametriniu grąžos svyravimų išgavimo metodu, surenkama informacija iš DJIA pasirinkimo sandorių rinkos. Šis darbas taip pat apžvelgia prognozių rinkas jų teorinius bei empirinius pagrindus.

Raktažodžiai: Prognozių rinka, Numanomas kintamumas, Pasirinkimo sandoriai, Neparametriniai modeliavimai

Contents

Conventional terminology and definitions

DJIA – Dow Jones Industrial Average

S&P100 – Standard and Poor's 100

DJX – options based on DJIA

B-S – Black-Scholes option pricing model

M-F – model-free (abbreviation used in discussions of implied volatility calculations)

Call option - gives the holder a right but not an obligation to buy an asset at a predetermined price (strike) at a predetermined future date (expiration date)

Put option - gives the holder a right but not an obligation to sell an asset at a predetermined price (strike) at a predetermined future date (expiration date)

In-the-money (ITM) options are options that would be exercised if the current asset spot price was observed at expiration

Out-of-the-money (OTM) options are options that would not be exercised if the current asset spot price was observed at expiration

At-the-money (ATM) options are located between ITM and OTM options with regard to strike price. In particular their strike price and current spot price are approximately equal (usually – option with the absolute lowest difference between strike price and spot price is called at-the-money)

Moneyness – ratio of strike price to underlying asset price, which describes how close an option is to being exercised. If moneyness is above 1, then a call option should be exercised.

CDF – cumulative distribution function. A CDF "describes the probability that a variate X takes on a value less than or equal to a number x" (Weisstein, 2011). See part *Probability density function and cumulative distribution function* for a more detailed description

PDF – probability density function. PDF maps variate X values to probabilities. See part *Probability density function and cumulative distribution function* for a more detailed description

Volatility – standard deviation of returns

For purposes of clarity and expedience, I use the following terminology to distinguish between different types of standard deviations:

> Predicted standard deviation – standard deviation calculated using prediction market data
>
> Implied volatility – volatility (standard deviation of returns) calculated using option data
>
> Black-Scholes (B-S) volatility – implied volatility calculated using B-S model from at-the-money options
>
> Model-free (M-F) volatility – implied volatility calculated from a range of option prices using Jiang and Tian implementation of Britten-Jones and Neuberger method
>
> Realised volatility – calculated from DJIA prices between some time *t* and the end of the sample

Introduction

There is a growing academic interest in prediction markets where contracts, contingent on uncertain future events, are traded (Wolfers & Zitzewitz, 2004). Traded future events range from sports results (e.g., Betfair) to more academically famous election results (e.g., Iowa Electronic Markets). These markets are different from simple betting in that there is no central bookmaker, that offers odds – market participants trade with each other.

Even though these markets have only recently become of interest to academic researchers, they were widely used for elections outcomes before the Second World War (Rhode & Strumpf, 2004). In this thesis a certain type of prediction contracts will be examined – forecasting Dow Jones Industrial Average (DJIA) value at a future date (operated by Intrade Inc.). The reason for this choice is twofold: first, the prediction can be compared to a much more liquid but similar market – options market, which will be also analysed. Second, most of the published research is concerned with elections and other events where there is one alternative source of predictions – subjective predictions of experts or general polls, therefore, this research will contribute to knowledge about prediction markets.

The aim of this thesis is to increase the knowledge about information available on different markets, specifically, what information on a particular equity could be extracted from options and prediction markets. The main objective of the thesis is to show that it is empirically possible to extract previously unrecovered information from prediction markets.

Predicted standard deviation of expected DJIA values at a future date will be extracted using prediction market data. Implied volatility estimates will be calculated from option markets and compared against realised volatility. Furthermore, the implied volatility measures will be evaluated for explanatory power of realisations of realised volatility using data from 2010.

This research will contribute to understanding how information available on different markets can be used and whether prediction markets bring any useful new information to volatility analysis.

Literature review section will give an overview of research done on the subject of prediction markets and a broader overview of knowledge on implied volatility calculations with specific insights by some authors. *Research problem definition* will define the research problem, present gaps in current knowledge and raise questions the thesis will try to answer. In *Research methodology* I will present the data used and actual proposed methods of extracting information from prediction and options markets. Next, *Empirical research results* will compare the results and check for accuracy using actual values of the DJIA. *Discussion* part discusses the results in the context of previous research. Finally, I summarise the paper.

Literature review

Literature review is structured in four parts: first, a short overview of risk-neutral probabilities, then a general overview of assumptions about asset price generating processes, followed by discussion of research on prediction markets, finally – more specific research on information in options markets.

Risk-neutral probabilities

Whenever empirically calculated probabilities are referred to in this paper they are risk-neutral probabilities, therefore I shortly present them in this part.

Risk-neutral probability measure assumes that investors are risk neutral with respect to return volatility, therefore, they optimise only risk and return trade-off. These probability measures are calculated by assuming that investors are risk-neutral and then finding probabilities that would solve for the observed asset prices – the result is risk-neutral probability (Jackwerth, 2000).

The problem with this assumption is that empirical research shows that investors are risk-averse (Jackwerth, 2000). Consequently, the risk-neutral probability measures will underestimate the likelihood of states of the world where out-of-the-money options are profitable, while overestimating in-the-money option probability of profit. This is because risk-averse investors will be willing to pay less than is “fair” (fair being equal to expected value under true probability measure) for deeply out-of-the-money options and pay more than is “fair” for options that are in the money, because of risk premiums attached.

However, Jackwerth (2000) shows that discrepancies between risk-neutral and real probabilities are small.

Asset price generating process and volatility

DJIA realised volatility for comparisons with implied volatilities will be calculated following a simple commonly used method, described by Figlewski (1997). This part will also describe the assumed asset price generating process.

Figlewski (1997) first describes the basic assumption on stock price movements used in the Black-Scholes model (discussed in more detail later in this section) – the lognormal diffusion model. The lognormal diffusion model postulates that continuously compounded returns follow a normal distribution:

$$\frac{dS}{S} = \mu dt + \sigma dz,$$

Equation 1

where *S* is the asset price, *dS* is the asset price change over an infinitesimal time interval *dt*, *μ* is the mean return at an annual rate, *dz* is a time independent random disturbance with mean 0 and variance *1·dt* (a stochastic process known as Brownian motion), and *σ* is the volatility, i.e., the standard deviation of the annual return." (Figlewski, 1997) Consequently, stock price levels follow a lognormal distribution (logarithm of *S* has a normal distribution) – *S* cannot fall below 0 and is positively skewed. Figlewski (1997) also shows that:

$\sigma_{annualised} = \sigma\sqrt{P}$, where *P* is the number of trading days in a year.

Equation 2

In short, if returns follow a normal distribution with constant mean and variance, then, given the daily continuously compounded return standard deviation *σ*, annual asset volatility can be found by multiplying *σ* by the square root of the number of trading days in a year.

Nevertheless, at this point it is reasonable to repeat empirical evidence against this process (taken from Figlewski (1997)):

1. Returns distribution varies with time (*σ* is not constant);
2. The process describes equilibrium prices, however, transaction prices are observed;
3. Price changes are correlated over time (especially at short time intervals);
4. Returns are not normally distributed – distributions display "fat tails" which are very unlikely if returns are normally distributed;
5. Volatility mean reversion (related to point number 1) – periods with high or low volatility are followed by periods with moderate volatility.

Despite the presented evidence against, assumption that stock price levels follow a lognormal distribution is a good estimate for the purposes of calculating Black-Scholes implied volatility.

Andersen and Bollerslev (1998) present evidence that realised volatility is best calculated using intra-day quotes at 5-minute intervals, however, this data is unavailable to me, therefore daily logarithmic returns will be used for estimation. In this work realised annual volatility until expiration is calculated following this formula:

$$\sigma_t = \sqrt{\frac{\sum_{i=t}^{T}(r_i-\bar{r})^2}{T-t-1} \times (T-t)},$$

Equation 3

where r_i is the DJIA daily logarithmic return and *(T-t)* is the time until expiration (*T* is expiration date, *t* is the current period), $\bar{r}$ is the mean logarithmic return over the period.

Problem of telescoping data

Hansen, Prabhala, and Christensen (2001) discuss a problem relevant for volatility estimations using options data. Their observation is that most of the studies using options data for volatility estimations (B-S volatility or implied binomial trees) try to use as much data as possible, which means using options with all strike prices and maturities. They find that using this data without proper techniques leads to non-converging regression coefficients and diverging *t*-statistics.

The cause of this problem lies in the use of the data – intervals to estimate realised volatility to regress against the implied volatility.

The cause of the problem is that almost the same data is used to calculate consecutive estimates of realised volatility. For example, let an option expire at $t = 252$, then current realised volatility (at $t = 0$) until the expiration of the option is calculated using returns for periods $t = 1$ *to* 252. Move forward one period ($t = 1$) – volatility is calculated using returns for periods $t = 2$ *to* 252. Continuing this procedure shows that realised volatility estimates are based on almost the same data if they are sampled daily for such a long interval; the "telescoping" part of the problem is due to the shortening of data series to calculate realised volatility.

Prediction markets

Three papers by Wolfers and Zitzewitz are presented in this section in addition to a paper by Rhode and Strumpf.

Wolfers and Zitzewitz (2004) present various prediction market contracts and shortly describe what information is available in those markets. The most simple prediction market contract is a “winner-take-all” contract where the contract costs $c and pays off $1 if a particular event occurs, e.g. DJIA is above 8,000 points on the 31st of December, 2010. For this type of contract the price $c is the risk-neutral probability that the event will occur – the aggregate market belief.

Index contracts pay $1 “based on a number that rises or falls, like the percentage of the vote received by a candidate” (2004) and they reveal market expectation of that number. Finally, Wolfers *et al.* (2004) explain the spread contract. In the context of elections, say, a contract costs $1 and pays out $2 if a candidate captures more than y% of votes, 0 otherwise. Market participants reveal what y would make them take the gamble with payoffs $2 and $0 – this contract reveals market expectation of the median. Median is a value such that 50% of values fall either side of the median; a risk-neutral market participant will only accept the spread gamble if he expects to win at least with 50% probability – the exact definition of the median. Any other percentile can be coaxed out of market participants by adjusting payoffs, e.g. spread contract costing $9 and paying out (in case of a “win”) $10 would reveal the 90^{th} percentile of the distribution.

Further in this paper the authors present evidence that prediction markets actually outperform other forecasts or that they are good predictors of the future. They draw on evidence from Iowa Electronic Markets, Saddam Security (discussed below), Hollywood Stock Exchange, and Economic Derivatives. They further find (in accordance with Rhode *et al.* (2004)) that arbitrage opportunities in prediction markets are, at least, difficult to discover. “Finally, the power of prediction markets derives from the fact that they provide incentives for *truthful revelation*, they provide incentives for research and *information discovery*, and the market provides an algorithm for *aggregating opinions*.”

Another paper by Wolfers and Zitzewitz (2006) presents a few simple models that show that this is not only true in practice, but also true in theory, less so "for prices close to $0 or $1, when the distribution of beliefs is either especially disperse, or when trading volumes are somehow constrained, or motivated by an unusual degree of risk-acceptance." (2006) The strength of prediction markets are that they are robust theoretically – prices giving market belief about probability of the event – and practically – prediction markets are more accurate than pre-election polls in determining the winner at all time horizons (Berg, Nelson, & Rietz, 2007).

Berg *et al* (2007) analyse US presidential and other elections since 1988 in order to evaluate the long-run efficiency of prediction markets. They use vote share contracts traded on Iowa Electronic Markets as prediction market forecast of the election result. Payoff of these contracts depends on the share of the vote captured by each candidate - $1 for every 1% of the share. As discussed by Wolfers *et al.* (2004), these index contracts reveal the market forecast of the vote share of each candidate. Berg *et al* (2007) compares these predictions and contemporaneous poll results (they mention a caveat that polls are designed to measure current voter sentiment, but they dismiss this because poll results are commonly used as election forecasts) against actual outcomes of the elections. They find that market prediction is better than poll forecast to the actual election outcome 74% of the time. They also find that for longer horizons (100 days to election at the margin) prediction markets are relatively more accurate than polls.

One criticism of this theoretical robustness is concisely presented by David Schneider-Joseph (2004) – if prediction markets are used as a form of insurance or hedging it might skew the prediction contract prices. For example, if a market participant believed that DJIA value at the end of 2010 on or above 10,000 negatively influenced his wealth (because he had entered a short position on the DJIA), he could purchase prediction contracts to insure against this potential loss (Schneider-Joseph, 2004). However, in practice this criticism is valid if trades are large enough to achieve this effect. Data suggests that this criticism is (at the moment and for this particular dataset) empirically unfounded: average trade over 2010 is $49, median – $26, largest trade – $4004 (second largest – $1360).

Most of the research in the field of prediction markets forecasts concentrate on non-financial events: elections (Rhode & Strumpf, 2004), wars (Leigh, Wolfers, & Zitzewitz, 2003), disease outbreaks (Nelson, Neumann, & Polgreen, 2006). However, election prediction markets were very important in the US until the Second World War and "betting activity at times dominated transactions in the stock exchanges on Wall Street" (Rhode & Strumpf, 2004). The authors find that "betting odds possessed much better predictive power than other generally available information" (these markets predated scientific polling) – decisive victories were predicted early, close races were anticipated. Furthermore, they find that betting odds were used as a source of information and outperformed massive non-scientific surveys, and that betting markets showed semi-strong market efficiency. Finally, they discuss how these markets lost popularity due to changes in legal environment (crackdown on organised betting and legalisation of pari-mutuel betting on horse races) and invention of scientific polls.

However, prediction markets have recently made a return with the advent of electronic markets. Rebranded as idea futures, event futures, information markets, or decision markets, prediction markets expanded in their scope to include Hollywood movie predictions (Hollywood Stock Exchange, 2011), weather and financial data (Intrade, 2011). As markets on financial outcomes are increasingly available – Intrade has 8 panels of contracts (a family of contracts that are based on the same event, but with different outcomes, e.g. EU unemployment rate to be 7 through 18 percent in 1 percentage point increments) on quantitative financial results – it becomes increasingly interesting to compare results from prediction markets and other sources.

A Leigh, Wolfers and Zitzewitz paper (2003) on prediction markets and Iraq war deal with interconnections between markets in a similar case as this thesis. In particular, they analyse how Saddam Security (contract depending on the event "Saddam Hussein is not President/Leader of Iraq by [Date]," traded on TradeSports) price varies with expert opinion, oil spot price and futures, equity spot prices and options. They also present evidence that prediction markets, even though not as "thick" (well-traded) as option markets, but, for Saddam Security in 2003, as well traded as Iowa Electronic markets, which Berg *et al.* (2007) showed to be empirically robust. It is not a great leap of faith to assume that since then the market volume and market participants' sophistication has increased, consequently, the prediction market data should be reliable.

The authors go on to use Saddam Security and other market data to establish various inferences about the effects of (at the time) potential Iraq war on oil prices and stock markets by sector and country in the short run and in the long run. In my opinion, these inferences are tenuous given that war would have multiple order effects on oil prices and stock markets, furthermore, using a single explanatory variable for variation of oil prices and stock prices is debatable. More detailed overview of research in this field can be found in Wolfers and Zitzewitz "Prediction Markets in Theory and Practice" (2006)

Options markets

Options are contracts that give the holder a right but not an obligation to buy or sell a particular security at a predetermined price and future date. For example, a call option on DJIA with a strike price of 10,000 and expiration date 2010.12.18 would give the holder a right but not an obligation to "buy" (options on DJIA are cash-settled – only money changes hands at expiration) DJIA at 10,000 on 2010.12.18. If, on this day, DJIA was higher than 10,000, then the holder would receive income from his investment, otherwise – he would lose the money he paid for the option (called premium), because it would not be advantageous to exercise it.

Black-Sholes model

In order to discuss implied volatility in more detail, first the Black-Scholes (B-S) option pricing model should be presented in broad strokes.

The whole Black & Scholes (1973) model is based on the idea of pricing the option by replicating it using the underlying asset and a riskless asset. In essence, the option is worth the same as a portfolio of assets, if the payoffs are the same. Theoretical derivation of the formula rests on this idea of perfect hedging.

First, the authors of the formula present their most important assumptions:

1. Known and constant short-term interest rate (same for borrowing and lending);

2. Asset return variance follows a normal distribution and has a constant variance;
3. Stock pays no dividends (this assumption can be relaxed, but is avoided in this study because dividend-adjusted closing prices are used);
4. Option is European style – it can be exercised only at expiration;
5. There are no transaction costs of buying or selling the option or the asset;
6. Assets are infinitely divisible;
7. No short-selling frictions.

Given these assumptions it is then possible to construct a hedged position, the value of which will depend on time and known constants:

1. Price of the underlying asset;
2. Strike price;
3. Risk free rate;
4. Time to expiration;
5. Volatility of the underlying asset.

Their solution to the problem of option valuation looks like this (for call options; put options can be valued by using the put-call parity (see part *Put-call parity*)):

$$C(S,t) = N(d_1)S - N(d_2)Ke^{-r(T-t)},$$

$$d_1 = \frac{\ln\left(\frac{S}{K}\right) + \left(r + \frac{\sigma^2}{2}\right)(T-t)}{\sigma\sqrt{T-t}},$$

$$d_2 = d_1 - \sigma\sqrt{T-t},$$

Equation 4

where *N()* is a cumulative distribution function of the standard normal distribution, *T-t* – time to expiration, *S* – spot price, *K* – strike price, *r* – risk free rate (continuously compounded), σ – annualised volatility until expiration date of the underlying asset.

The most important finding in Black & Scholes study (for this thesis) is that given that the first four factors (price of the underlying asset, strike price, risk free rate, time to expiration) in the formula and the option price are observable, volatility implied by the model can be calculated. This is an annualised measure of how the market participants expect the asset price to vary until expiration of the option.

Put-call parity

Put-call parity was used in the discussion of the Black-Scholes formula above and will be used further on in the paper, therefore I think it is useful to describe the concept in short. Put-call parity for the modern finance (there were some earlier derivations of the parity, however, they were forgotten and rediscovered only recently (Hafner & Zimmermann, 2004)) was derived by Hans R. Stoll (1969) – the following is an overview of Stoll's contribution. Put-call parity refers to the relationship between put and call prices – it is an arbitrage condition, and does not depend on any model of option pricing. The required conditions for put-call parity to hold are: it is possible to trade puts and calls, the underlying asset and earn interest. The put-call parity formula for European options with identical expiration dates and strike prices:

$$C(t) - P(t) = S(t) - K \times B(t,T),$$

Equation 5

where *t* is time, *C(t)* is call option price at time *t*, *P(t)* – put option price, *S(t)* – underlying asset price, *K* – strike price, *B(t,T)* – time *t* price of zero-coupon bond that pays $1 at time *T*.

$$B(t,T) = e^{-r(T-t)},$$

Equation 6

where *B(t, T)* – is time *t* price of zero-coupon bond that pays $1 at time *T*, *T* is time to bond maturity, *t* is current period and *r* is the interest rate.

In this study put-call parity will be used to calculate B-S implied volatilities (by restating Equation 4 to use put option prices instead of call options), and in model-free implied volatilities (see Equation 10).

Another important aspect of the put-call parity is used by Ait-Sahalia and Lo (1998) for capturing information in less liquid options. They use the parity to replace in-the-money call option prices with prices calculated using out-of-the-money put option prices, because the former are illiquid, while the latter are liquid. The authors explain that this way "all the information contained in liquid put prices has been extracted and resides in corresponding call prices via put-call parity; therefore, put prices may now be discarded without any loss of reliable information." (1998) This method would improve the power of the analysis because more data would become available; however, it is not usable in this research because in order to yield good results not only strike prices and maturities must be matched in the parity formula, but also the time of the option quotes.

Furthermore, Ait-Sahalia *et al.* (1998) also calculate implied index values using the put-call parity (see Equation 5 – $S(t)$ is the implied index value, if the formula is supplied with call and put option prices ($C(t)$ and $P(t)$, respectively), strike price (K), and zero coupon bond price ($B(t,T)$)). Using implied index values takes into account market expected dividend payments, which could also improve the analysis.

Unfortunately, time of the option quotes is not available to me at this time, consequently, I cannot match the required put option and call option prices; this means that implied index value cannot be calculated and information in illiquid options cannot be replaced with information from liquid options. However, Jiang *et al.* (2005) test the robustness of their analysis and find that discrepancies between implied and actual index values are very small – "[t]he mean and median of the absolute percentage difference are 0.12 and 0.08%, respectively" (2005), moreover, the percentage differences are symmetrically distributed.

B-S implied volatility

Implied volatility literature was reviewed by Mayhew (1995), where he presents various algorithms and weighting schemes for calculating implied volatility using B-S. Mayhew finds that the best implied volatility for direct use in econometric analysis should be calculated using near-the-money options, because the method is as good as using all options and some weighting scheme.

On the other hand, Mayhew (1995) also reviews research on the topic of “historic volatility versus implied volatility”. His findings are that historical volatility is less useful than implied volatility for forecasting purposes, with particularly promising results (and potential results) showed by time-series models that incorporate both types of data for forecasting purposes.

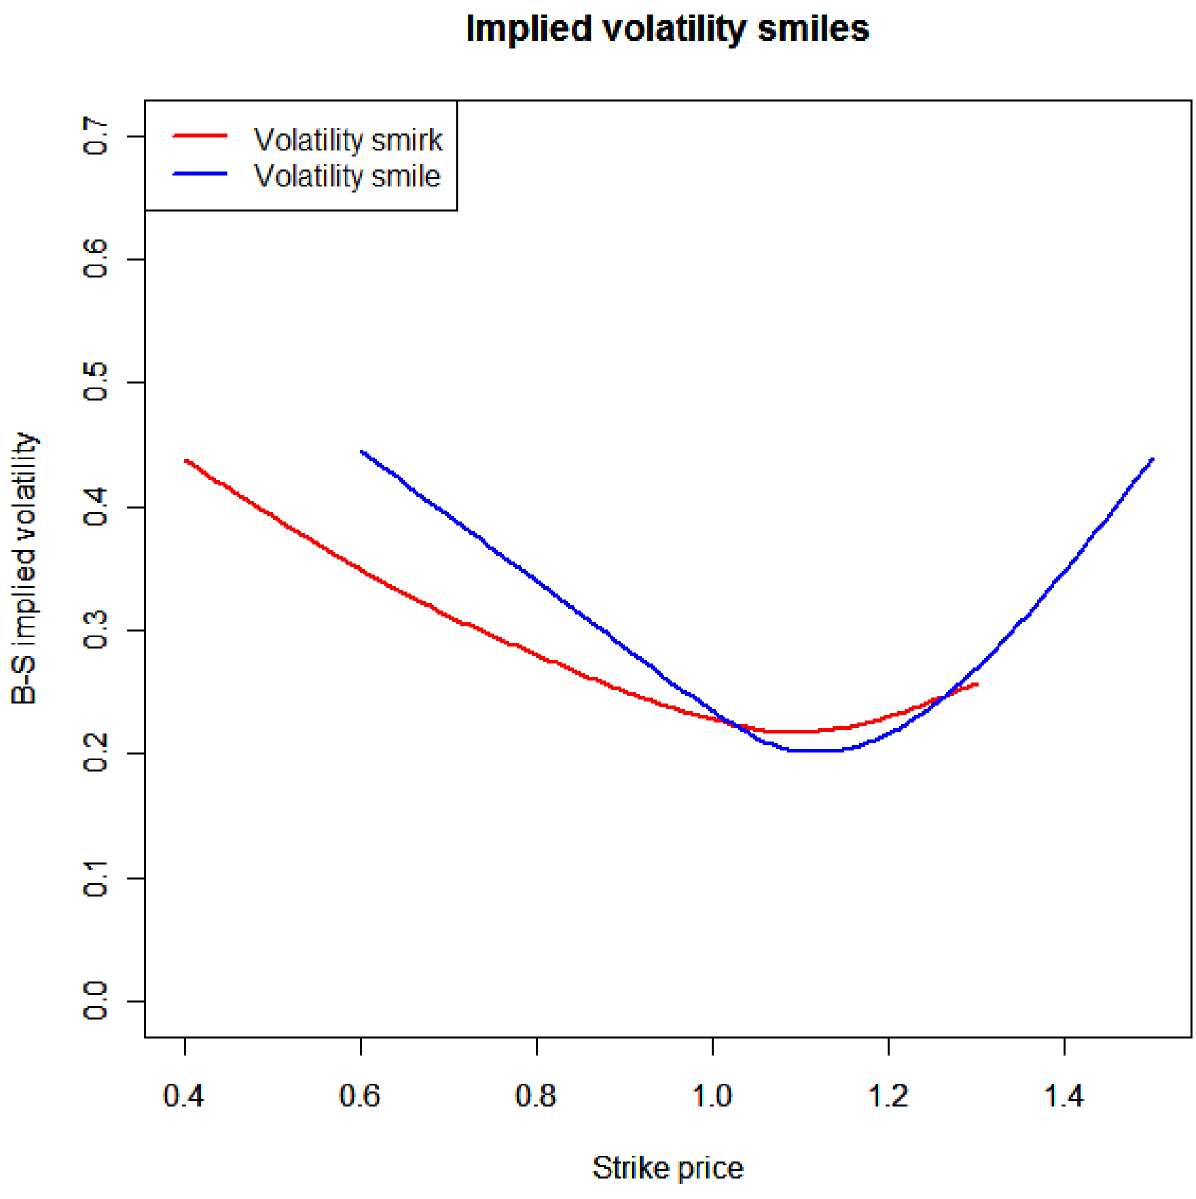

Figure 1

Furthermore, Mayhew discusses "the schizophrenic notion of the volatility smile" (1995). Volatility smile (Figure 1) refers to the empirical observation of the patterns of implied volatilities plotted across strike prices and times to maturity. The smile is observed when implied volatility (calculated using B-S model) is lowest for at-the-money options, symmetrically rising for out-of-the-money or in-the-money call options. Related concepts to volatility smile are volatility sneer and smirk – the former is a pattern when implied volatility from in-the-money call options is higher than for out-of-the-money; the smirk is the opposite of sneer (put options can be used to construct these "facial patterns" using put-call parity). Mayhew calls this a "schizophrenic notion" because the B-S model first assumes that volatility is constant and then is used to calculate different volatilities for the same asset using different strike price and exercise date options. Jiang *et al.* (2005) find that using at-the-money options to calculate implied volatility provides the most correlation with realised volatility and model-free implied volatility (discussed below). As an alternative solution to the volatility smile paradox Mayhew presents implied binomial tree option pricing models (discussed below).

Fleming (1998) discusses more recent practical application of B-S implied volatility calculations. He analyses American style (they can be exercised at any point until maturity) options on S&P100 index for the period October 1985 through April 1992. His findings are that B-S "implied volatility is an upward biased forecast" (Fleming, 1998), which means that it could be a reasonably good estimate, provided a linear model corrected for the bias. Furthermore, the author finds that B-S volatility is a superior (to historical volatility) forecast of realised volatility.

Fleming (1998) also discusses the underlying assumptions in greater detail. The most important assumption is that the market correctly forecasts the volatility of an underlying asset – this is corroborated by research testing forecasting accuracy of market estimates of volatility, e.g. Blair, Poon, and Taylor (2000) use VIX implied volatility index to test whether it correctly forecasts future realised volatility and find that indeed it does. Given that even B-S implied volatility model gives a better forecast than models based on historical volatility, we assume that market is better informed than econometric models and make a short leap from this to an assumption that market knows the "true" future volatility. Based on this assumption Fleming

(1998) tries to discern the market forecast through B-S implied volatility. He assumes two things: options market is efficient – i.e. the “true” (unbiased) forecast is revealed there, and that B-S model holds. These two assumptions are intertwined because the “true” forecast of asset volatility cannot be observed, so we try to model it; in this case – using the B-S option pricing model. However, the B-S model, as explained above, has numerous drawbacks – the volatility smile/smirk/sneer when plotted against strike prices and time to expiration; periods where observed option prices diverge from model predicted by more than at other times (compare B-S goodness-of-fit before 1987 crash and after). I try to avoid this problem by also using model-free implied volatility (discussed below).

Implied binomial tree volatility

Rubinstein (1994) developed a different approach to option pricing from the B-S model. His approach is to start from the market prices of options and find the risk-neutral probabilities equalising the outcomes and prices, then, using these probabilities, a “unique fully specified recombining binomial tree” (Rubinstein, 1994) is constructed. Solving the entire tree provides a wealth of information about the probable paths of asset prices, including the volatility. The main advantage of this binomial volatility is that it does not have to hold any assumed form and can vary over the life of an option.

Dumas, Fleming, and Whaley (1998) examine how the binomial tree implied volatility holds up against B-S implied volatility and find that binomial implied volatility is very good in-sample (because it allows volatility function to take any form), but deteriorates when tested out of sample. The authors find that more parsimonious models produce lower errors out-of-sample and B-S implied volatility is more reliable because of its simplicity. Therefore, I do not attempt to replicate the implied binomial tree volatility in this research.

Model-free implied volatility

Britten-Jones and Neuberger (2000) present a model for extracting market forecast of volatility that is theoretically more general than the implied binomial tree model. The problem with binomial tree volatility is that, even though it is much less strict than B-S model, it still assumes deterministic volatility. Britten-Jones *et al.* (2000) describe deterministic volatility as an assumption that “instantaneous volatility is a function solely of the current stock price and time and therefore implies that all claims can be hedged by the underlying and the riskless assets” (2000). They further refer to evidence that deterministic volatility is an assumption not upheld by empirical evidence.

Moreover, as discussed by Jiang *et al.* (2005), model-free (M-F) volatility has a few other attractive features compared to previously commonly used B-S implied volatility. First, it does not require any assumptions about the option pricing model (hence, model-free). Second, M-F volatility calculation is based on all strike prices of the options, whereas best B-S estimates of implied volatility are generated based on at-the-money options (as discussed in *B-S implied volatility*). A third related advantage is that testing M-F volatility means directly testing market efficiency, as opposed to a joint test of a model and market efficiency.

Britten-Jones *et al.* (2000) method is based on calculation of risk-neutral probabilities using the results derived by Breeden and Litzenberger (1978) – risk neutral density is the second order derivative of the call price with respect to strike price (given that it exists and the call price function is continuous).

Britten-Jones *et al.* (2000) are quite theoretical in their approach, because it requires some data which is not available in the real world, e.g., continuum of strike prices from minus infinity to positive infinity. However, based on their work Jiang *et al.* (2005) describe an approach that is quantitatively feasible – data is available, assumptions are empirically tenable. Their solution is the basis for implied volatility calculations in this work and their method is presented in more detail in *Model-free risk neutral implied volatility calculation method*.

Jiang *et al.* (2005) also generalise Britten-Jones *et al.* (2000) approach – they show that it is not necessary to require that stock prices follow a diffusion process (prices follow a continuous path with no breaks), but that the same results can be achieved even if prices follow a diffusion with jumps process (continuous change disrupted in some places by discontinuous breaks).

Another very important result in the Jiang *et al.* (2005) paper is that they produce theoretical bounds of the errors arising from the actual implementation of the method. Specifically, they derive formulae for truncation errors (errors arising from relaxing the requirement for strike prices to range from minus infinity to plus infinity) and discretization errors (numerical integration is the source of these errors). They also transform the original Britten-Jones *et al.* (2000) specification requiring forward prices to use spot prices of the options.

Furthermore, Jiang *et al.* (2005) test their method on S&P500 index options[1] (these are very similar options to DJX – traded on CBOE, index options, high liquidity; for more information please see the original paper). They calculate volatility estimates using B-S method, model-free and realised volatility. In this case the benchmark is the realised volatility, because it is the realisation of the volatility that both methods are trying to forecast.

The authors find that both measures consistently overestimate realised volatility (this could be related to the fact that both forecasts assume risk-neutrality), but the relevant correlation[2] coefficients for B-S volatility and model-free (M-F) volatility with realised volatility are 81.4% and 85.7% respectively. This shows that M-F volatility is a better forecast than B-S volatility, and further tests by the authors reveal that once M-F volatility is included in the model, B-S volatility becomes statistically insignificant explanatory variable. Comparing univariate regression results of realised volatility on B-S, M-F and lagged realised volatility, the authors find that M-F has the greatest explanatory power (R^2 is between 73 and 74% for regressions of volatility, variance and logarithm of variance), followed by B-S (R^2 between 60 and 64%), lagged realised volatility has the least explanatory power (R^2 from 11 to 20%).

Authors also test various multivariate specifications of the regression involving M-F, B-S and lagged realised volatility factors. Their results suggest that M-F volatility "fully subsumes information contained in both the B–S implied volatility and the lagged realized volatility and is a more efficient forecast for future realized volatility."

---

[1] Symbol: SPX

[2] Reported correlation coefficients are Pearson product-moment correlation coefficients.

Research problem definition

Review of the relevant literature reveals a gap in research of prediction market financial events forecasts – I am unaware of research focused on empirical analysis of predictions of financial events. As presented in *Literature review*, prediction market research is mostly concerned with elections and other non-financial events – this is so probably because there are established financial markets dealing in financial predictions – options and futures markets – which are well researched. As a result, I believe this under-researched field is very promising because of two reasons: first, prediction market predictions can be compared and complemented by predictions generated in other markets (e.g., options); second, prediction markets at the moment are small relative to options markets (2010 trade values in US dollars $156 thousand versus $651 million, respectively), which reduces their attractiveness as a hedging instrument. Given the arguments discussed above, research problem – determining market predictions of future financial events in prediction and options markets – leads to research question "what information about future financial events is available in prediction and options markets." In order to answer this question a few hypotheses will be tested:

1. Stock, prediction and option market prices follow random walk. This is important because random walk is one of the basic indicators of market efficiency – essential assumption for most financial calculations.
2. Implied volatility is an estimator of realized volatility – test whether any information is revealed in implied volatility.
3. Prediction markets expect DJIA values to follow Gaussian distribution – testing whether a commonly used assumption (Figlewski, 1997) is shared by the market.

Research objectives are:

1. Calculate predicted standard deviations, implied and realised volatilities from acquired data;
2. Construct and analyse probability density functions of DJIA values from prediction market data;
3. Test how much information implied volatility contains about realised volatility;

4. Determine whether market predictions (options and prediction contracts) differ in available information.

In summary, this research will contribute to understanding how information available on different markets can be used and whether prediction markets bring any useful new information for standard deviation analysis. Furthermore, this thesis aims to enhance understanding of relationship information available on stock markets, options markets and prediction markets – what the expected values of stocks at certain times are and what the distributions of these values are. Due to the fact that prediction markets are the least researched and the newest in this case, this paper will place more emphasis on them.

Research methodology

This section deals with research methodology – how and what should be done to arrive to the goal of the thesis. First, two concepts used in the calculations are presented. Then data sources and main characteristics are overviewed. Third, the building blocks of the research are presented: how prediction market data will be analysed and what is expected to be uncovered; option market data analysis tools are discussed and how they will be used; and lastly – how results from prediction markets and option markets will be analysed to find out what information is available and if it is complementary.

Probability density function and cumulative distribution function

This part will shortly discuss the nature and concept of probability density functions and cumulative distribution functions because they are often used in the analysis and discussions of results.

Basically, a probability density function (PDF) is a mapping of probabilities to values of a random variable. For example, a throw of a fair uniform die has a discrete uniform probability density function – each of the 6 potential outcomes has an equal probability. The situation is presented graphically in Figure 2 – red dots represent probabilities of different outcomes. Because all the outcomes are equally likely, the dots are on the same level. The information shown on the graph can be expressed as a function – the probability density function.

Cumulative distribution function (CDF) is very similar to probability density function – it maps probabilities that a random variable will take a value equal to X or less to values of random variable X. For example, Figure 2 shows that the probability that a roll of a die will produce a three or less is equal to 50%. Actually, these two functions are two sides of the same coin – area under a PDF up to some X is the CDF value at X.

It is useful to generalise from discrete cases to continuous cases because they better represent reality and are easier to manipulate mathematically. As discussed in part *Risk-neutral probabilities*

Whenever empirically calculated probabilities are referred to in this paper they are risk-neutral probabilities, therefore I shortly present them in this part.

Risk-neutral probability measure assumes that investors are risk neutral with respect to return volatility, therefore, they optimise only risk and return trade-off. These probability measures are calculated by assuming that investors are risk-neutral and then finding probabilities that would solve for the observed asset prices – the result is risk-neutral probability .

The problem with this assumption is that empirical research shows that investors are risk-averse. Consequently, the risk-neutral probability measures will underestimate the likelihood of states of the world where out-of-the-money options are profitable, while overestimating in-the-money option probability of profit. This is because risk-averse investors will be willing to pay less than is "fair" (fair being equal to expected value under true probability measure) for deeply out-of-the-money options and pay more than is "fair" for options that are in the money, because of risk premiums attached.

However, Jackwerth shows that discrepancies between risk-neutral and real probabilities are small.

Asset price generating process and volatility it is common to assume continuous processes (consequently, continuous probability density functions) when dealing with financial time series, even though discrete observations (discrete in time and accuracy – e.g. two significant digits for most prices) of financial data would suggest discrete probability density functions. However, it is unreasonable to assume that prices of financial assets exist only at times when they are observed (a transaction occurs).

In short, probability density functions and cumulative distribution functions describe how probable it is for a random variable to take some value.

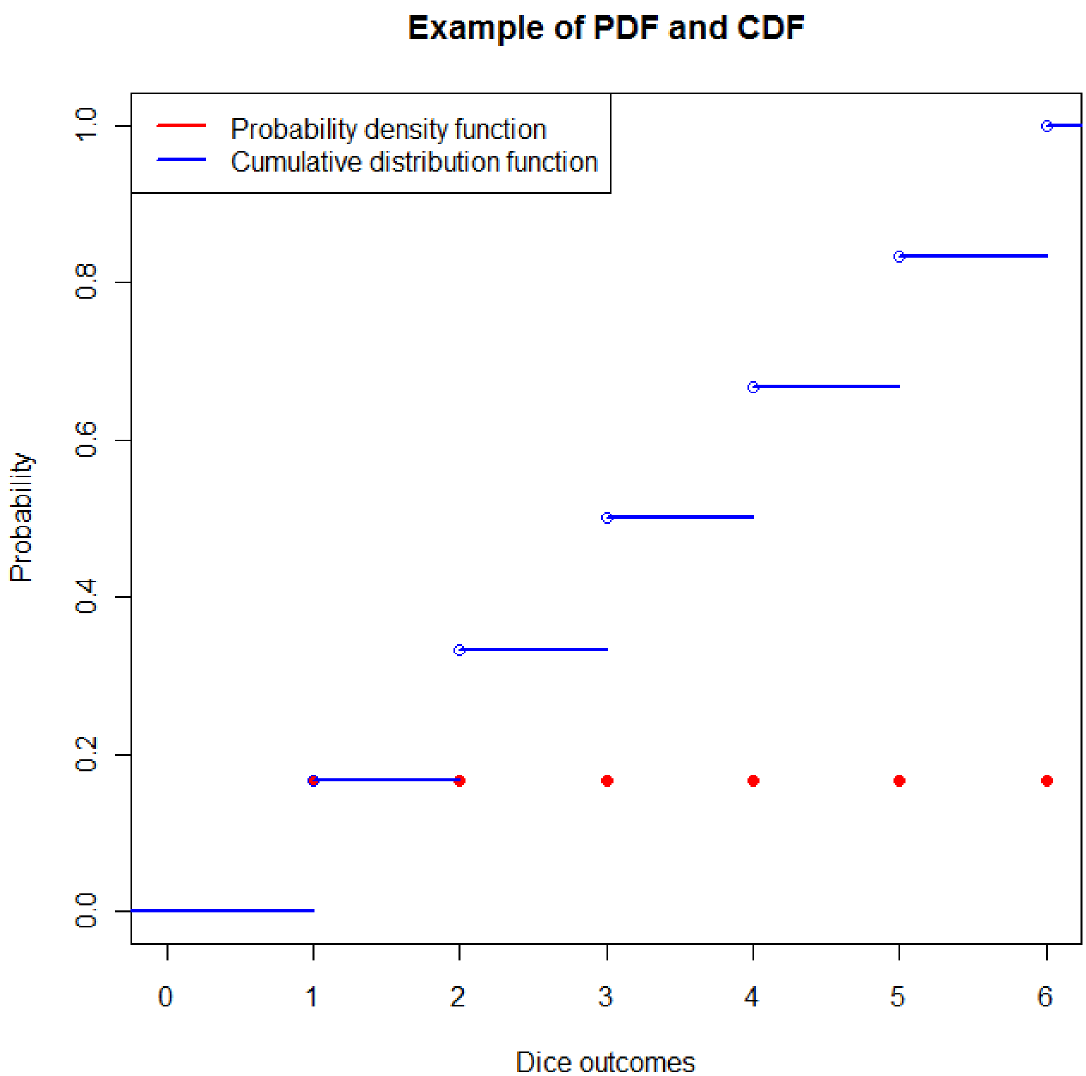

Figure 2

Cubic spline

Cubic splines are used in several steps of this research therefore this short section is dedicated to presenting the rationale behind the method and its choice.

Historically, a spline is an engineering concept of fitting an elastic ruler over some set of points – this physical process is replicated mathematically. The general principle of cubic spline interpolation is that for two adjacent observed data points a cubic (or quadratic) function is fitted so that both data points are the endpoints of the function. Further requirement is that both first and second order derivatives of adjacent splines at the observed data points are equal. This procedure ensures a smooth and continuous curve is fitted through all observed data points (McKinley & Levine, 1998).

In this research I use two types of spline interpolation. First, I use monotonic smoothed quadratic spline interpolation (functions created by Meyer (2010)) to interpolate the cumulative distribution functions (CDFs). “Smoothed” means that the splines are fitted not exactly to the data points, but as a best fit line – balancing a smooth curve and a good fit for the data points. Monotonicity is required because CDFs are by definition monotonically increasing.

Spline interpolation is also used in model-free volatility calculations (more on that further on) – in this case I use smoothed cubic spline method because monotonicity is not required and in order to reduce the number of restrictions on the shape of the function.

Other spline interpolation methods were also tested and compared against – non-smoothed and non-monotonic – but were discarded because they decrease the smoothness of the curves without greatly improving other attributes and also sometimes create negative probability density functions (PDFs).

Different interpolation methods were also evaluated but were discarded. Smoothed cubic spline interpolation method was chosen above piecewise constant interpolation and linear interpolation because of its smoothness property, which seems theoretically more reasonable than others (rarely do natural phenomena follow linear or non-continuous functions). Polynomial interpolation of higher orders (11^{th} order polynomials would be required to interpolate the cumulative distribution function from prediction market data) have an unfortunate property of high variation between any two data points (McKinley & Levine, 1998), also known as Runge’s phenomenon (Fornberg & Zuev, 2006), therefore, it was also passed over.

The fitted curve is then used to interpolate unobserved values – calculate values that were not directly observed in the market.

Data

Sample period for this study is year 2010 – there are two reasons for this. First, a recent period is more interesting to analyse. Second, the choice of a sample period is constrained by available prediction market data because more liquid contracts are available for end of the year dates and they are traded somewhat often since the end of the preceding year.

All three analysed markets are shortly presented and main data characteristics laid out later on.

Interest rates

Before further discussions of the markets, main interest rate data will be presented, because interest rates are used in various formulae. We use two main interest rates for calculations: risk-free rate is required by the B-S model and bond rate is required by the M-F volatility calculation procedure.

By definition, risk-free interest rate is the rate of interest paid on assets that have no uncertainty – cash flows from these assets are known with 100% certainty (Brigham & Houston, 2008). However, in the real world this strict condition cannot be fulfilled because all issuers of assets can default, therefore, we use the next best thing – assets that have as little risk as possible. Despite recent negative outlook, United States government is considered the most reliable debtor in the world due to the economy's size, strength and historical track record of meeting obligations to its creditors, consequently, US short term Treasury bills are considered as a risk-free asset in this work. Yield on these treasury bills is a proxy for the risk-free rate.

Figure 3 shows that the risk-free rate is quite variable. For example, 28 days rate over 2010 varied from as low as 0.01% to as high as 0.18% - a difference of 18 times (the more relevant 360 days rate varied from 0.2% to 0.48% - more than twice). As explained in part *Black-Sholes model*, the model assumes that risk free rate is constant; however, Figure 3 is in contradiction to this assumption.

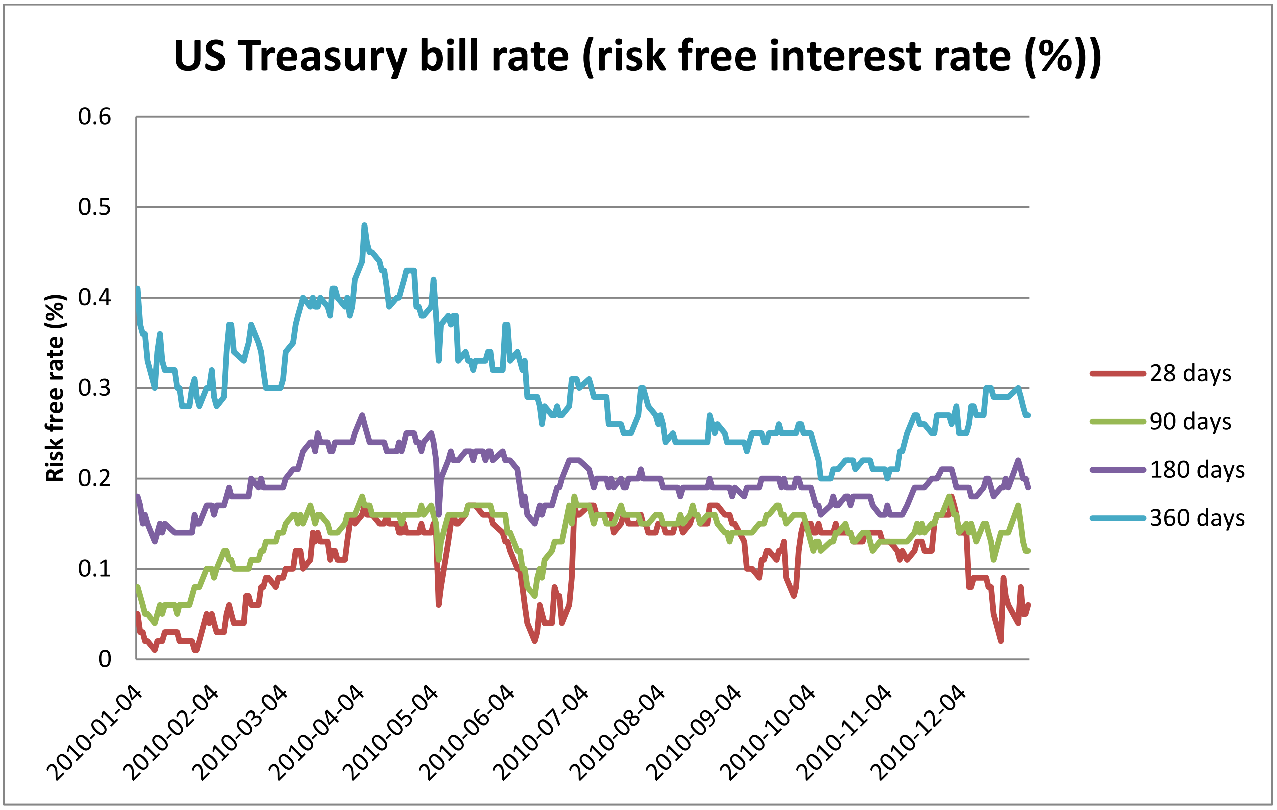

Figure 3

US Treasury bill secondary market yields (The Federal Reserve Board, 2011) for 1-year maturity will be used as the risk free rate. Secondary market was chosen because primary market yield publications were discontinued by the Federal Reserve Board (FRB). 1-year maturity was chosen, because the B-S model uses annual risk free rate in the formula, which is then adjusted for the time to expiration of the option. Another possible approach would be to use all available Treasury bill maturities – 4-week, 3-month, 6-month and 1-year maturities – to interpolate (possibly using cubic splines) for an exact time to option expiration, annualise this risk-free rate and use it in the B-S formula. This approach deals with time to maturity in a more robust way, but as time to expiration approaches, the annualised risk-free rate becomes more unstable due to small fluctuations leading to large changes when annualised. Consequently, I use 1-year maturity Treasury bill yield as a proxy for the risk-free rate. The data series is presented in Figure 4.

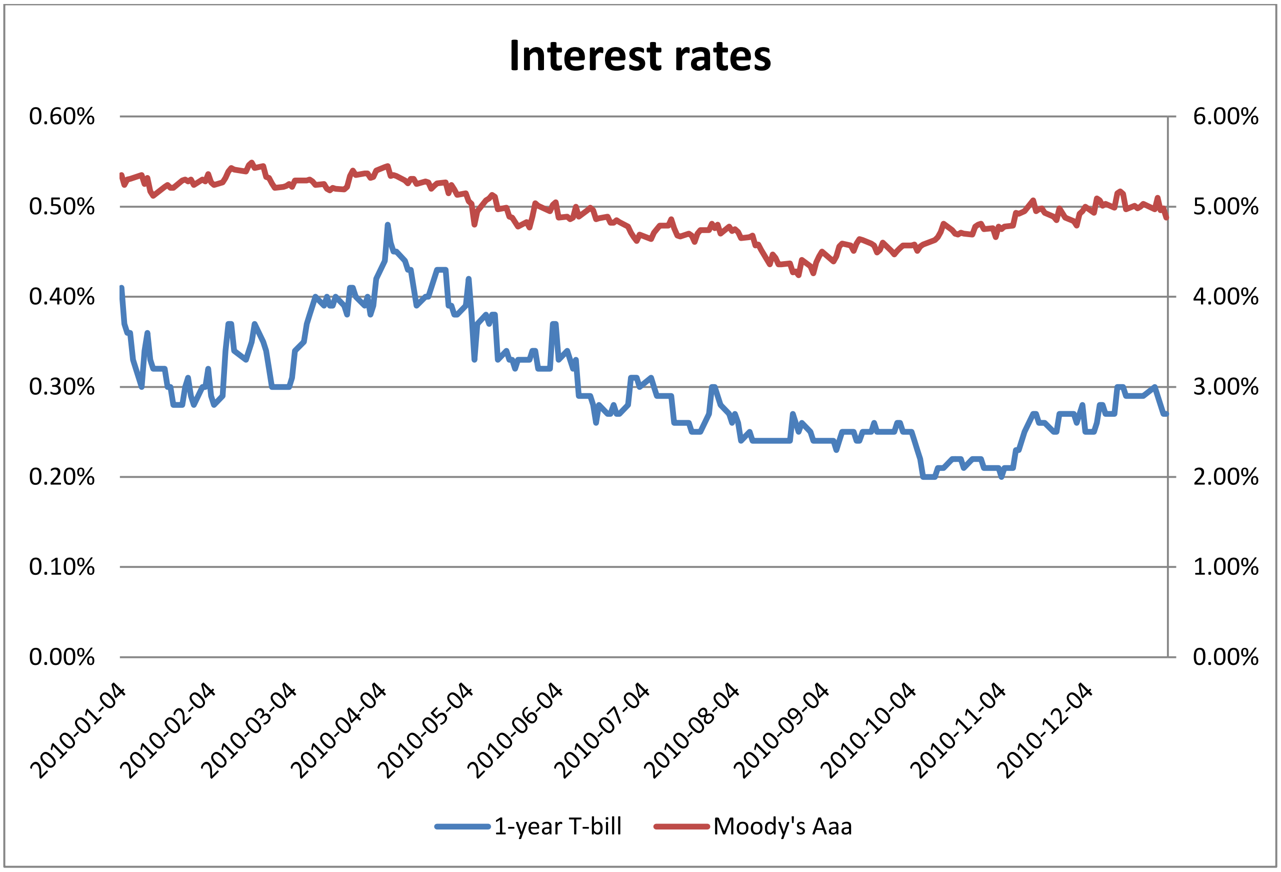

Figure 4

M-F volatility calculations also require corporate bond rate due to transformation from using forward prices to spot prices. Given that DJIA is composed out of the best performing companies in the US, Moody's Aaa corporate bond yields were chosen to represent this rate (see Figure 4, axis on the right). Moody's Aaa corporate bond interest rates were acquired from the Federal Reserve Board website (2011).

Risk-free rate and Moody's corporate rate do not diverge; however, T-bill rate is much more volatile than the corporate bonds rate – relative standard deviation of 21% against 6%, respectively. Risk-free rate varies from a low of 0.2% to a high of 0.48% (average – 0.3%), corporate – 4.24% to 5.49% (average – 4.94%).

Stock market - Dow Jones Industrial Average

Daily DJIA data was retrieved from Yahoo Finance (2011) – 252 data points – one for every trading day of 2010 (NYSE Euronext, 2011). Dataset contains the following data: date, opening, high, low and closing prices, trading volume.

Figure 5 shows the path of DJIA adjusted closing price over 2010 – starting at 10,584 points DJIA ended 2010 at 11,578. There were two periods of stable growth – February to late April and September to year-end. Over the summer DJIA suffered relatively high "negative" volatility – successive decreases in DJIA price (price increases, which also increase volatility, are usually not associated with increased risk because they are beneficial for the investor). This happened possibly due to turmoil in Europe concerning the future of the Eurozone. As these worries declined due to determination by the EU to support the monetary union, DJIA regained strong growth.

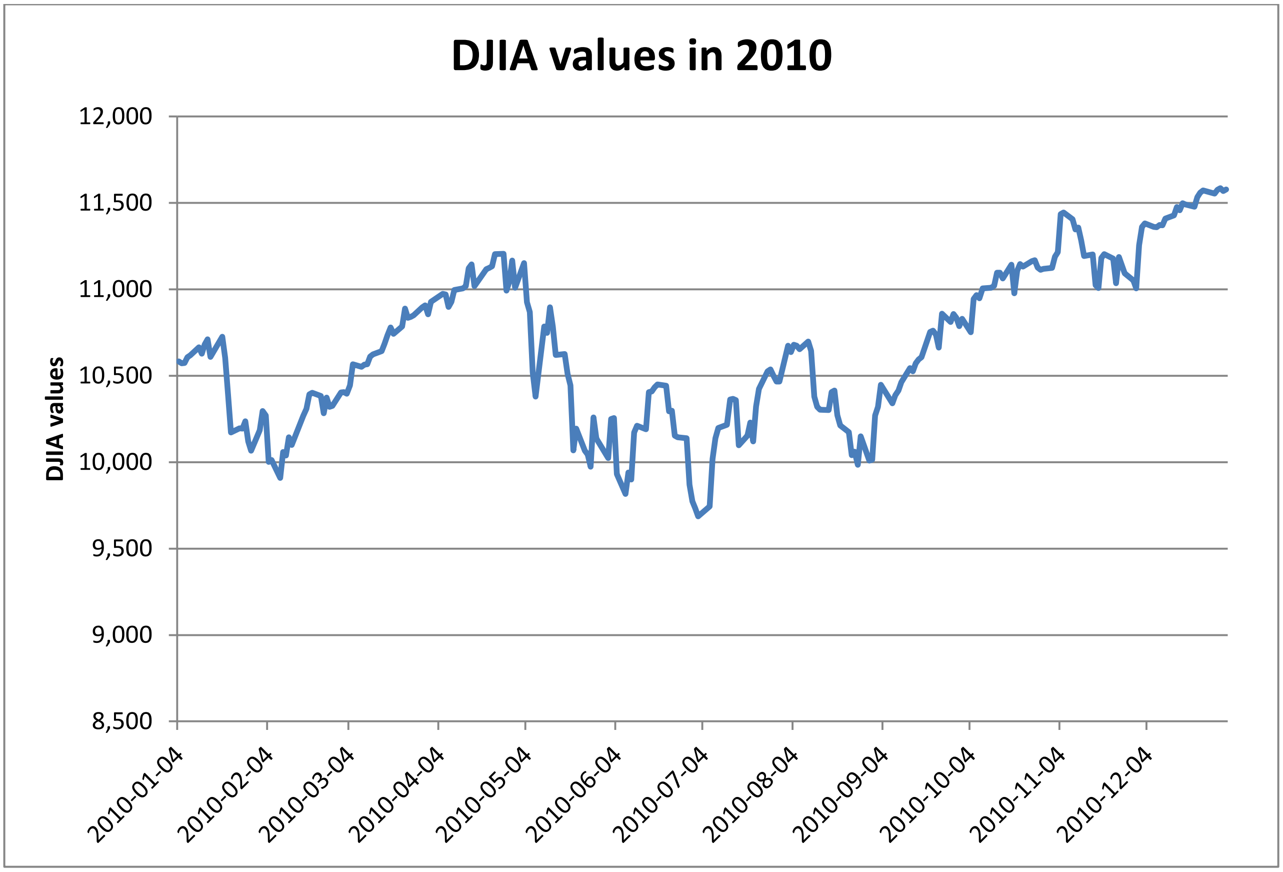

Figure 5

DJIA was chosen as a subject of this thesis for several important reasons. First of all, it "is the most-quoted market indicator in newspapers, on television and on the Internet" (CME Group Index Services, LLC, 2011) in the world (other widely quoted market indicators are S&P 500, FTSE100, NASDAQ, NIKKEI). Consequently, it is always very well traded – over the sample period of 2010 average daily trading volume on the composite stocks was approximately 4.6 billion shares (Figure 6 shows how DJIA trading volume fluctuates over the sample period), maximum – 10.6 billion shares traded on 2010.05.06, minimum – 1.3 billion on 2010.12.22 (just before a major holiday, when trading usually declines (Chordia, Roll, & Subrahmanyam, 2001)). A further very important factor in this choice is that DJIA is available on both options market and predictions market.

According to official DJIA website (CME Group Index Services, LLC, 2011) DJIA was the first indicator of overall stock market – Charles Dow picked 12 stocks, added their prices and divided by 12. That was in 1886, a hundred years later the calculation is different. In order to "preserve historical continuity" (CME Group Index Services, LLC, 2011) the sum of composite stocks prices is divided by a divisor, which is adjusted when a stock split or a merger occurs for the constituent companies. This divisor is currently at 0.13212949 (one dollar change in composite stock price moves the whole index by 7.56833316 points (1/0.13212949)) – last (and the only change in 2010) change was on 2010.07.02, when "Verizon Communications Incorporated spun off New Communications Holdings Inc. (Spinco)" (CME Group Index Services, LLC, 2010). Because the purpose of this divisor is to preserve historical continuity of DJIA series, this does not have any impact on the analysis of the data.

Another point of interest of the DJIA is its rebalancing strategy – new stocks are from time to time substituted for old stocks in the index. Consequently, out of 12 original constituents at the start of the index, only one company remains with the same name – General Electric (for full list of current constituent stocks see *Appendix 1 – DJIA component stocks*). Fortunately, for the sample in question – 1st of January, 2010, to 31st of December, 2010 – the DJIA did not change its composition.

On the other hand, DJIA is not without its flaws. Most of the criticisms arise from the fact that it is an "average" and not an "index". This means that the resulting figure is constructed by adding up all the constituent prices and dividing by the divisor (or number of constituents, as it was done in the beginning). While this calculation method is very simple and straightforward, its drawback is that if a stock A, valued at $10, increases to $11 (a 10% increase), while stock B, valued at $100, decreases to $99 (1% decline), the average, ceteris paribus, will remain at 55 points. This is undesirable, because, obviously, shares that have high nominal prices have a much greater influence on DJIA. For example, AIG adjusted closing stock price dropped 94% from 8th of September, 2008, to 27th of October, 2008, which led to an almost 3000 point drop in DJIA.

Other strong criticism, presented by John Mauldin (2009), is concerned with the arbitrary choice of the composite stocks for the average. Because the stocks are chosen by a committee to represent "the industrial prowess of the US economy," they may be influenced by public opinion to replace "declining" companies with rapidly growing companies. Mauldin (2009) gives an example of Goodyear and Chevron being replaced by Microsoft and Intel in 1999 (almost at the peak of dot-com bubble), but then market situation changed and Chevron outperformed both these stocks. Analysis by Mauldin (2009) shows that if DJIA was left "as is" since 1928 (when DJIA grew to 30 stocks) it would have performed much better than the actual version.

A study of the DJIA by Shoven and Sialm (2000) enumerates the three biggest flaws as: price weighting instead of market capitalisation weighting, arbitrary composition of the average, and exclusion of dividends. Two of these three flaws were discussed already; therefore, I will shortly explain the authors' position on dividends in the DJIA. The authors explain that DJIA is not a total return index which would incorporate investor's return from dividend payments which underestimates the total return from holding constituent assets of the average. Most of the 30 constituents of the DJIA pay quarterly dividends, giving the average annual dividend yield for the whole life of DJIA of 4.87% (they examine period 1928 through 1999) - this is an enormous difference given that without accounting for dividends the average grew at around 7% annually. The authors find that if the dividends were added to the average to find the total return, the resulting DJIA value on 31st of December, 1998 would be 233,060 (also assuming value weighting instead of price weighting – price weighted estimate is not available in the paper) instead of 9,181 (as it actually was). This is a non-trivial difference – 25 times, however, in this research, this flaw of the average is actually an advantage, because option and prediction contract prices do not have to be adjusted for dividends. This is because neither option nor prediction contracts give the holder the right to receive dividends, so dividends are not included in the actual valuation of either contract. What is even more important is that because DJIA is constructed excluding dividend payments, its value fluctuates depending on dividend payments (DJIA price increases prior to dividend payment and then sharply falls by the amount of the dividend after it is paid). The calculation of the index without adjusting for dividends (except large dividend payments, which are adjusted for through the divisor – not done in 2010) increases the overall volatility of the average, but this should be adequately reflected in the market implied volatilities.

In summary, DJIA is a very liquid market indicator; however, it has flaws concerning its calculation, composition, and dividends. Nevertheless, the criticism that it would have been "outperformed" by not adjusting its composition is not important in this discussion, because a good market indicator is not necessarily a good portfolio. The argument about calculation is more worrying, however, in this case, DJIA and predictions of it are analysed, therefore, it is assumed that the market predictions take into account the information on the calculation problems of the DJIA. Moreover, the dividend calculation flaw is actually an advantage in this case.

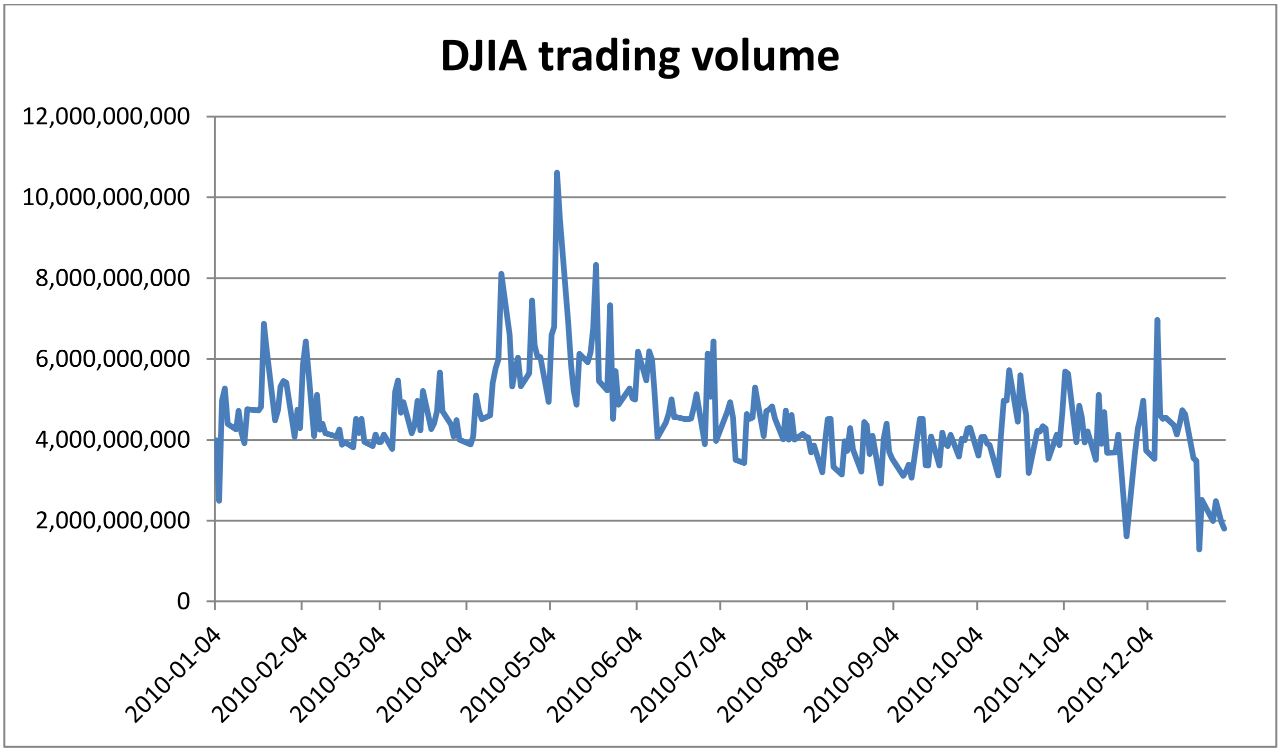

Figure 6

Prediction market – DJIA prediction contracts

Prediction market used in this research is operated by Intrade Inc. It is an exchange where many different types of prediction contracts are traded – from American Idol winners to U.S. presidential election nominees.

Prediction market contracts used in this thesis are based on the realisation of Dow Jones Industrial Average (DJIA) value on 31st December, 2010, which were traded on Intrade prediction exchange (Intrade, 2011). There was a panel of 17 contracts each concerning DJIA value of the form “DJIA to be on or above value X on 31st December, 2010” where X ranges from 7,000 to 15,000 in 500 points increments. Volume traded on these contracts over 2010 – 19,473 contracts.

Contracts were traded for 266 days over 2010. There is no information on trading days on the official Intrade website, however, from the data it appears that all days with trading activity are included in addition to non-holidays (according to Republic of Ireland law). Paragraphs *Differences in data series* contain more information on how differences in data series length and availability were resolved.

Dataset procured from Intrade includes this data: date, contract symbol, contract id, session closing price, session trading volume, open interest (volume of long (short) positions – see Figure 7).

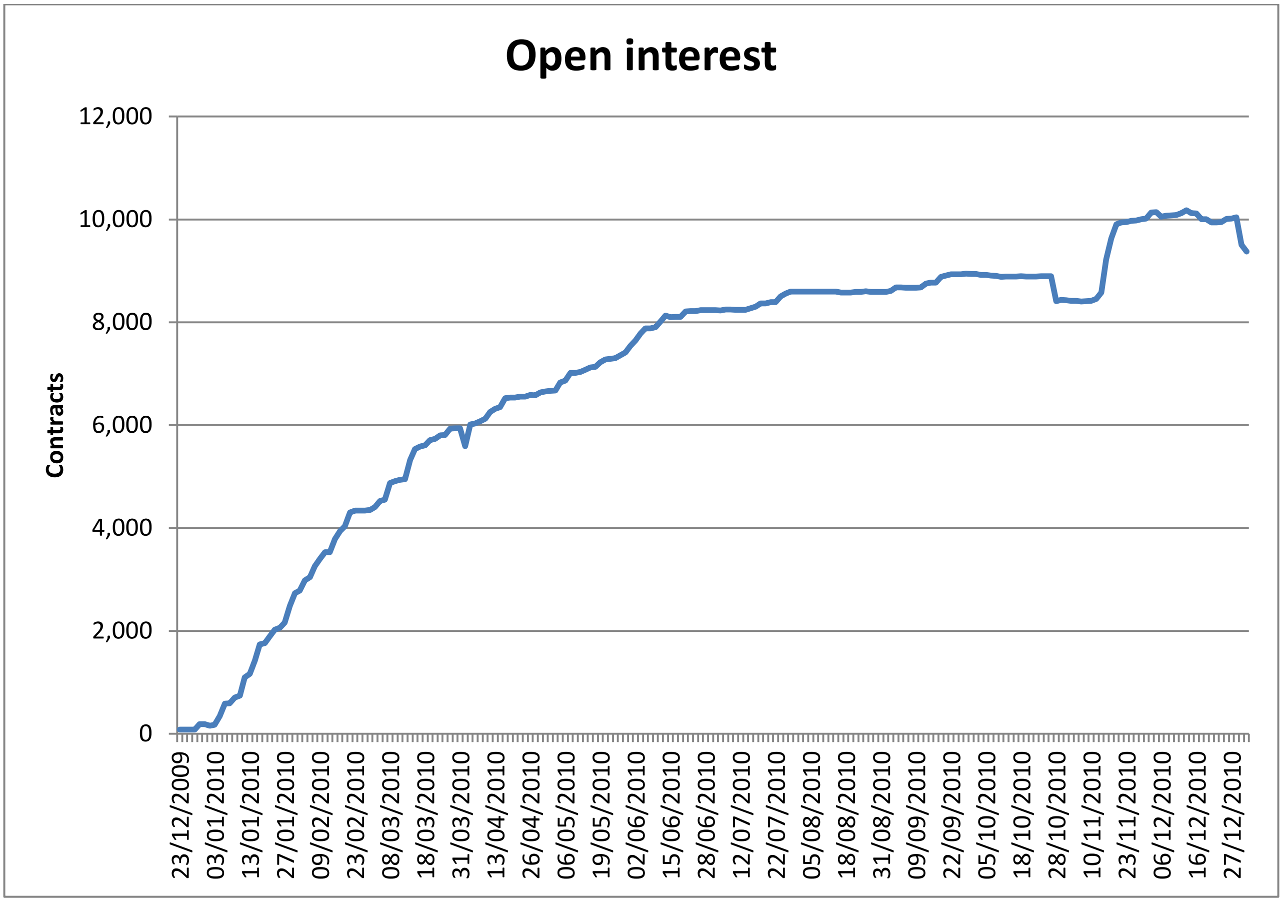

Figure 7

Out of the full list of 17 contracts, 12 contracts were selected (7,000 to 12,500) because the upper tail has such an undesirable property as low trading volume and, more importantly, prices indicating arbitrage opportunities – contract "DJIA to be on or above 15,000 on 31st December, 2010" was sometimes more expensive than "DJIA to be on or above 14,000 on 31st December, 2010". This market inconsistency would make it very difficult to make inferences about market expectations. Furthermore, due to cost considerations, data on upper tail contracts (DJIA at 13,000 and higher) was not purchased.

Taking into account exclusions discussed in the previous paragraph, there are 3,183 price points in the data set. Further taking into account that there was very little trading happening over the period 2009.12.23 to 2010.01.03 (1.2% of total traded volume over 2010) and this period is less interesting because neither of the other datasets contains any data on this period, the data was also excluded. This leaves the prediction market data set at 3,155 price points. Trading volume of these contracts is presented in Figure 8.

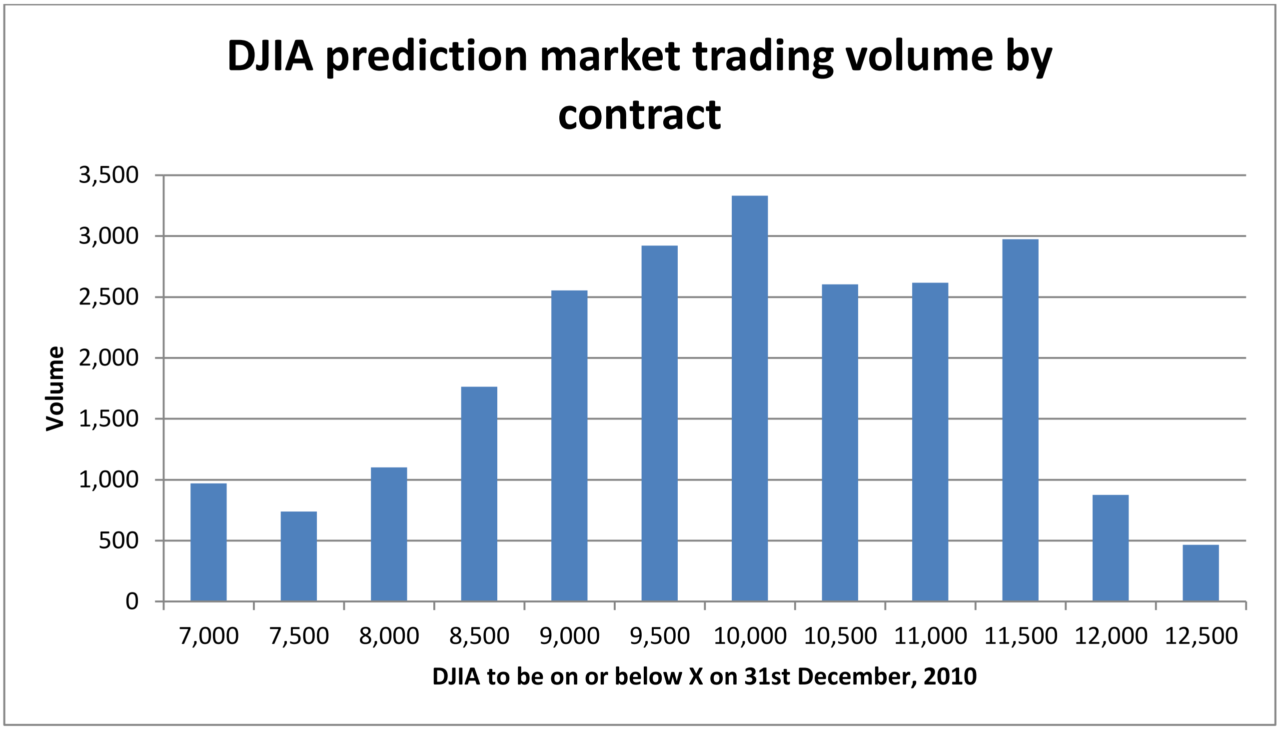

Figure 8

Finally, to show the general idea of the available data I present some market price sets transformed to represent cumulative distribution functions. Transformation is required because the contracts are formulated as "DJIA to be on or above X on the 31st of December, 2010". The price for such a contract gives the probability that DJIA will be greater or equal to X ($P(DJIA \geq X)$). In order to convert this measure to a more conventional way of expressing probability as DJIA to be smaller or equal than X ($P(DJIA \leq X)$), I subtract the actual price (or probability) from 1 to arrive at the data points presented in Figure 9.

Figure 9 shows the evolution of market probabilities over the year. Middle of the year – June to September – is omitted because this period exhibits significantly lower trading volumes (see Figure 10). This effect is probably due to holiday season (similar effect is unobservable to the naked eye in the DJIA data).

Presented figures show that (at least for these dates) prediction market prices do not violate boundary conditions – probabilities are strictly increasing in DJIA values. It should be noted that this CDF boundary condition is not always satisfied because of low trading volumes and high transaction costs (monetary and time-effort costs) – one of the reasons why contracts above 12,500 were excluded was that they often did not satisfy boundary conditions.

Figure 9 clearly shows how market becomes more and more certain of the DJIA outcome as time passes by – each successive CDF is steeper and steeper. Admittedly, these properties are not exhibited over the whole time series but, as expected, the general trend remains.

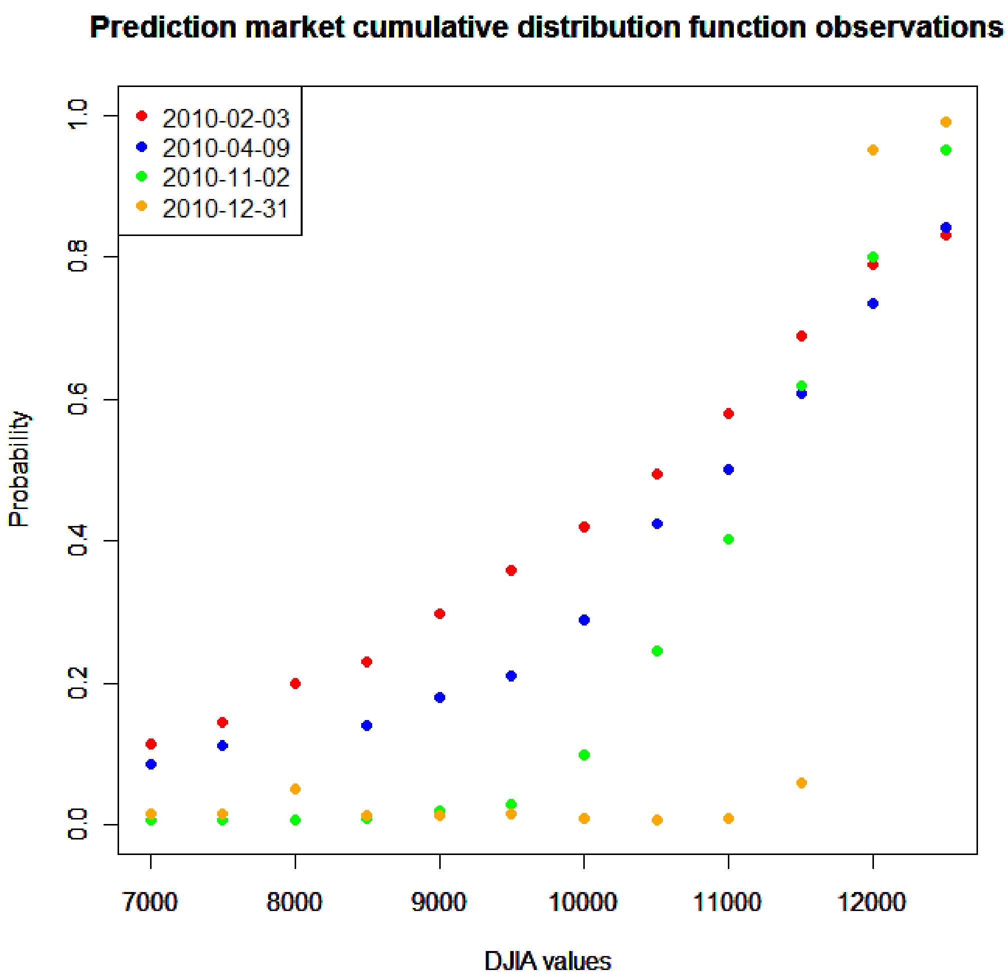

Figure 9

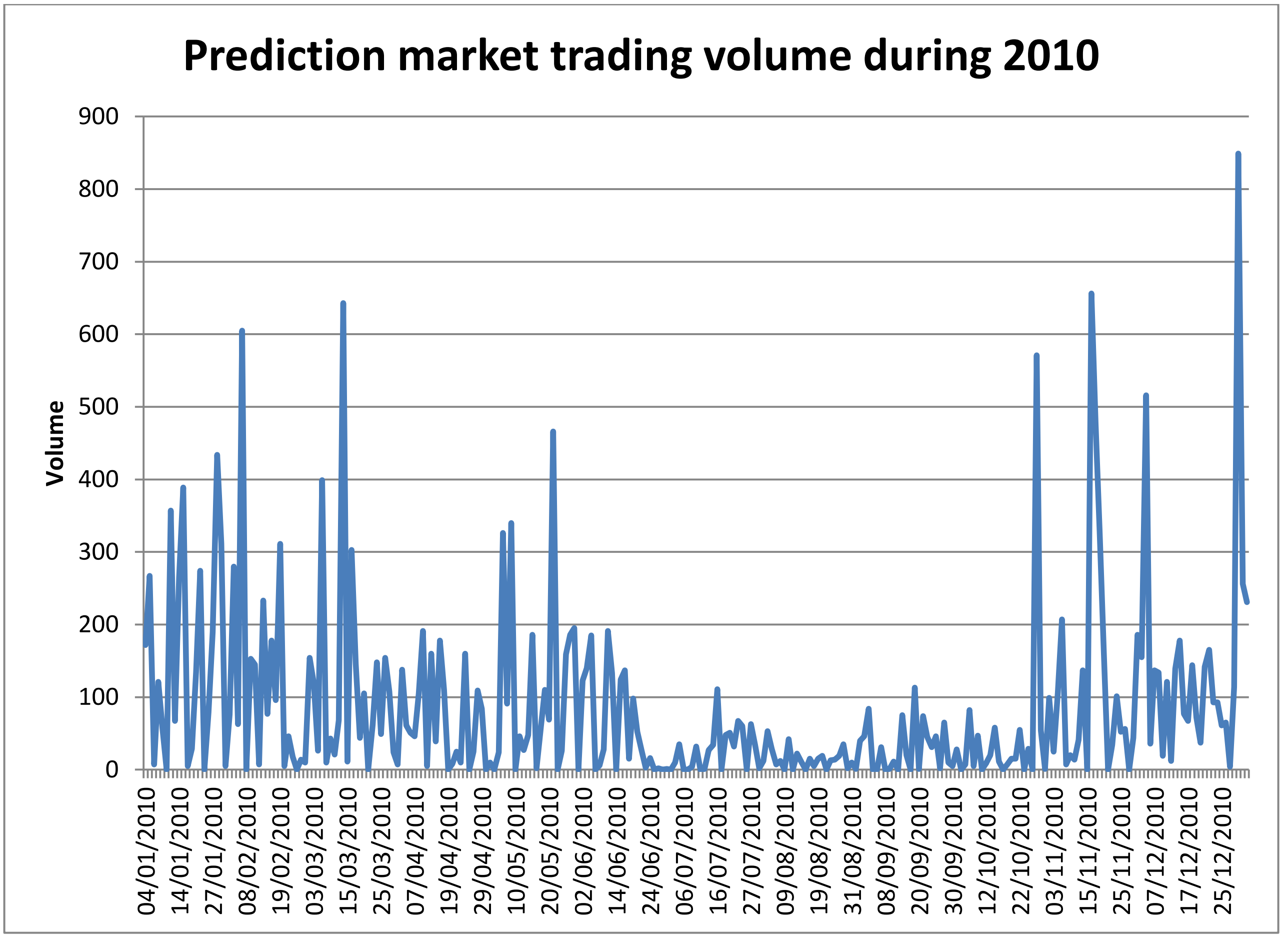

Figure 10

Options market – DJX options

Options used in this research are based on DJIA (the underlying) – for practical purposes "the DJX index option contract is based on 1/100th (one one-hundredth) of the current value of the Dow Jones Industrial Average." (Chicago Board Options Exchange, Inc., 2011) These options are traded on the Chicago Board Options Exchange (CBOE) since 1997 (this and following information about the options were found on official CBOE website, unless otherwise indicated).

One of the most important features of these options is that their exercise style is European. This means that they can only be exercised on the expiration date as opposed to American style options which can be exercised at any point up to and including expiration date. This is very important for two reasons – first of all, pricing of these options is much simpler because the model does not need to take into account "early exercise" value. Second, prediction market contracts are, in essence, European because they are settled based on the DJIA value at the end of the term, meaning that prediction and options market results can be compared.

Data on all DJX options for the year 2010 was kindly provided free of charge by Livevol Inc. Out of this dataset, options expiring on the 18th of December, 2010 are used because this is the closest traditional option expiry date to the expiry date of prediction market contracts. It is expected that by using these options it will be possible to compare and draw conclusions regarding these two markets (10 trading days separate 2010.12.18 and 2010.12.31). DJX options expiring 18th December, 2010 (hereafter: DJX options or DJX) are traded at 111 strike prices – from 40 to 150 in 1 increments (taking into account that DJX is based on 1/100th of DJIA this corresponds to 4,000-15,000 DJIA interval in 100 point increments).

Data available in the dataset: quote date, option type (call or put), strike, open, high, low, closing prices, trading volume, open interest, volume-weighted average price; and the following data at 15:45 and end-of-the-day – bid size and price, ask size and price, underlying asset bid and ask prices.

Trading volume for DJX options over 2010 is 232,527[3] contracts (see Figure 11), open interest at the end of the option life – 133,494 (see Figure 12). Note on the trading days: options are traded every day DJIA is traded; however, in this particular data series there are 242 trading days because 10 days are "lost" due to option expiring on the 18th of December.

[3] Options are traded in bundles of 100, so the actual number of options on the underlying is 100 times higher.

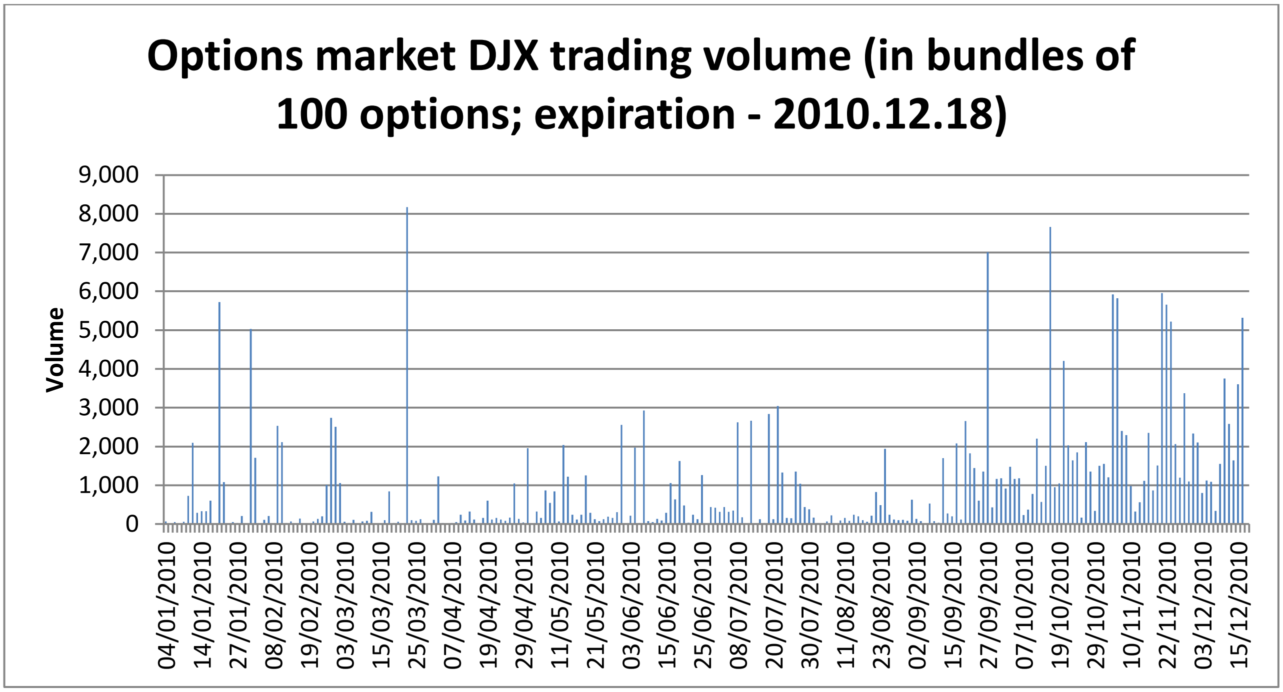

Figure 11

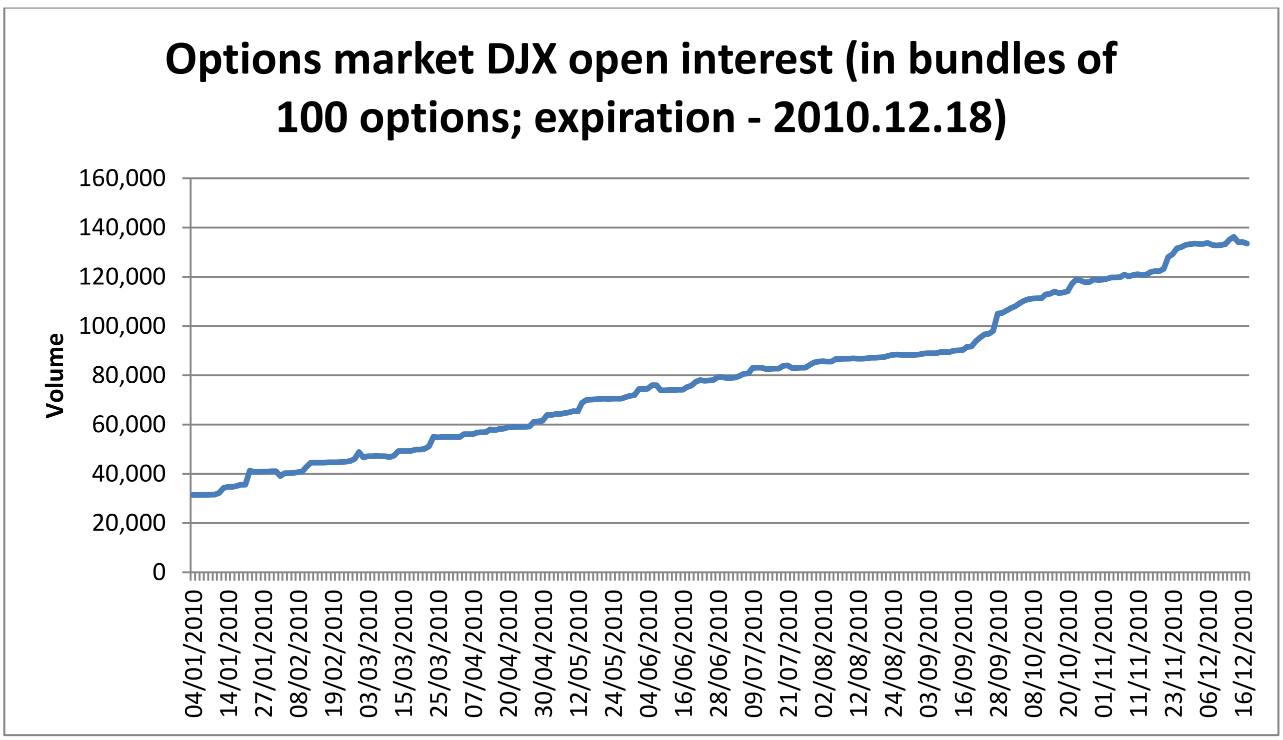

Figure 12

Differences in data series

As already mentioned, using three different data series for comparison creates a problem that not all markets are open on all dates. In this case prediction market data contains 266 trading days, options market – 242, stock – 252. These differences arise for two reasons: first, options and stock market have a 10 day difference because the relevant options are not traded over the last 10 trading days of 2010; except this there are no other differences in trading days between these two series. Second, prediction market and stock market have differences because trading on the stock market occurs on official market trading days (NYSE Euronext, 2011), whereas prediction market may be traded every day but it appears that all days with trading activity are included in the data series in addition to non-holidays (according to Republic of Ireland law).

I resolve this complication by using same day values of the different data series. This method allows comparison of values without loss of economic meaning and avoids unnecessary complications of interpolation.

A further small inconvenience is caused by trading times, because options are traded in the United States during official trading hours (13:30 to 20:15 GMT (Chicago Board Options Exchange, Inc., 2011)), whereas prediction market contracts are traded around the clock. Data, however, shows that 79% (see Figure 13) of trading on Intrade on the contracts in question occurs between 12:00 and 22:00 server time (Intrade is an Irish company, which would lead to conclusion that its server time is based on the time zone of Republic of Ireland – GMT in winter and GMT+1 in summer due to daylight savings) – this roughly corresponds to CBOE trading hours. In short, this inconvenience is negligible, because most trading occurs simultaneously with options market and any news should be aggregated on a similar time scale in both markets.

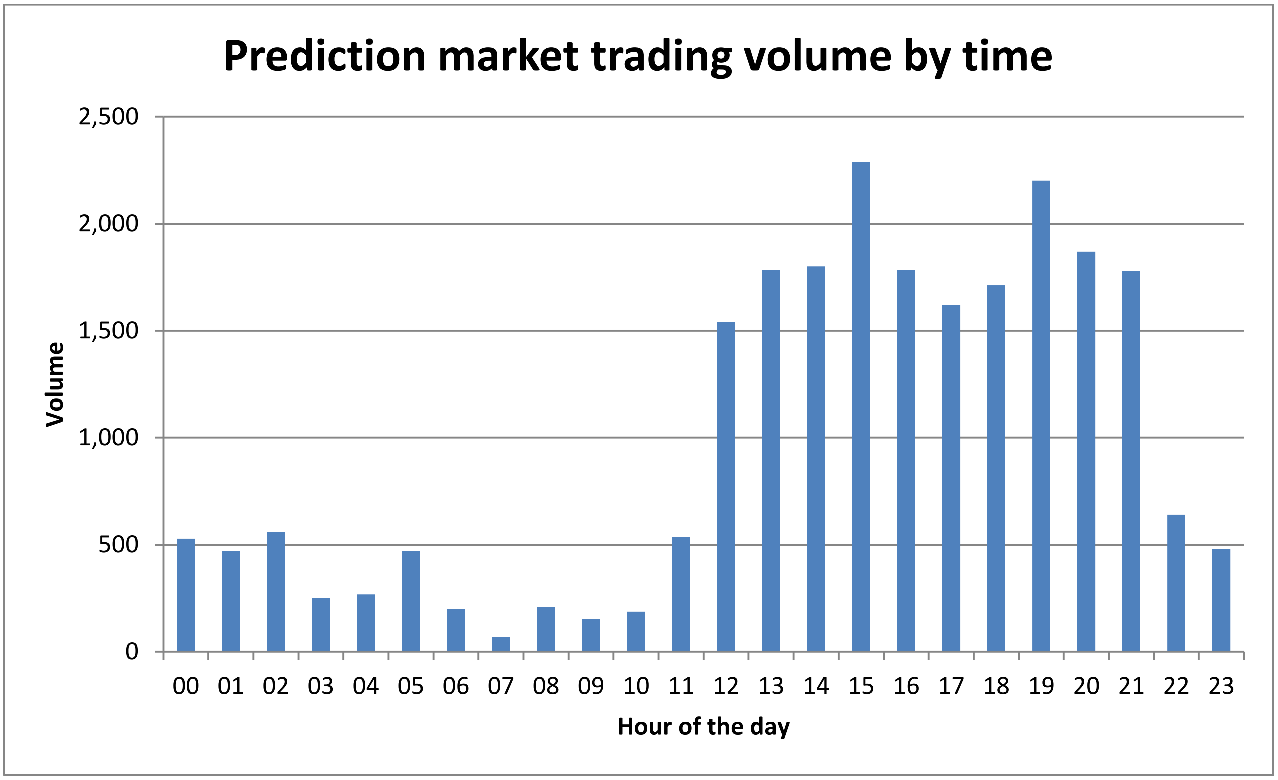

Figure 13

Another problem with comparing prediction and option markets arises because of differences of expiration date – prediction market contracts are conditional on DJIA value on the 31st of December, 2010, whereas options are conditional on DJIA value on the last quadruple witching of the year (3rd Friday of the last month of the quarter is the day when four types of options and index futures, namely "stock-index futures, stock-index options, stock options and single stock futures" (Phillips, 2010), expire; in this case it is the 18th of December, 2010). In short – prediction market and options market predict two slightly different values. The difference is small – 10 trading days – but significant close to the expiry, because, in the extreme, on the 18th options are settled while prediction contract settlement is still 10 days in the future.

Prediction market

In order to ensure that the results of this study are reliable, the prediction market contract prices have to follow random walk. This is first examined graphically, then appropriate stationarity tests are performed: Augmented Dickey-Fuller and KPSS tests.

Probability density function of DJIA expected values

Given the panel of DJIA contracts it is possible to construct a frame of the cumulative distribution function of expectations of DJIA values. Interpolation using splines for values in between available market probabilities will yield data for constructing a smooth probability density function. In order to avoid negative probability density functions (which are, by definition, not possible) I use smoothed monotonic quadratic splines at this step to ensure that the PDFs are always non-negative.

Smoothed cumulative distribution function can be used to numerically derive a probability density function in the following way. First, endpoints are defined – these will be DJIA values of 7,000 and 12,500, because data is available for this interval. Next step is to use the fitted values of the quadratic spline to find values of the probability density function. I propose to do this by iterating the algorithm CDF *(i + 1) - CDF (i) = PDF (i + 1)* for *i = 7001* to *12500*, where *CDF* is cumulative distribution function, and *PDF* is probability density function. By iterating the aforementioned algorithm for every date in the dataset, probability density functions for every day are generated.

A caveat of this method is that it is expected that due to the structure of the prediction market – one fixed panel of contracts for the whole year – as the „expiration“ date approaches and the value of DJIA approaches one particular side of the distribution (in this particular case DJIA ended at 11,578 points (adjusted closing price) on the 31st of December, 2010), one side of the distribution is „lost“, because there are only 2 contracts for values higher than 11,578 points. Furthermore, as the expiration date approaches, the predicted standard deviation becomes much more difficult to determine, because the distribution of expected values is sampled at regular intervals (500 points) and as it contracts (due to shorter horizon) the samples become more unreliable because of bid-ask spreads and other transaction costs.

One method to determine the distributions is visual comparison to computer generated distributions; another option is statistical analysis by using Kolmogorov-Smirnov test. Kolmogorov-Smirnov test checks whether two data series were drawn from the same distribution. It is expected that the logarithm of DJIA expected values should either be close to Gaussian distribution or, possibly, a Levy distribution (Kou, 2008). Kou (2008) suggests Levy distributions as an alternative to Gaussian distributions, because they allow for extreme changes in asset prices that are observed in the markets but are extremely unlikely if Gaussian distribution is assumed. However, 2010 was a relatively good year for DJIA and these extreme events were not observed, moreover, a short time frame makes extreme values unlikely, therefore, testing whether returns follow a Levy distribution would be difficult.

Expected values and volatility

Most interesting information that could be gleaned from prediction market data is the mean (or expected value) of the distribution and the standard deviation of DJIA expected value.

These statistics will be discovered by using the previously discussed probability density functions. Because the probability density functions will be created from the observed data points using interpolation, there will be 5,500 discrete pairs of DJIA values and probabilities for every day of the sample. Interpolation generates these 5,500 pairs because a corresponding probability is calculated for every integer in the range 7,000 to 12,500.

In order to extract the mean and standard deviation without assuming anything about the distribution of the expected values (nonparametrically), random sampling of the interpolated probability density functions will be used. To ensure that results are reliable, a probability weighted sample of 1,000 data points is taken without replacement to construct the standard deviation and mean of the distribution.

It is expected that this analysis would produce a measure of standard deviation of the expected DJIA values at the maturity of the contract and mean expected value of DJIA at the end of 2010. However, this expected value should be equal to current DJIA value discounting influence of transaction costs, bid-ask spreads and time value of money. The rationale is this: if market participants expect DJIA to be above the current price at some certain point in the future, then trading would occur at the current price in order to profit from the expected price increase. The observation that DJIA price and DJIA prediction market expected value are highly correlated and similar in magnitude would indicate that the prediction market is efficient in this sense.

Options market

Data selection

Following Jackwerth (2000), out of the full dataset of option prices some of the options have to be omitted because reliability issues due to low volume and low sensitivity to volatility. Jackwerth (2000) discards options that have bids less than 1/16th and moneyness ratios ($\frac{\text{strike price}}{\text{DJIA level} \times 100}$) is outside the interval 0.84 to 1.12. He also excludes options that have not been traded during the day.

Jiang *et al.* (2005) (they follow Bakshi *et al.* (1997) in their selections, who provide more clarity as to why these particular filters are required) use somewhat different filters on the option data. They remove options with quotes of less than 3/8^{ths}, and options with less than a week to expiration. They further cull option quotes that were time-stamped later than 15:00 Central Standard Time (authors use time and sales data instead of closing prices), because they cannot match these prices to index values. Another filter the authors use is to discard in-the-money options – they discard call options with moneyness of less than 0.97, and put options with moneyness above 1.03, because these options are more expensive and less liquid. Finally, they remove options that violate the boundary conditions (more on these can be found in a paper by Galai (1978)). Out of this reduced set they pick option quotes between 14:00 and 15:00 because trading is usually higher during this hour than for other times of the day.

Given these two approaches and the practical problem of small data set (I use only one option maturity instead of a range), I discard options based on trading volume, boundary conditions, and availability of risk-free interest rate (if risk-free interest rate is unavailable, then implied volatility calculation cannot be performed). Option has to be traded during the day, in order to use the mid-point of the bid-ask spread in the estimation (bid-ask spread is used following Bakshi *et al.* (1997)). Before the aforementioned restrictions there were 26,973 put and call option entries each; after – 1,544 put option entries and 785 call (trade volume restriction reduced the numbers to 1,569 and 842, respectively).

A note on different information

It is important to note that prediction markets and option markets reveal fundamentally different information about the DJIA. As discussed before, prediction markets reveal market belief about the distribution of DJIA values at some particular date (in this case – 2010.12.31); whereas options market reveals information on the asset (DJIA) price generating process. When B-S model is used, logarithmic returns on the DJIA are assumed to follow a normal distribution with some mean and standard deviation. Implied volatility is actually a measure of this deviation and thus is a characteristic of the asset price generating process. As previously discussed, implied volatility can also be calculated without assumptions about the asset price generating process. Furthermore, using a similar technique it should be possible to construct returns distribution but this is beyond the scope of this work. Using the returns distribution it is then possible to construct a distribution of DJIA values at some future date. Nevertheless, this thesis will not attempt this, but will present information that requires the least assumptions – DJIA values distribution from prediction market and DJIA returns standard deviation (volatility) from options market.

Black-Scholes implied volatility

B-S implied volatility was briefly discussed in *Literature review*. Here is presented the step by step method of deriving implied volatility in practice.

Implied volatility can be calculated either from call or put options – the general method is the same, but the B-S formulae employed are slightly different. It will be assumed for ease of discussion that the calculations are performed on a call option.

The first, and probably most important task, is collecting the required inputs: call price, spot price of the asset, strike price, time to expiry, annual risk free rate. In this particular case, all data except risk free rate were included in the dataset provided by Livevol Inc. Risk-free rate data characteristics were presented above.

The final step is to calculate the actual implied volatility. This is done by using numerical optimisation methods to find the volatility that gives the lowest absolute error between actual option price and B-S option price. This calculation is performed in R statistical software package

(R Development Core Team, 2011), which uses a numerical optimisation method to find the implied volatility.

Model-free risk neutral implied volatility calculation method

The following methodology is presented in George J. Jiang and Yisong S. Tian paper *The Model-Free Implied Volatility and Its Information Content* (2005).

The authors aim to implement the Britten-Jones and Neuberger (2000) approach to model-free implied volatility. Their approach is to use all available strike prices and maturities in order to calculate implied volatility of S&P500 index, however, I will use all available (subject to restrictions discussed in *Data selection*) strike prices at only one option maturity (2010.12.18) in my implementation, for two reasons: first, this is the goal of this thesis – to estimate implied volatility for a particular horizon; second, due to lack of data given the budget of this research.

The authors begin by restating integrated variance formula, which is the basis of all further analysis:

$$E_0^F\left[\int_0^T\left(\frac{dF_t}{F_t}\right)^2\right] = 2\int_0^\infty \frac{C^F(T,K) - \max(0, F_0 - K)}{K^2}\,dK,$$

Equation 7

where T is a future date, F - forward asset price (under forward probability measure), $C^F(T, K)$ – forward option price (under forward probability measure), K – strike price. The formula uses call prices, but can be simply restated to use put option prices.

Variance (or volatility, or standard deviation) is an unobservable property of a price process, but it is of great interest to academics and investors alike, therefore, many techniques are available for its estimation and forecasting (see Andersen, Bollerslev, Diebold and Ebens (2001) for an outstanding example of such a study). Realised or historical variance is determined by observing returns at some discrete points in time and then using these observations to estimate the volatility of returns, however, asset prices follow continuous processes, therefore this method is only approximating the actual variance. Variance of the underlying continuous process is known as integrated variance, market forecast of which is estimated, in this case, using market data on option prices. In order to be consistent, I use integrated standard deviation (square root of integrated variance) when estimating implied volatility.

Implementation of Equation 7 in this form requires analytical integration which is practically not possible because required inputs are available only at certain points. However, a result can be arrived at using numerical integration, but certain errors arise. Because strike prices are available only for a certain interval, truncation errors occur (truncation error and other error bounds are derived analytically by Jiang *et al.* (2005)). Truncation error means that Equation 7 requires the integral to be evaluated in the interval of 0 to infinity with respect to strike prices, however, strike prices are available only in a small part of this interval. Consequently, the interval has to be "truncated."

Jiang *et al.* (2005) practically implement the integrated variance formula by using numerical integration using the trapezoidal rule:

$$2\int_{K_{min}}^{K_{max}} \frac{C^F(T,K) - \max(0, F_0 - K)}{K^2} dK \approx \sum_{i=1}^{m} [g(T,K_i) + g(T,K_{i-1})]\Delta K\,,$$

$$\Delta K = \frac{K_{max} - K_{min}}{m}$$

$$K_i = K_{min} + i\Delta K, for\ 0 \leq i \leq m$$

$$g(T,K_i) = \frac{C^F(T,K_i) - \max(0, F_0 - K_i)}{K_i^2}$$

Equation 8

where *m* is the number of intervals to approximate integration on. For a finite *m* (or $\Delta K > 0$), the numerical integration leads to discretization errors. Discretization error, in more detail, arises because Equation 7 requires integration, which, in turn, requires a continuum of strike prices. This continuum of strike prices is not available, because strike prices are usually set in advance at certain convenient prices. Furthermore, as discussed in the part on data selection, many of the data points are excluded on the basis of trading volume, which leads to gaps in the strike price continuum. Consequently, as per Equation 8, numerical integration using trapezoidal rule has to be used, which leads to discretization errors. Fortunately, the size of this error can be mitigated by using a large *m* at the expense of computing time and resources (I propose to use *m = 100*, because it provides reasonable accuracy – negligible differences when compared against *m = 1000* – and computational efficiency).

Another extremely useful addition to M-F implied volatility calculations is the formula for substituting spot prices for forward prices (which are not as readily available as spot prices) in the integrated variance:

$$E_0^F\left[\int_0^T \left(\frac{dS_t}{S_t}\right)^2\right] = 2\int_0^{\infty} \frac{C\left(T, \frac{K}{B(0,T)}\right) - \max(0, S_0 - K)}{K^2} dK,$$

Equation 9

where *S* is spot price, *B(t, T)* – is time *t* price of zero-coupon bond that pays $1 at time *T*. Interest rate used is the corporate interest rate (see *Interest rates*) – this could create an error in the estimate of the "underlying" process because the market could be using a different interest rate.

Equation 6 presents the formula for the price of zero-coupon bond in continuous time – time to option maturity and corporate interest rate are plugged in to give the required value for transforming integrated variance formula from using forward prices to using spot prices.

Equation 10 gives the full formula for integrated variance using underlying asset spot prices and put option prices.

$$E_0^F\left[\int_0^T\left(\frac{dS_t}{S_t}\right)^2\right] \approx \sum_{i=1}^{m}\left[\frac{P\left(T,\frac{K_i}{B(0,T)}\right)+S_0-K_i-\max(0,S_0-K_i)}{K_i^2} + \frac{P\left(T,\frac{K_{i-1}}{B(0,T)}\right)+S_0-K_{i-1}-\max(0,S_0-K_{i-1})}{K_{i-1}^2}\right]\frac{K_{max}-K_{min}}{m}$$

Equation 10

Outline of the process, where previously described formulae are used:

1. Using B-S model calculate implied volatilities for all strike prices.
2. Fit a smoothed cubic spline on these implied volatilities.
3. Use B-S model on the interpolated volatilities to calculate option prices.
4. Using Jiang *et al.* (2005) formula for spot prices and trapezoidal rule of numerical integration find expected return variance.

It should be noted that even though this process uses B-S model in steps 1 and 2, this does not affect the model-free nature of the implied volatility. First the option prices are translated into B-S implied volatilities – introducing the B-S model assumptions into the process. Next, these implied volatilities are interpolated to produce an equidistantly spaced and smooth implied volatility function. Jiang *et al.* (2005) explain that it is better to use cubic spline interpolation on implied volatilities than directly on option prices, because better fit to the actual option prices is achieved. Finally, implied volatilities are translated back into option prices which removes the B-S assumptions – even though B-S model is used, it has no effect on the results. The last step of the process is to apply the Jiang *et al.* (2005) formula on the option prices to construct the model-free implied volatility (Equation 10).

DJIA realised volatility

Practical implementation

The following method of calculating realised volatility is based on Figlewski (1997). The daily standard deviation σ_t of continuously compounded returns (calculated as $r_t = ln\ (S_t/S_{t-1})$) is calculated for the period until respective contract expiration date – 2010.12.18 for options, and 2010.12.31 for prediction market contracts. In order to annualise the σ_t, it is then multiplied by the square root of remaining trading days in the period. This method of calculation results in a monotonically decreasing function of realised volatilities; furthermore, the estimates become less reliable as the expiration dates approach, because less and less returns are available for standard deviation estimation. However, this is the correct approach, because implied volatility measure is intended to estimate the volatility of the asset over the remaining life of the option/contract.

Empirical research results

Testing for random walk

Theoretical foundation

Before presenting the tests for random walk on all three markets, I discuss the theoretical rationale for requiring random walk and the fundamental process generating random walk behaviour.

Schiller and Perron (1985) describe random walk in a financial context as a process, where “price P_t can never be described as 'too high' (i.e., that it can be expected to fall in the future) or 'too low' (i.e., that it can be expected to rise in the future).” Essentially, a random walk is a time series process, which cannot be predicted using historical data. If we find that market prices follow a random walk process that means that the market is at least weak-form efficient (Fama, 1970).

It is useful at this point to shortly describe the concept of stationarity. A very broad overview of this concept is provided by Nau in his lecture notes (2005). Nau (2005) describes a stationary process as such a time series that its properties – mean, variance, autocorrelation, etc. – do not change with time. Unless a process is stationary, it is difficult to construct meaningful sample statistics. Two types of stationarity should also be mentioned – trend-stationary and difference-stationary processes. A process is trend-stationary if, after using the regression to remove a deterministic time trend, the process becomes stationary. Difference-stationary processes lose their non-stationary properties after taking 1^{st} or higher order differences of the series.

We test stationarity using two different tests: Augmented Dickey-Fuller test (henceforth, ADF) and Kwiatkowski-Phillips-Schmidt-Shin test (KPSS). ADF null hypothesis is that the time series is non-stationary, whereas KPSS null hypothesis is the opposite – series is stationary.

There are many different expressions of the random walk process, but in this case we will define a random walk data point as a sum of previous random walk value and an independent identically distributed error. Formally this is stated as $P_{t+1} = P_t + \varepsilon_{t+1},$ where P_t is the price (or any other variable) at time t and ε_t is the independent (does not depend on t or previous values of ε)

identically distributed (meaning that the distribution of errors is constant across time and prices) random error. Consequently, taking first differences of the series will remove the non-stationarity of the series. Differencing the series would reduce the previously described process to $P_{t+1} – P_t = P_t + \varepsilon_{t+1} - P_t = \varepsilon_{t+1}$ – differenced series should exhibit no autocorrelation patterns (required by assumption of error independence) and no deterministic trend (identically distributed errors).

It should be noted that we do not test for random walk with drift (expressed as $P_{t+1} = P_t + \varepsilon_{t+1} + \mu$, where μ is the drift parameter) explicitly, because differencing would also remove the drift parameter. However, based on common knowledge about stock prices, it is reasonable to assume that at least stock market prices follow a random walk with drift process.

Concepts for analysis

Autocorrelation function (ACF) is a function that shows correlation of a variable with its own lags. An independent random variable would not be significantly autocorrelated at any lag – presence of autocorrelation indicates that a relationship can be discovered.

Partial autocorrelation function is a function similar to ACF except for one important difference – linear dependence of lag $k + 1$ on lag k is removed.

Stock market

Figure 5 in *Research methodology* shows the path of DJIA in 2010. Random walk in stock market prices will be analysed in three steps: stock prices, natural logarithms of stock prices and differences of natural logarithms of stock prices will be tested for random walk.

First, Figure 14 panel 2 shows ACF function for the series – autocorrelation of almost one for lag 1 and decreasing in higher lags, panel 3 shows PACF of almost 1 for lag 1; this is evidence that stock price at time t can be explained to a large degree by stock price at time $t – 1$ and PACF shows that lags of a higher degree have almost no explanatory power once P_{t-1} is included (blue dashed lines indicate 95% confidence intervals). ADF and KPSS tests are in agreement with each other – series is not stationary.

Given these ACF and PACF graphs it is not necessary to present the tests (they were performed but did not differ materially from previously discussed results) whether taking logarithms of stock prices will solve the problem of non-stationarity, therefore, we go ahead to testing whether logdiffs (differences of logarithms) will exhibit non-stationary behaviour. It is important to note that logdiffs of stock price time series are continuously compounded returns on the DJIA.

Figure 15 presents the DJIA returns series (logdiffs) – visual observation of panel 1 shows no linear trend in the series (regression line superimposed on the scatter plot in red), however, the histogram (panel 2) suggests that returns did not follow a Gaussian distribution (at least over this sample period) – the peak is too high and the tails of the distribution are too wide (this is discussed more formally in part DJIA returns). (A short side note on kernel density (based on notes by Duong (2001)) – it is a density estimator which is constructed by adding up Gaussian PDFs (with area *1/n*, where *n* is the number of observations) centred at the data points; the standard deviations of said Gaussian PDFs are determined by an optimisation process. The main advantage of kernel density compared to histogram is that it is a smooth function with objectively determined “bins”).

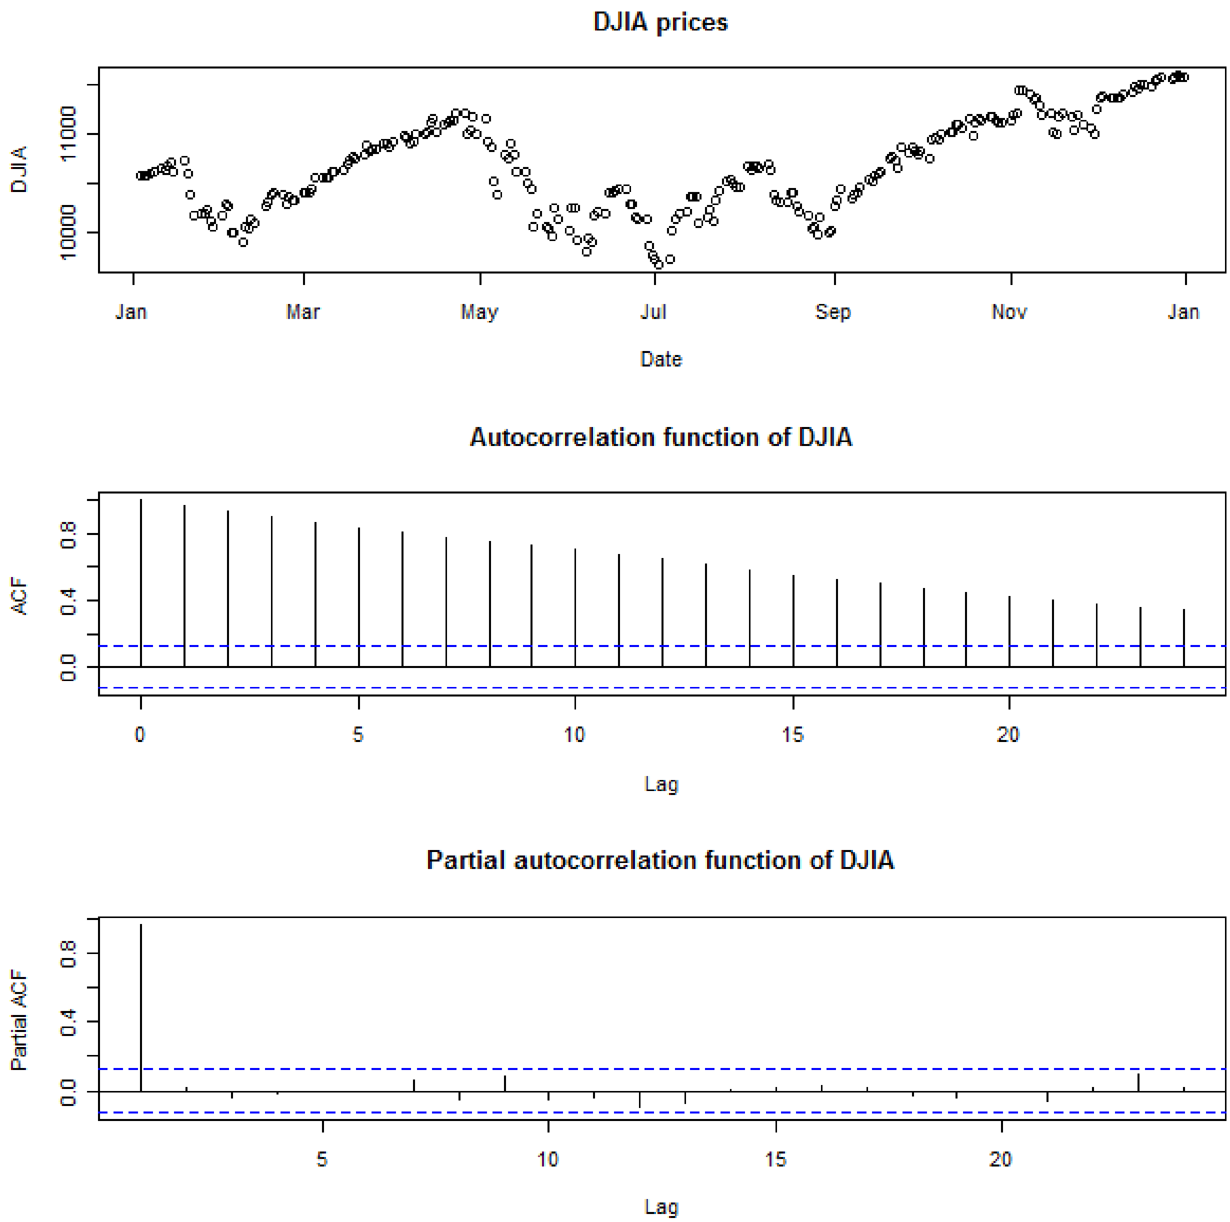
DJIA prices
DJIA
10000
11000
Jan
Mar
May
Jul
Sep
Nov
Jan
Date
Autocorrelation function of DJIA
ACF
0.0
0.4
0.8
0
5
10
15
20
Lag
Partial autocorrelation function of DJIA
Partial ACF
0.0
0.4
0.8
5
10
15
20
Lag

Figure 14

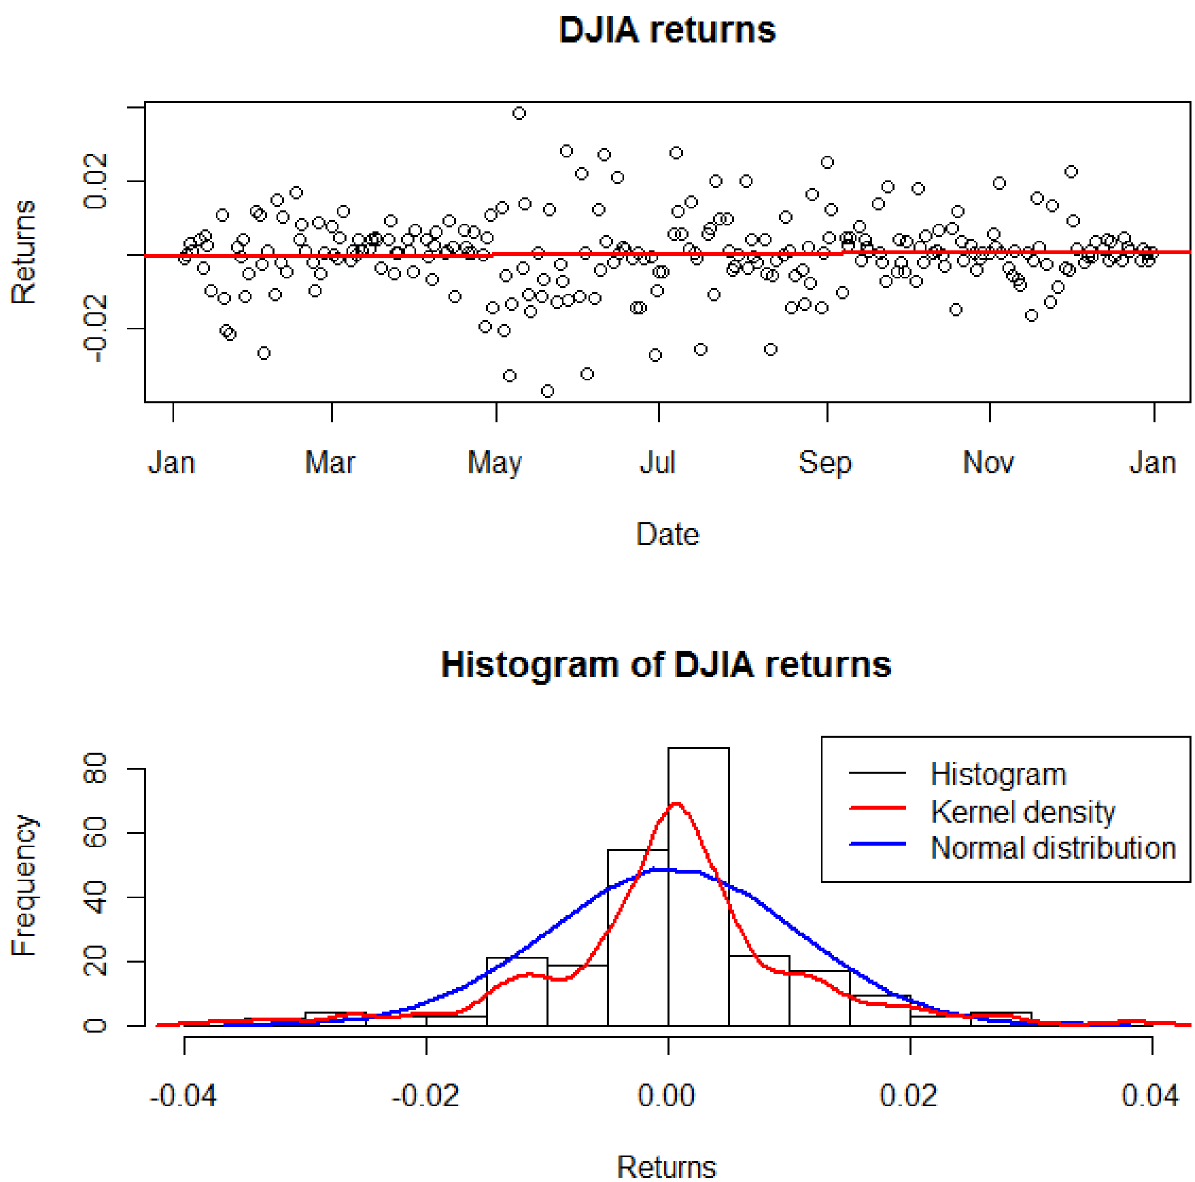

Figure 15

We continue with analysis of random walk of logdiffs of stock prices with Figure 16. No autocorrelations remain in the data for any lag. Furthermore, ADF and KPSS tests concur with the conclusion that logdiffs of DJIA prices do not exhibit any non-stationary behaviour, therefore, we assume that stock prices follow a random walk and are difference-stationary.

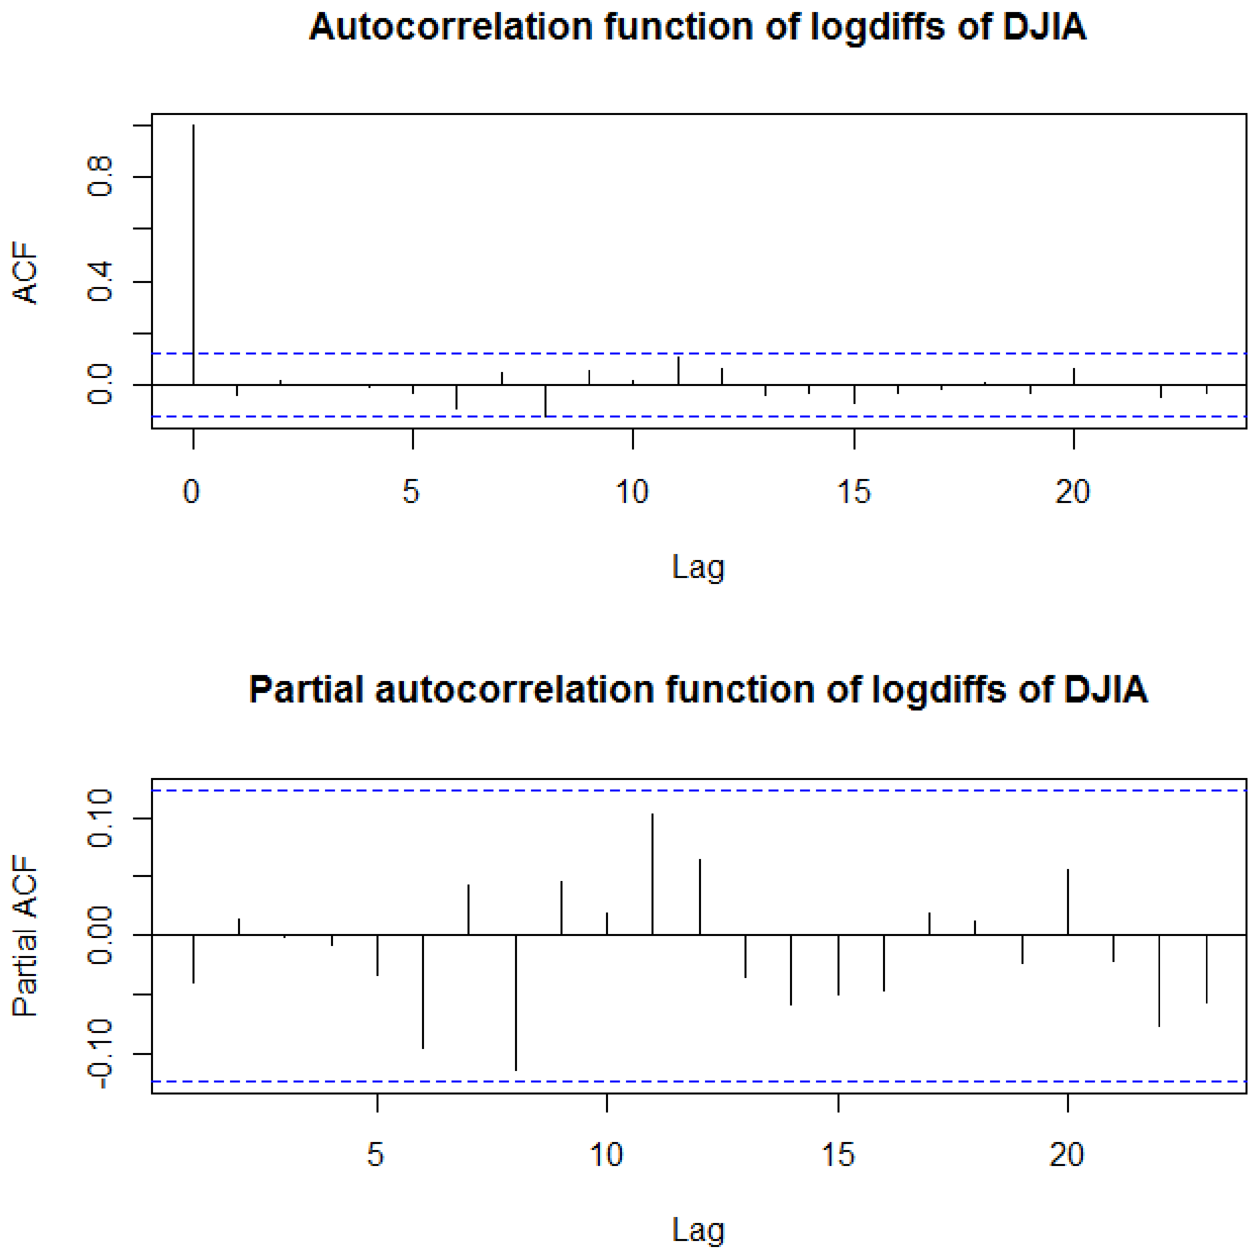

Figure 16

Prediction markets

Similar testing for random walk was also applied to prediction market contracts – the results follow.

All 12 prediction market contracts follow random walk, so, for the purposes of clarity, a representative contract price will be formally tested and presented here. In particular, I have chosen prediction market contract stipulating that DJIA value will be less than 10,000 points at maturity. This contract was chosen because it is one of the more liquid contracts; therefore, it has more variability in prices. However, analysis shows that results are not materially different for the other contracts.

Figure 18 panel 1 shows the path of the contract price over 2010 – the year was started at 50.4% probability that DJIA will end up lower than 10,000 by the end of the year (or price of 50.4 cents for 1$ contract – probability and price are used interchangeably, when talking about winner-take-all prediction market contracts[4]), because at the same time DJIA was at 10,584 points. A first guess would predict that the contract price should have been lower, however, it must be taken into account that prediction market trading at this date was still very low, furthermore, transaction costs increased the "actual" price of the contract. Contract expired at 1%, because DJIA ended the year at 11,578, making this contract worthless. Transaction costs (including time) prevented the contract from closing even lower.

It is also important to note that the contract does not follow DJIA values exactly but very closely, meaning that the prediction market contains information not contained in the DJIA; more information is contained over the full range of prediction contracts. For instance, correlation coefficient between DJIA and prediction market price is -92%, linear model of DJIA with contract price as explanatory variable gives adjusted R^2 value of 85% (see Figure 17 for the scatter-plot and best-fit-line of the linear model).

[4] See Section I for justification.

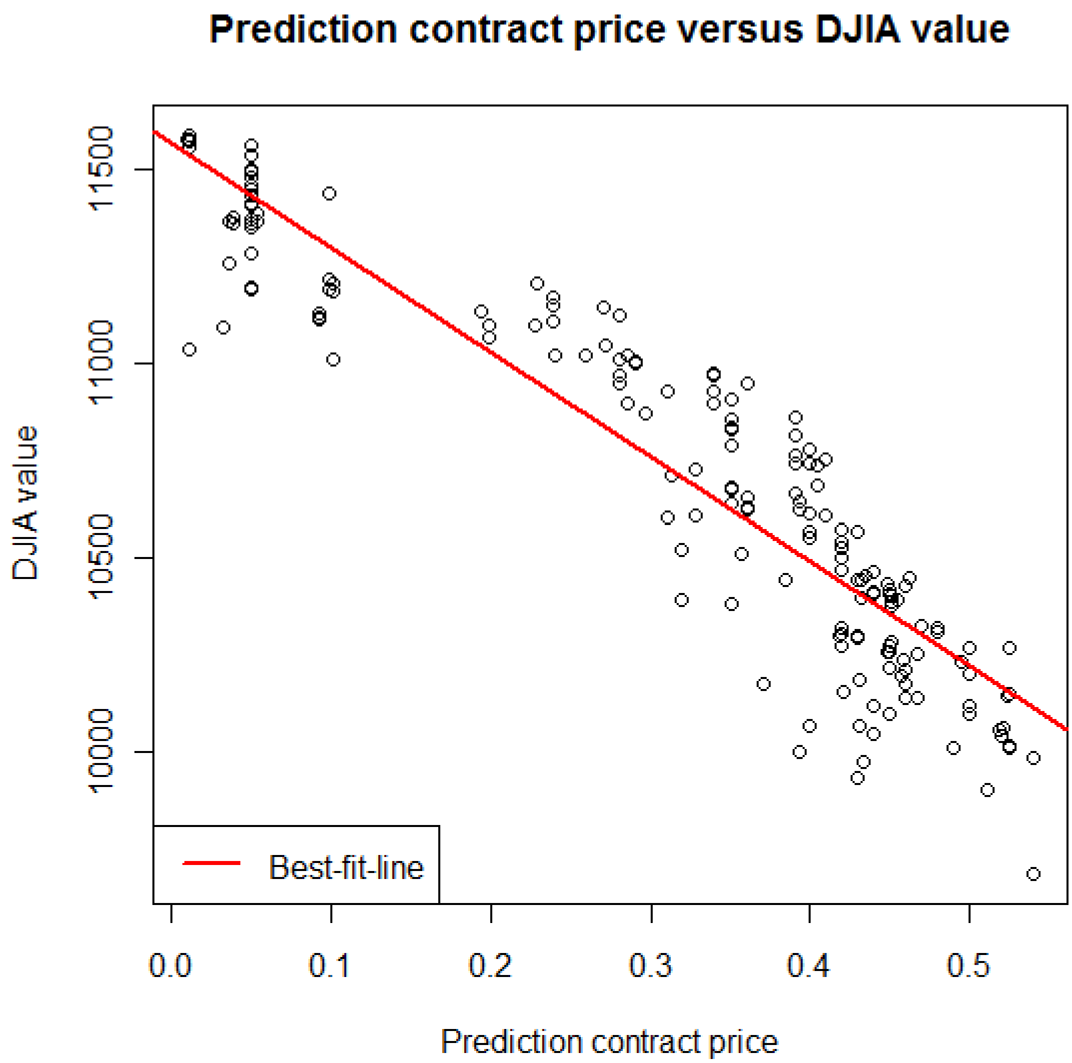

Figure 17

Another important aspect of this time series is that even though prediction market contracts are well traded over the sum of all 12 contracts, this individual contract shows frequent "plateaus," where contract price remains unchanged because of lack of trades for those dates (see Figure 18). This behaviour is not observed in options and stock markets except for the least liquid assets; whereas this contract has the highest trading volume of all 12 contracts (see Figure 8).

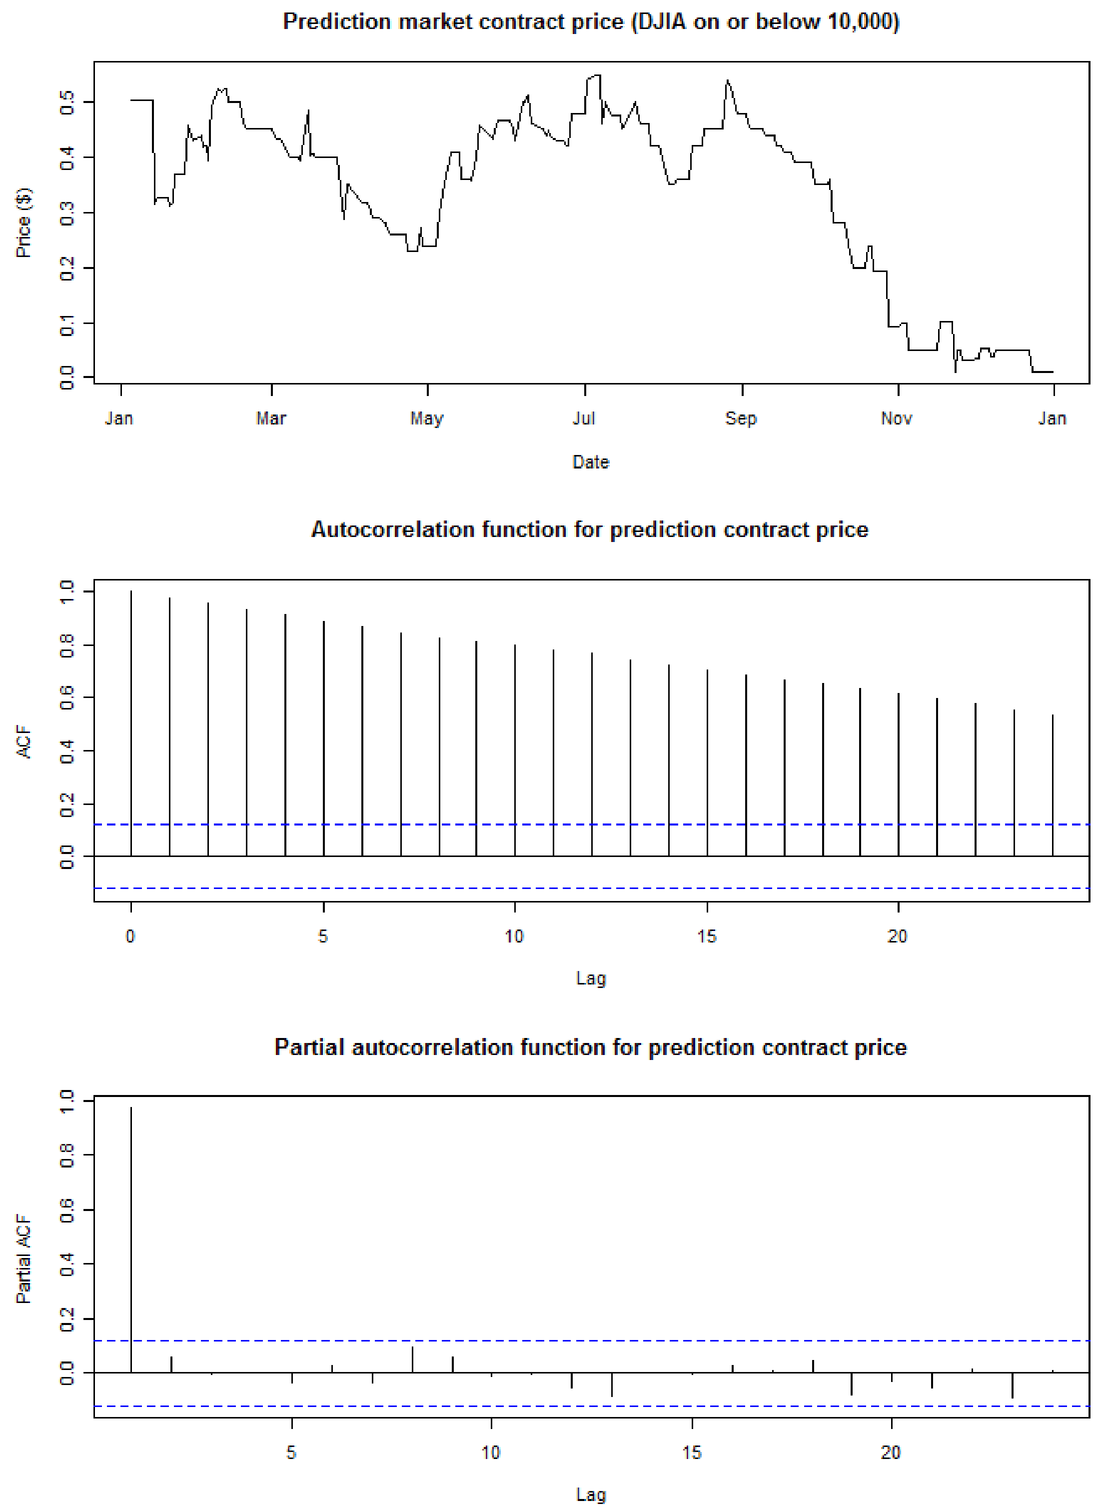
Prediction market contract price (DJIA on or below 10,000)
Price ($)
0.0
0.1
0.2
0.3
0.4
0.5
Jan
Mar
May
Jul
Sep
Nov
Jan
Date
Autocorrelation function for prediction contract price
ACF
0.0
0.2
0.4
0.6
0.8
1.0
0
5
10
15
20
Lag
Partial autocorrelation function for prediction contract price
Partial ACF
0.0
0.2
0.4
0.6
0.8
1.0
5
10
15
20
Lag

Figure 18

Next we examine the second and third panels of Figure 18 – autocorrelation (ACF) and partial autocorrelation (PACF) functions of the contract price. ACF indicates that the prices are highly correlated with each other (autocorrelation at lag 1 is almost 1), however, PACF being significantly non-zero only at lag 1 (and lag 9, but this is most likely an accident – at 95% significance, it is very likely for such a fluke to happen out of 24 lags) confirms the suspicion that autocorrelations at higher lags are only propagations of autocorrelation at lag 1 (Nau, 2005). As discussed above, differencing should remove the non-stationarity of the series – this is corroborated by ADF and KPSS tests (using ordinary least squares regression to detrend the process did not result in a stationary series, therefore, we discard testing for trend-stationarity).

Figure 19 shows the results of differencing: the first panel shows the scatter plot of the differences (excepting cases where differences are zero, to avoid the mentioned problem of price "plateaus" due to lack of trading – this does not influence the results) and the line of best fit (linear least squares regression on index of the difference); the second panel shows the histogram of the differences (excluding differences exactly equal to zero makes this graph more informative without sacrificing much). Observing the scatter plot and the linear regression line it seems that no pattern remains after taking the first difference. Checking for the effect of a consistently downward trend of the contract price from September to year end reveals that even exclusion of all those data points does not change the basic shape of the histogram.

Further testing for random walk requires ACF and PACF of differences to be examined – it is expected that they will not show any statistically significant autocorrelations at any lag. This expectation is fulfilled – evidence presented in Figure 20.

Formal tests of random walk (ADF and KPSS tests for stationarity) support the presented evidence. Augmented Dickey-Fuller test uses the null hypothesis of non-stationarity and alternative hypothesis of stationarity. P-value for this test is less than 1% (R does not provide a more accurate assessment), which means that we reject the null hypothesis. KPSS is formulated the other way round – null hypothesis of stationarity and alternative hypothesis that the process is non-stationary. In this test R provides a p-value as "above 10%," which does not allow us to reject the null hypothesis.

In summary, analysis of the prediction market prices time series shows that the prices follow a random walk, which is expected, considering that these are financial data. Furthermore, all performed tests are in agreement with each other. Also, prediction market prices are difference-stationary.

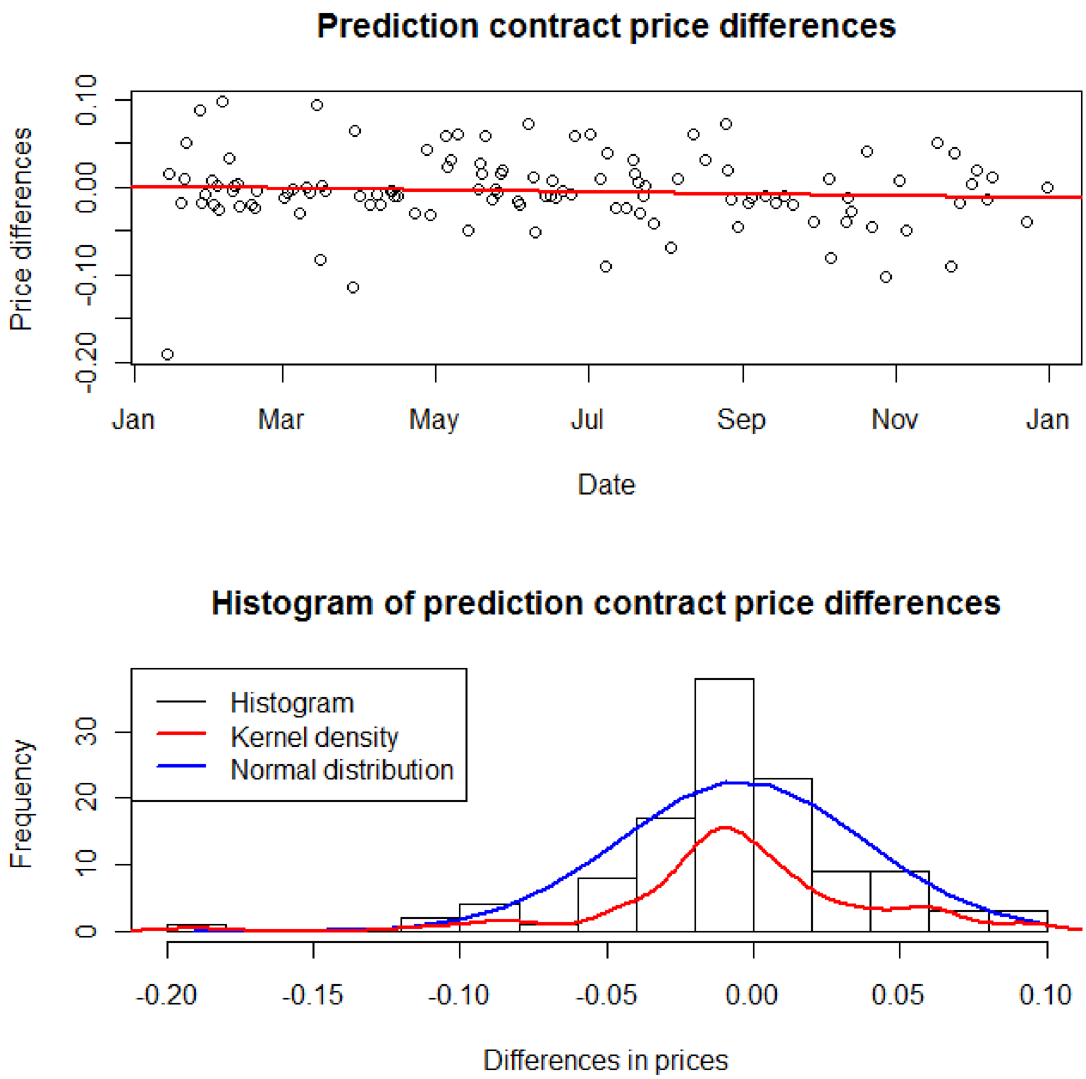

Figure 19

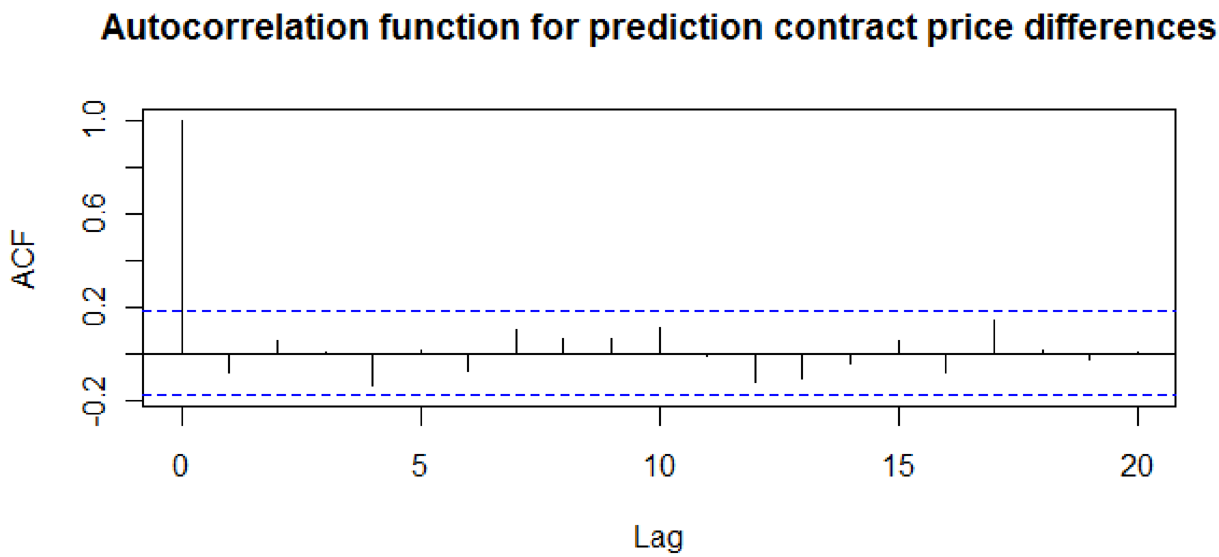

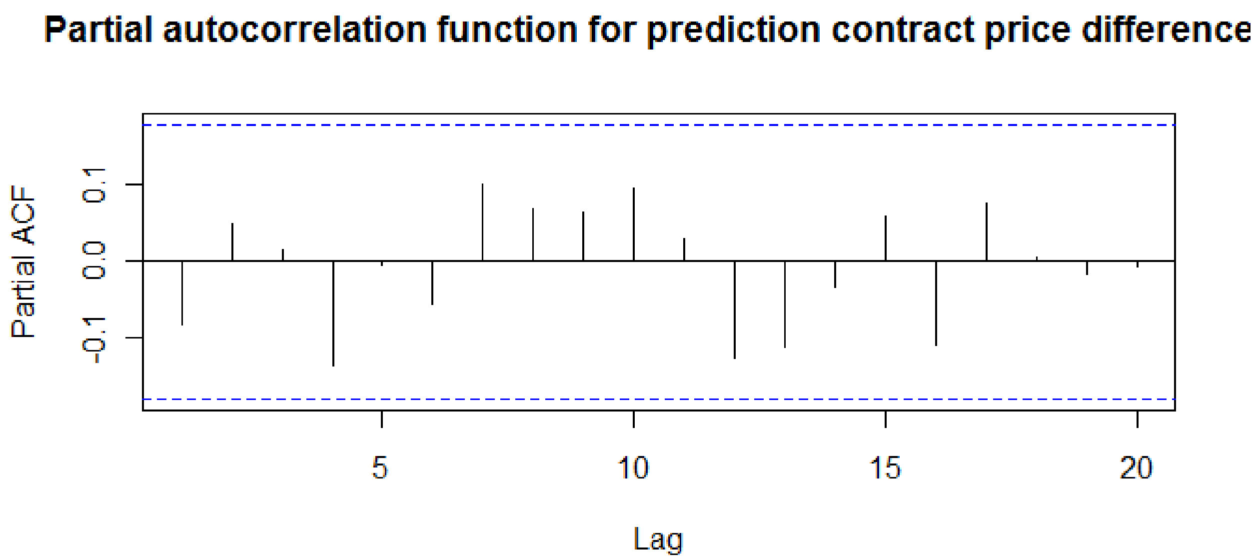

Figure 20

Option prices

In order to test for random walk in option prices I arbitrarily (and with the view to compare this to prediction market contract price) choose a well traded option roughly in the middle of the strike range – at a strike price of 100 (translating to 10,000 DJIA points) – volume on this option over 2010 is 15,069 options. This choice provides robust results because of many observations, while the same random walk properties translate to other strike prices as well.

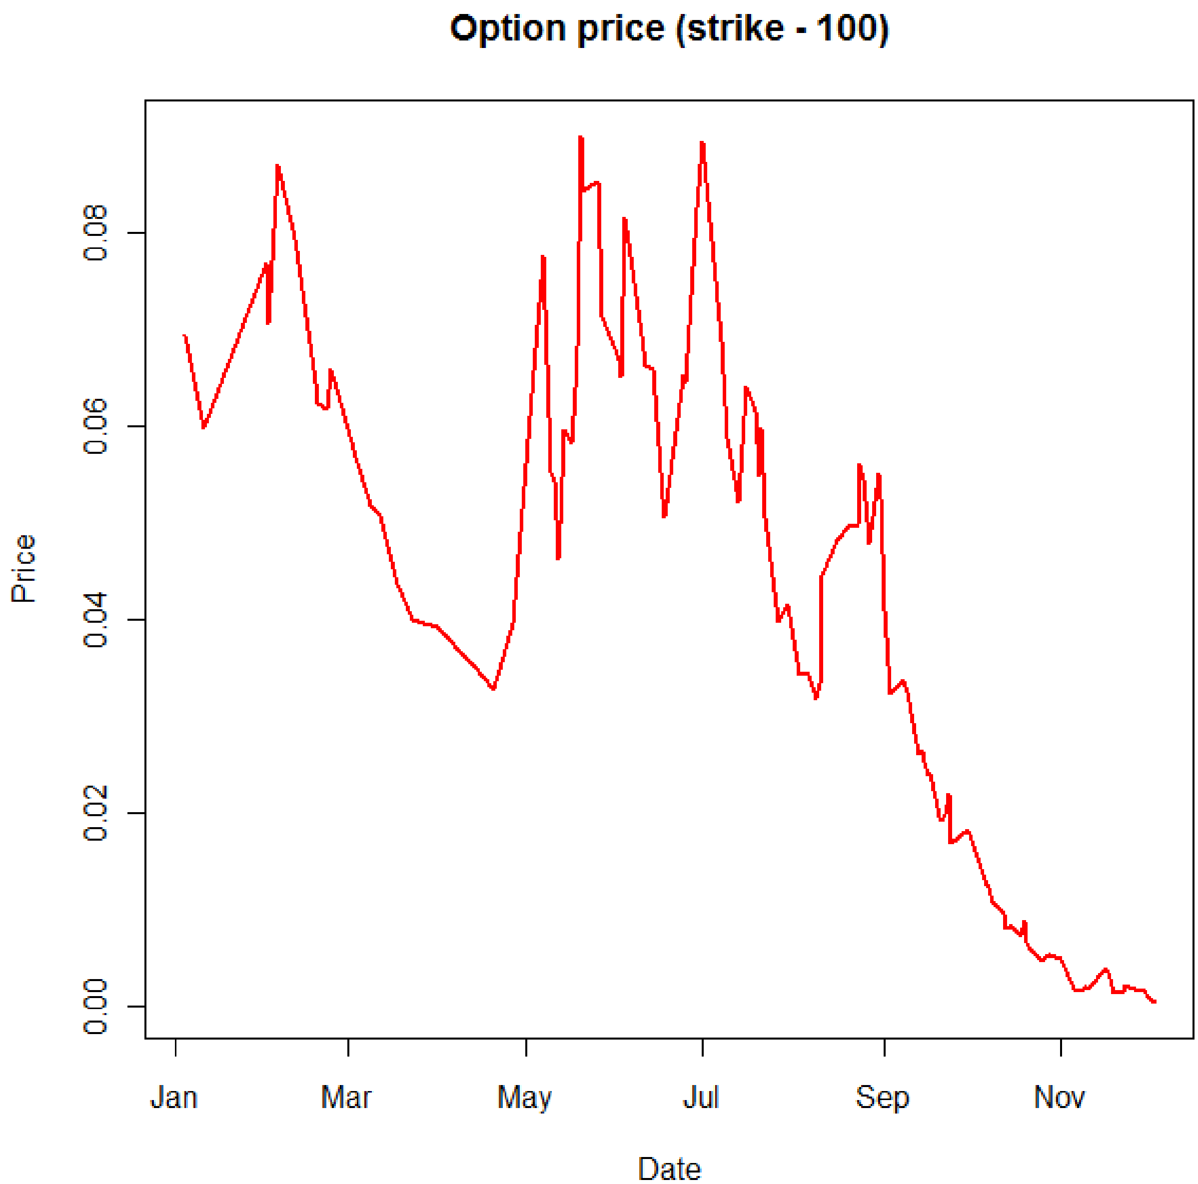

Figure 21

Figure 21 shows the path of the option price (midpoint of the bid-ask range at the end of the day) over 2010. Compare this to Figure 18 panel 1 – the general shape of the plot is similar. For both types of contingent securities prices decline up to May, followed by a period of higher prices and higher volatility until September; however, option price decreases more smoothly from September through December – probably due to greater trading volume of the contract. All in all, it is encouraging to observe similar patterns in both markets, especially because both contracts are contingent on very similar events – DJIA to be less than 10,000 on the 18th or 31st of December (put option with strike 10,000 becomes valuable only if DJIA is less than that) for options and prediction contracts respectively.

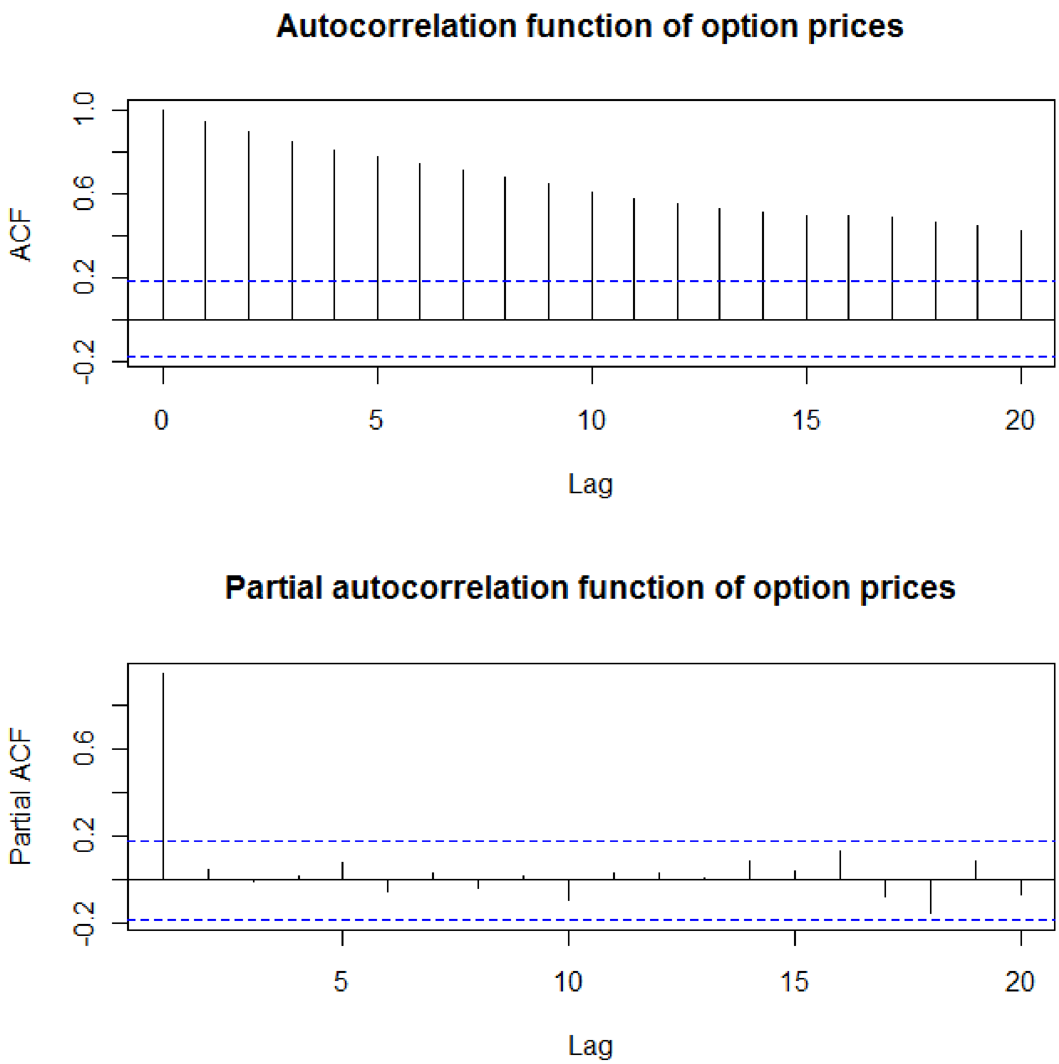

Figure 22

Figure 22 shows ACF and PACF graphs for the option price – they are very similar to both DJIA and prediction market ACF and PACF graphs. Evidence that option prices are a random walk is supported by both ADF and KPSS tests.

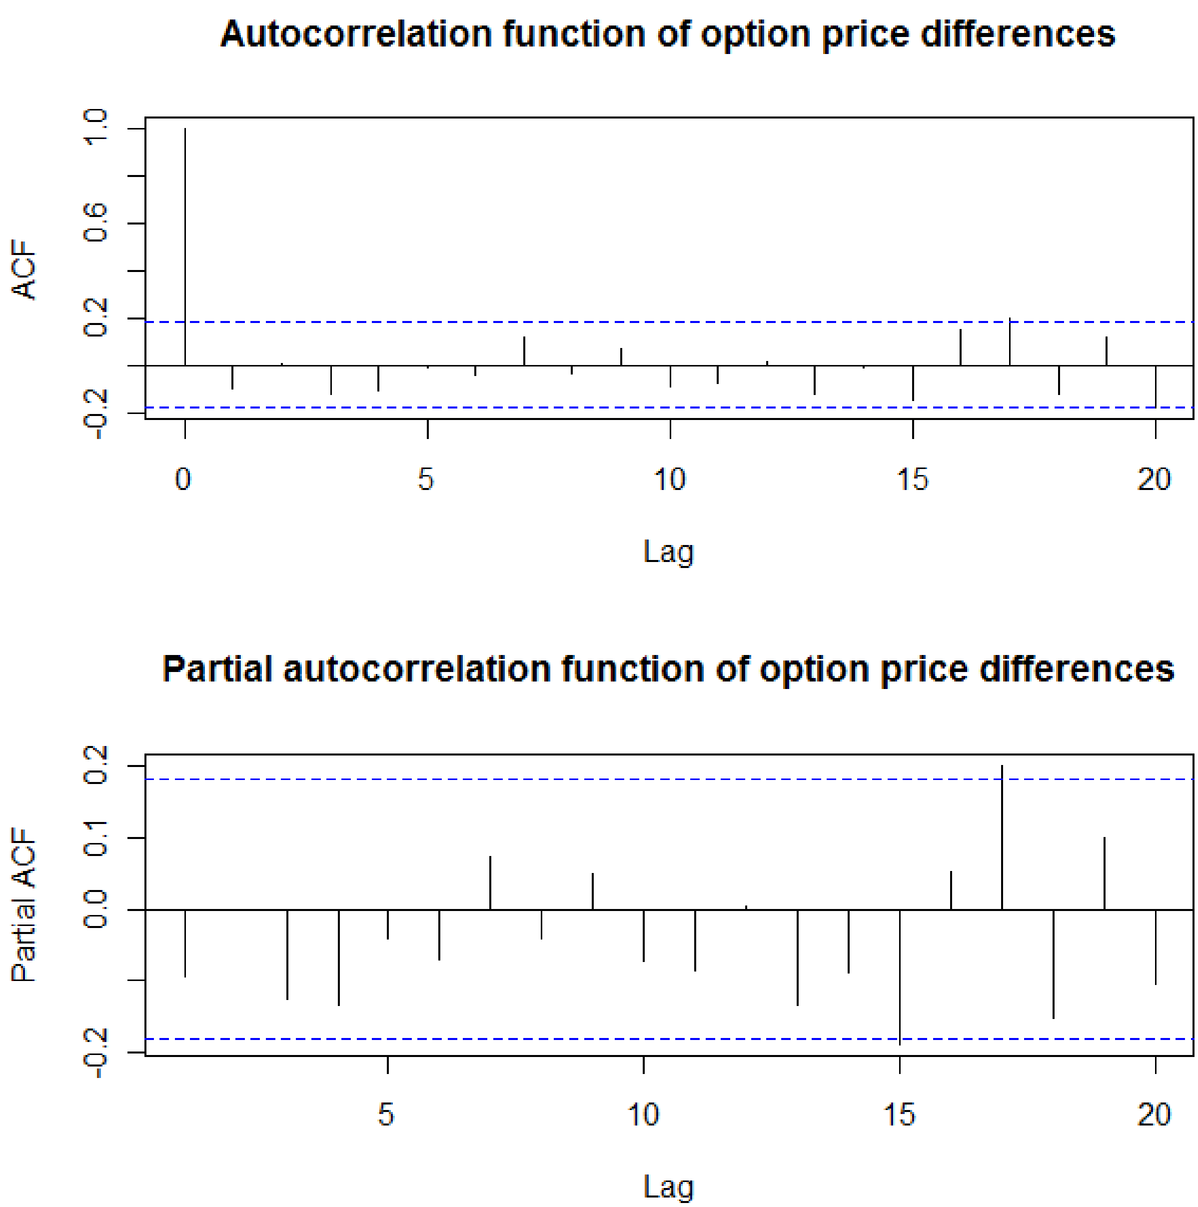

Figure 23

Figure 23 shows ACF and PACF graphs for differenced series of the option price. ACF shows no predictability, meaning that future changes in option price cannot be extracted from past changes in price. PACF demonstrates statistically significant spikes at lags 15 and 17; however, there are a few reasons why these can be dismissed as statistical flukes. First, spikes are not very significant – just over the threshold of 5%. Second, there is no theoretical reason why lags 15 and 17 should be important in forecasting future behaviour of option prices. Last, due to the nature of the time series –not traded every day – lags 15 and 17 fall randomly above the calendar interval of 19 and 21 days, respectively. Finally, ADF indicates that null hypothesis of non-stationarity should be rejected at less than 1% significance in preference of stationarity; moreover, KPSS does not reject stationarity (at more than 10% significance). In short, I find that option prices also follow a random walk.

Summary of random walk testing

To summarise: all three time series – stock, prediction contract and option prices – follow random walk and the series are difference stationary. This indicates that all three markets are weak-form efficient – future prices cannot be forecast using historical data and a simple model.

Expected values

It is probably useful at this point to take a small detour to discuss empirical evidence of predicting expected value of DJIA as discussed in Section III (*Expected values and volatility*). Using the presented method (and subject to noted software constraint that some of the sample dates are discarded due to sometimes negative PDF (because of data and interpolation)) I find that actual DJIA values and prediction market expected values for the contract expiry are correlated (correlation coefficient of 49%). Figure 24 presents the findings in a simple graph – the series appears to be more correlated in the second half of the sample than in the first. Some of the discrepancies can be explained by transaction costs associated with the “execution” of the contract – an investor in a contingent asset requires a positive return after transaction costs. Also a very important reason in this case is the time value of money – I use Moody’s Aaa interest rates as a proxy to make necessary adjustments, whereas the market could be using an entirely different interest rate. Table 1 presents the basic characteristics of both time series.

Table 1

	Observations	Mean	Standard deviation	Median	Min	Max	Skew	Kurtosis
DJIA actual value	165	10,699	425	10,680	9,686	11,499	0.05	-0.96
Expected DJIA	162	10,620	378	10,751	9,738	11,268	-0.37	-1.00
Standard deviation of E[DJIA]	162	1,094	199	1,173	506	1,309	-1.70	1.84
Standard deviation of E[DJIA] (%)	162	10.76%	2.21%	11.62%	4.48%	13.26%	-1.58	1.44

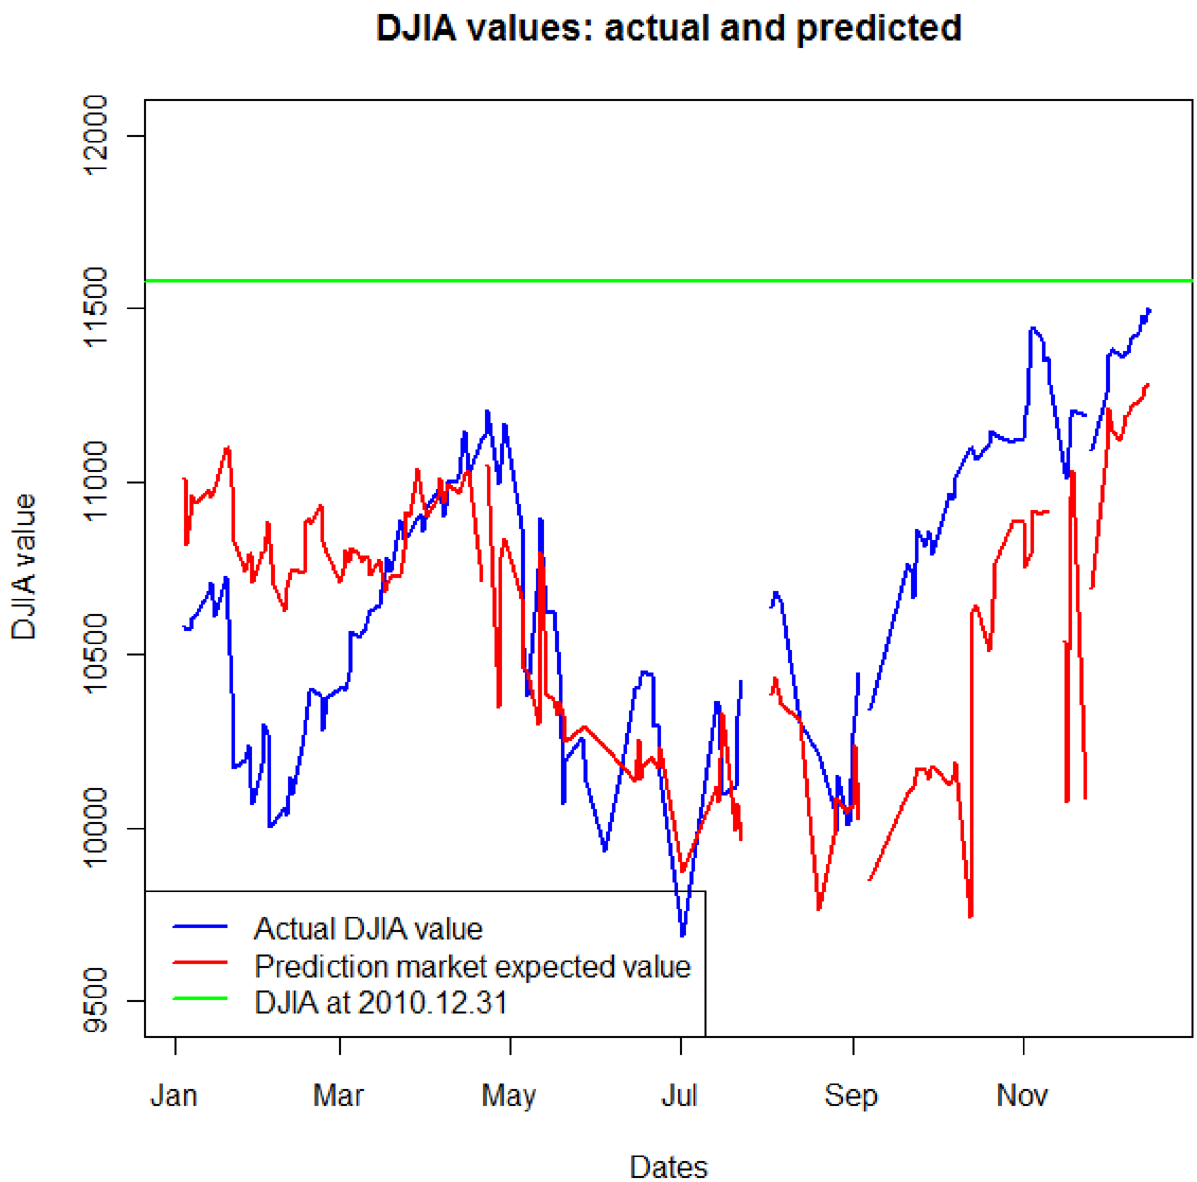

Figure 24

Using the more formal method of least squares linear regression I find that current DJIA value explains 60% of the variation in prediction market expected DJIA value. Furthermore, current DJIA value is a statistically significant explanatory variable of DJIA expected value at 1%

significance level. If the current DJIA value and predicted DJIA value move together I also expect that the residuals of the regression will have a mean of zero. I test this hypothesis using the t-test and find that H_0 hypothesis (residuals of mean zero) cannot be rejected (p-value is reported as 1).

Figure 25 provides the best fit line for the data discussed; 45 degree line illustrates the negative bias associated with the relationship. In short, prediction market data does not provide information on expected DJIA values in the future due to reasons discussed in Section III.

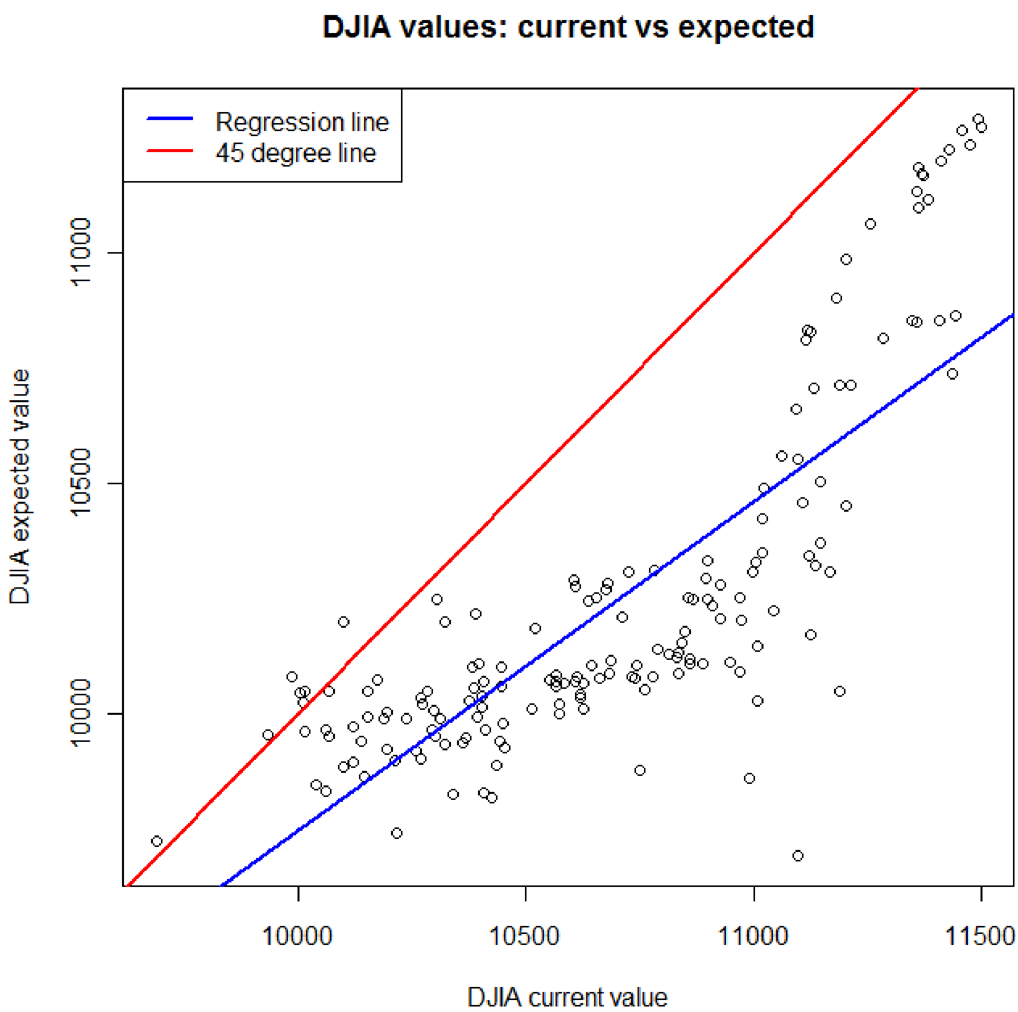

Figure 25

DJIA returns

Before further discussions of volatility, I present the DJIA returns series which is the object of implied volatility series. As discussed before, returns are calculated as $r_t = ln\ (S_t/S_{t-1})$, where S_t is DJIA value at time t; this method is used because it gives the continuously compounded rate of interest – asset prices change continuously with new information but are observed only at certain intervals, e.g. during trading hours. Also, as discussed in *Asset price generating process and volatility* part, the returns should be normally distributed with some mean (usually positive) and constant standard deviation. This is one of the assumptions of the B-S model of option pricing; therefore it is useful to test returns normality, because accuracy of B-S model is the basis calculating B-S implied volatilities.

Figure 26 graphically presents evidence that, at least during 2010, DJIA logarithmic returns do not follow a normal distribution. The blue line shows how a normal distribution with the actual mean and standard deviation of the returns series should look like, the bar graph shows the histogram of the series and, finally, the red line shows kernel density (assuming Gaussian distribution) – essentially how the smoothed histogram would look like. It is obvious observing the simulated normal distribution and the kernel density graphs that returns do not follow the normal distribution. Also examining Q-Q plot of the returns (Figure 27) provides more evidence against returns normality. If the returns were indeed normally distributed, the scattered points would fall along the diagonal line – this is clearly not the case. I further test normality using the Shapiro-Wilk test (R Development Core Team, 2011): null hypothesis – sample is normally distributed. Test p-value is 3.831e-07, which means null hypothesis is rejected.

Calculating the mean and standard deviation of the returns I find that they are 0.04% and 1%, respectively. Testing null hypothesis that the true mean is zero results in p-value of 52% - do not reject null hypothesis. However, even if the mean of daily logarithmic returns is equal to zero, this gives annualised returns of 11% (actual change over the 2010 – 9.4%).

In summary, evidence strongly suggests that DJIA returns are not distributed normally; therefore, another one of the assumptions of B-S model is violated – B-S implied volatilities should not be used to forecast volatility.

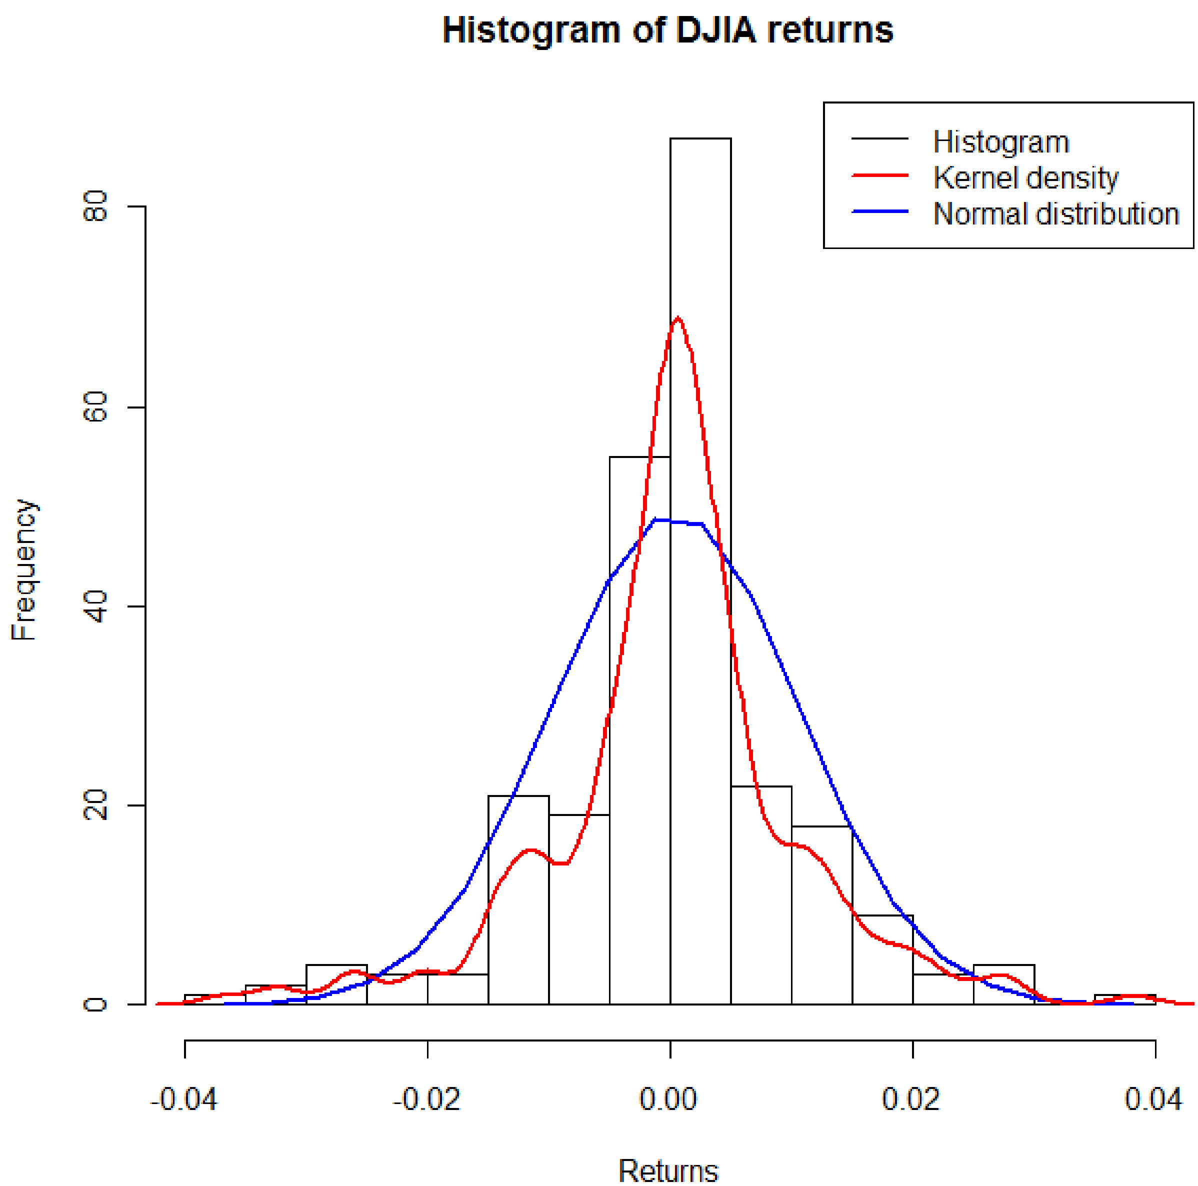
Histogram of DJIA returns
Histogram
Kernel density
Normal distribution
Frequency
0
20
40
60
80
-0.04
-0.02
0.00
0.02
0.04
Returns

Figure 26

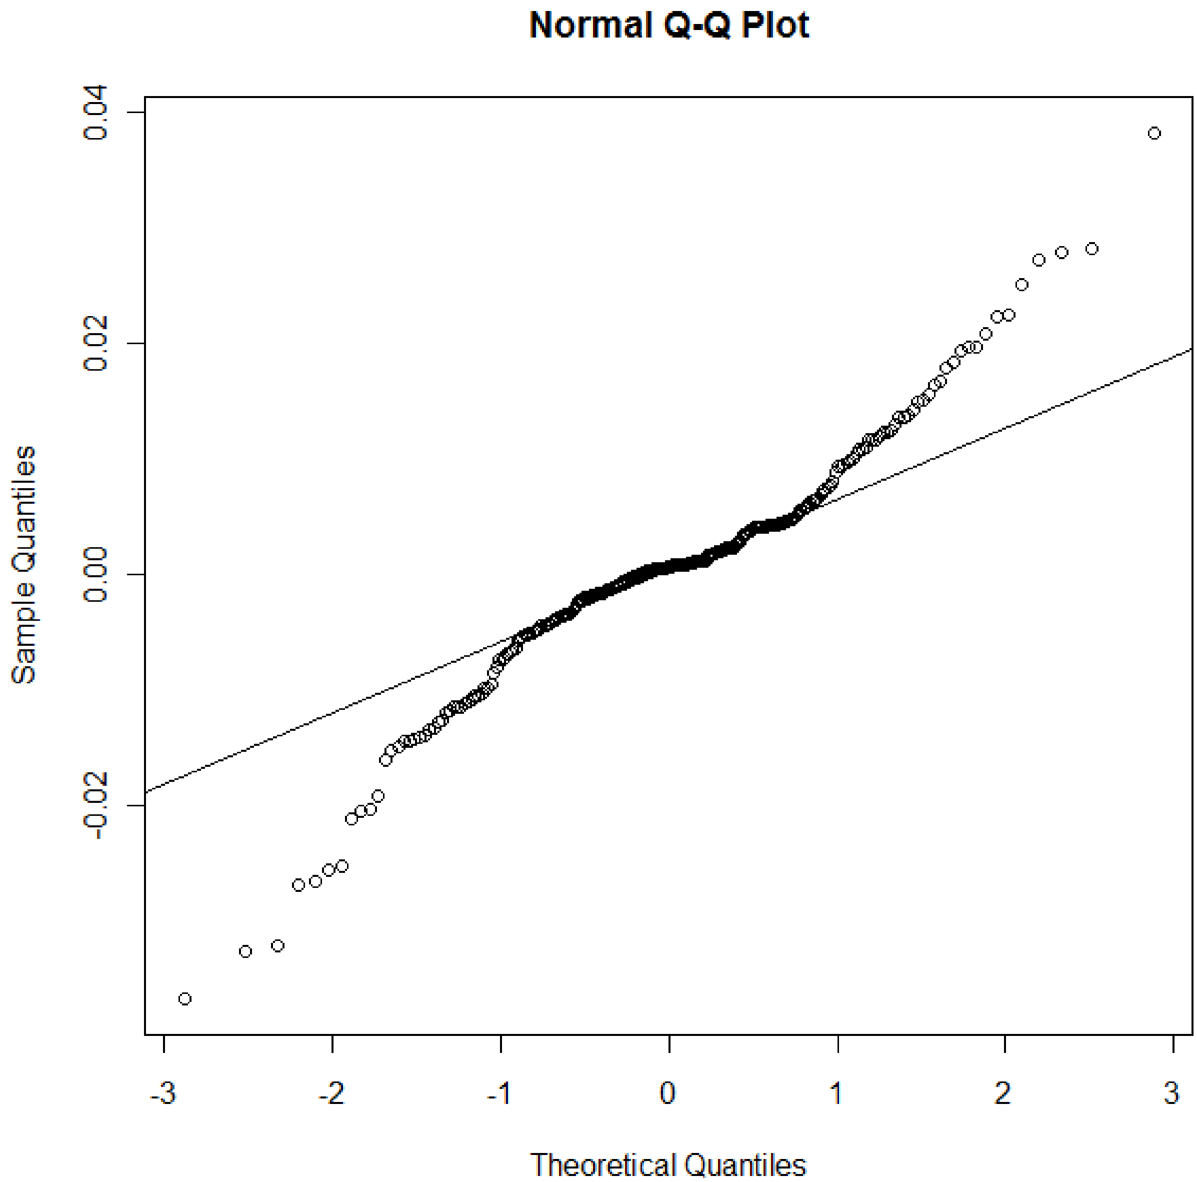

Figure 27

Standard deviations

In this part I present the standard deviations calculated using previously described techniques.

Note that all volatilities are calculated annualised and are matched by date for all three markets.

Predicted standard deviations

Probability density functions

Figure 28 shows the plot of the areas under the prediction market probability density functions (after interpolation using cubic splines) for all trading days – PDF maps the values to probabilities that those values will materialise. These areas show how likely it is that DJIA value will fall between 7,000 and 12,500 points. The red lines illustrate arbitrary cutoff points – area under the PDF should not be less than 0.5, because that shows that a large part of the distribution is outside the observed interval (7,000 to 12,500); whereas area under PDF greater than 1 indicates problems with interpolation (contract prices are sampled at only 12 points, so there is a lot of leeway for PDF in between those points), because the sum of probabilities cannot be greater than 1.

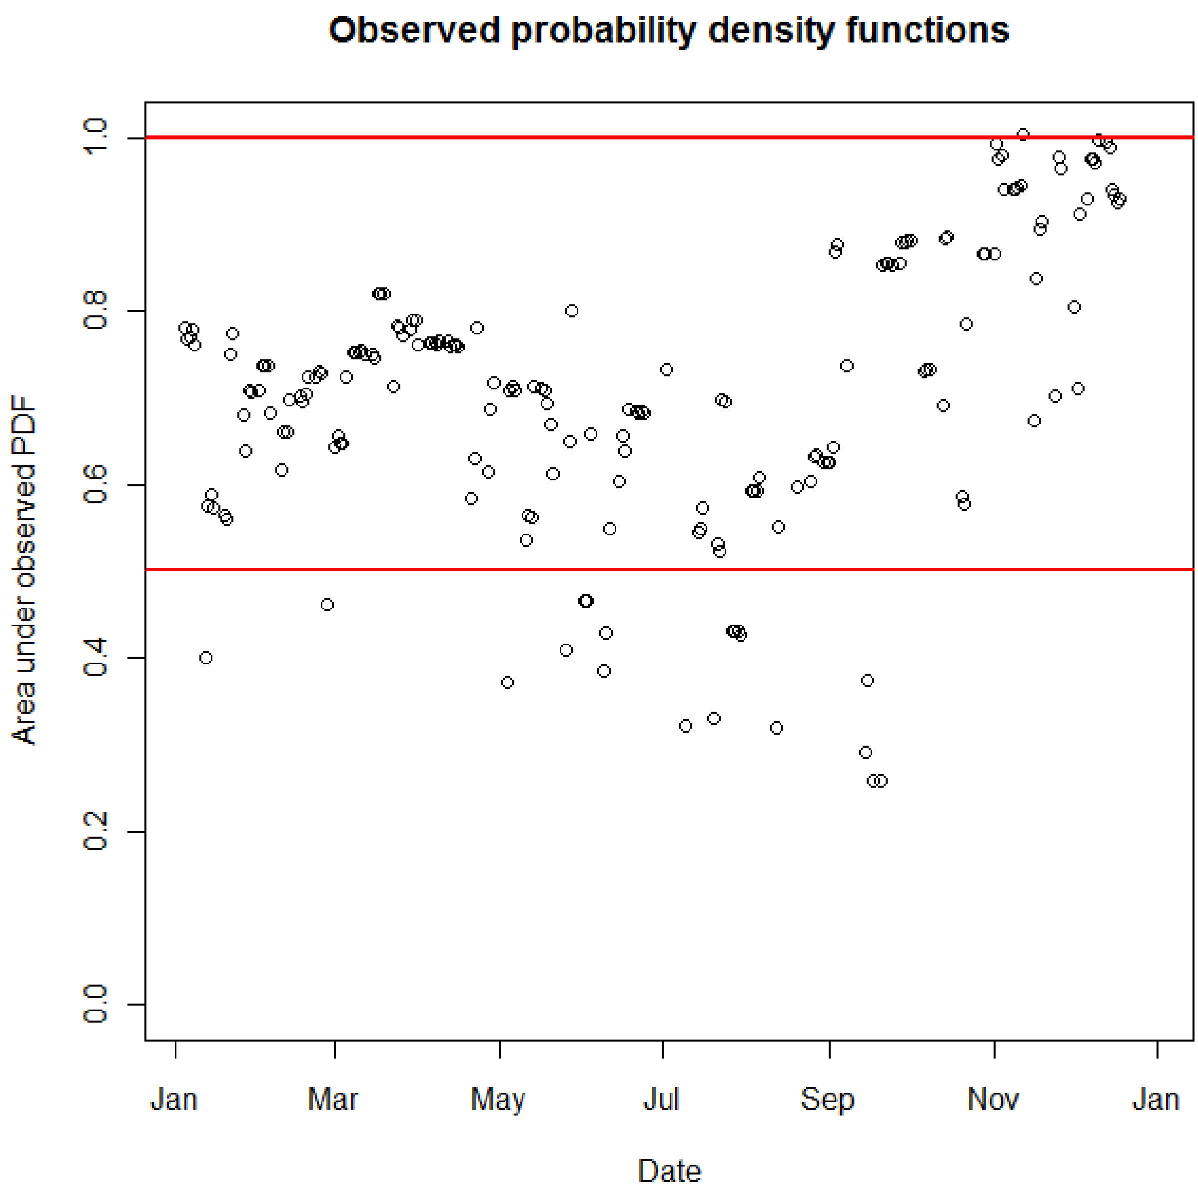
Observed probability density functions
Area under observed PDF
0.0
0.2
0.4
0.6
0.8
1.0
Jan
Mar
May
Jul
Sep
Nov
Jan
Date

Figure 28

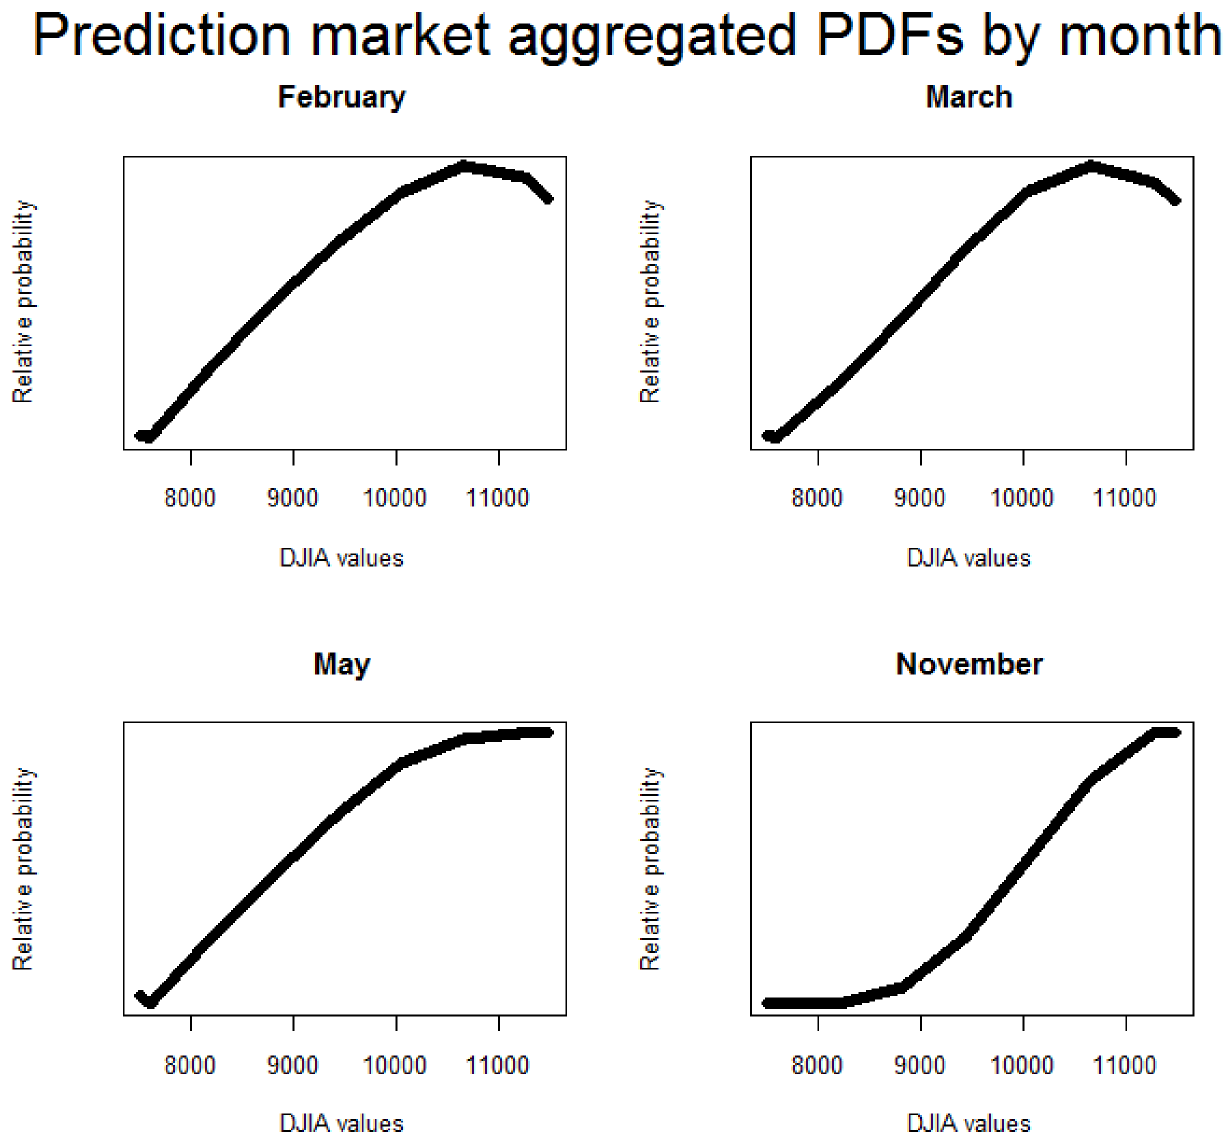

Figure 29

Figure 29 shows a sample of probability density functions for DJIA values calculated from February, March, May, and November values (these four months were picked because they include dates where trading occurred on all 12 contracts – this improves the calculations). As expected, the further from the prediction, the higher dispersion of the distribution. On the other hand, most of the PDFs are left unobserved (market contracts went up to 15,000 but the data were not purchased due to low liquidity). It appears that 38% of the total probabilities are left outside the graphs. It can also be observed that PDFs evolve over time (Figure 29 provides 4 snapshots) to become more concentrated – market becomes more certain about the expected value as contract horizon approaches because less time is available for the price to deviate from the current price (this is also explained in *Literature review*).

Finally, using Kolmogorov-Smirnov test for whether two data sets were drawn from the same distribution, I find that individual PDFs of expected DJIA values are not normally distributed (H_0: E[DJIA] PDF follows a normal probability density function). However, this result is not very surprising because the Kolmogorov-Smirnov test rejects the null hypothesis even when testing large artificially generated samples with a known distribution. Consequently, formal testing does not verify that expected DJIA values follow a normal distribution (lognormal – even less so).

Predicted standard deviations

Figure 30 shows time series of standard deviations of the expected values of DJIA on 2010.12.31. This is a nonparametric estimate of how the market believes the DJIA expected values to be spread out. Naturally, the further away the forecast horizon, the less market is certain of the expected value – the higher the dispersion of expectations. Surprisingly, the standard deviation fluctuates at a stable level of around 12% (using absolute standard deviation does not yield materially different results), and then suddenly drops as the expiration date approaches. This is most likely associated with a large increase in prediction market trading volume (see Figure 10). Table 1 shows that the average predicted standard deviation over the whole sample is 10.8% (median is 11.6% - skewness due to dramatic fall in standard deviation in October). In short, the standard deviation of DJIA expected values predicted by prediction markets is quite stable over the sample period.

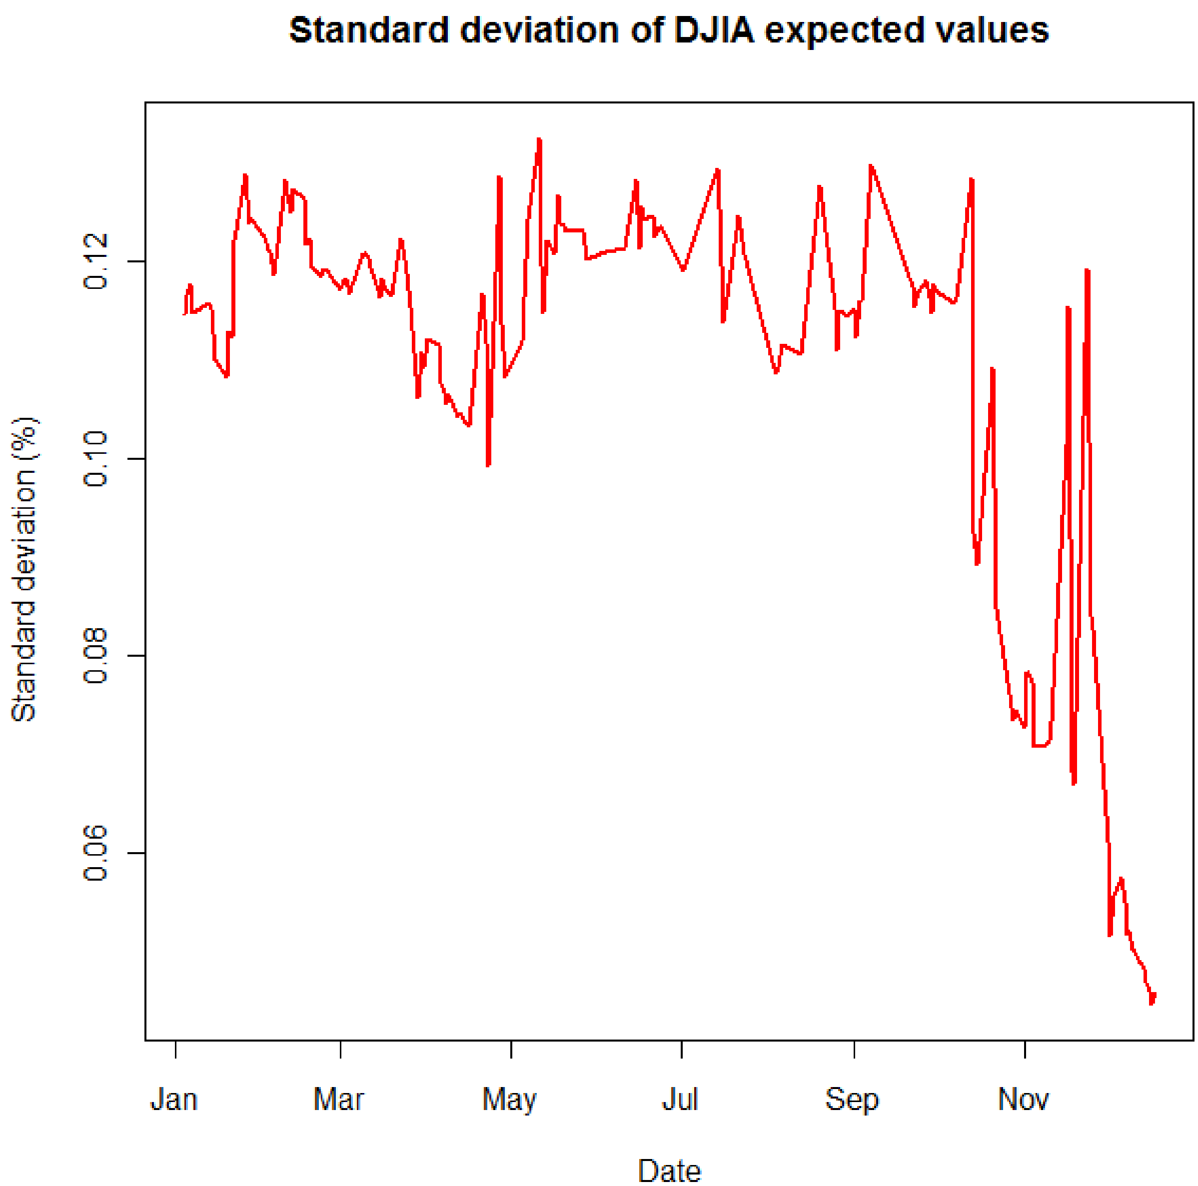

Figure 30

Implied volatility

Implied volatility was calculated using two different methods – Black-Scholes formula and model-free volatility method. I first present these two volatility measures separately, then I analyse them together.

Table 2

	Observations	Mean	Standard deviation	Median	Min	Max	Range	Skew
DJIA	252	10,669	460	10,627	9,686	11,585	1,899	0.13
DJIA log returns	252	0.04%	1.02%	0.07%	-3.67%	3.82%	7.49%	-0.18
Realised volatility	252	9.61%	4.98%	9.74%	0.00%	16.16%	16.16%	-0.30
M-F implied volatility (puts)	174	13.66%	7.02%	13.70%	0.00%	32.96%	32.96%	0.08
M-F implied volatility (calls)	88	9.37%	5.70%	7.77%	0.00%	26.43%	26.43%	1.15
B-S implied volatility (puts)	236	22.25%	4.86%	22.32%	1.53%	41.95%	40.42%	-0.12
B-S implied volatility (calls)	192	16.72%	3.56%	16.21%	0.03%	28.20%	28.18%	-0.09

Table 2 presents general descriptive statistics of DJIA, DJIA logarithmic returns and calculated volatilities. Implied volatility series have fewer observations because I require trading to calculate implied volatility (see *Data selection*) – this is most apparent on M-F volatilities calculated from call options because they rarely traded, therefore they have few traded strike prices per day; M-F volatility cannot be estimated without at least 4 traded options on a particular day (due to interpolation method), which leads to greatly smaller number of observations (B-S implied volatility can be calculated from one traded option).

B-S volatility

Figure 31 shows implied volatility measures calculated using B-S model on at-the-money options from put and call options, respectively. It is apparent that the two measures give significantly different implied volatilities, e.g. correlation coefficient for put and call implied volatilities is 56% (after eliminating non-paired observations). Combining these two measures by taking some average would lead to unpredictable results; therefore henceforth B-S implied volatility will be calculated using put options. Put options are used in preference to call options because they are better traded – more liquid market and more dates with trades.

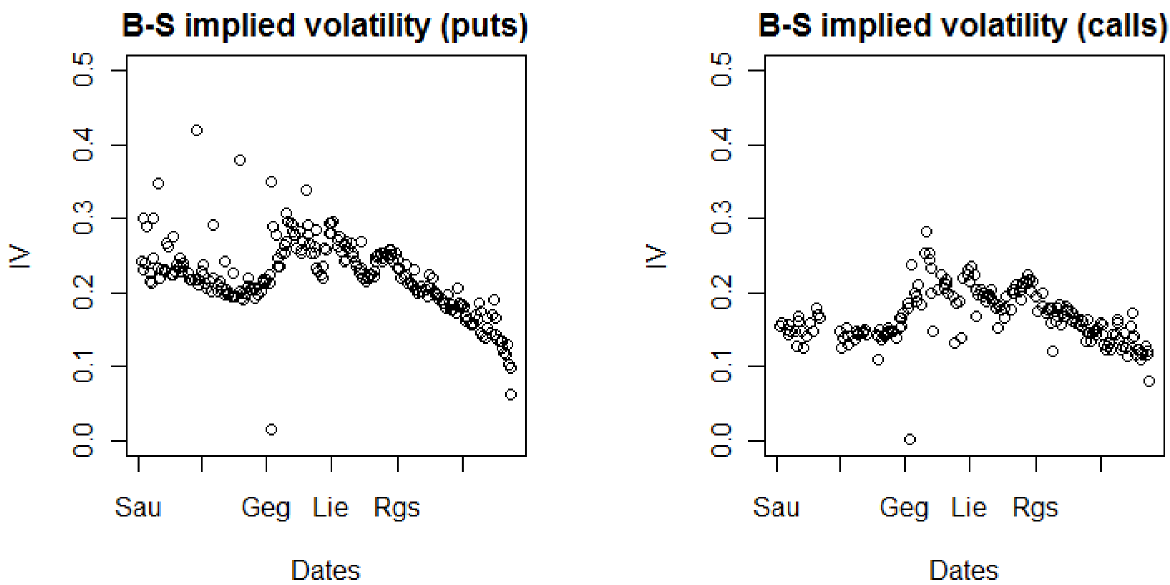

Figure 31

M-F volatility

M-F implied volatility evolution over 2010 is plotted in Figure 32. Again, as in B-S graph (Figure 31), calculations based on puts and calls produce different results (e.g. correlation coefficient for put and call implied volatilities is 27%) which are, nevertheless, similar in magnitude (33% and 26% maximum implied volatilities for puts and calls, respectively). However, it is again apparent that M-F volatilities from puts are superior to calls, because of more puts data and less dispersion (this is desirable, because it can be argued that "real" implied volatility should not fluctuate greatly from day to day). Consequently, I discard implied volatility data calculated from call options (both B-S and M-F) and focus further analysis on put option information.

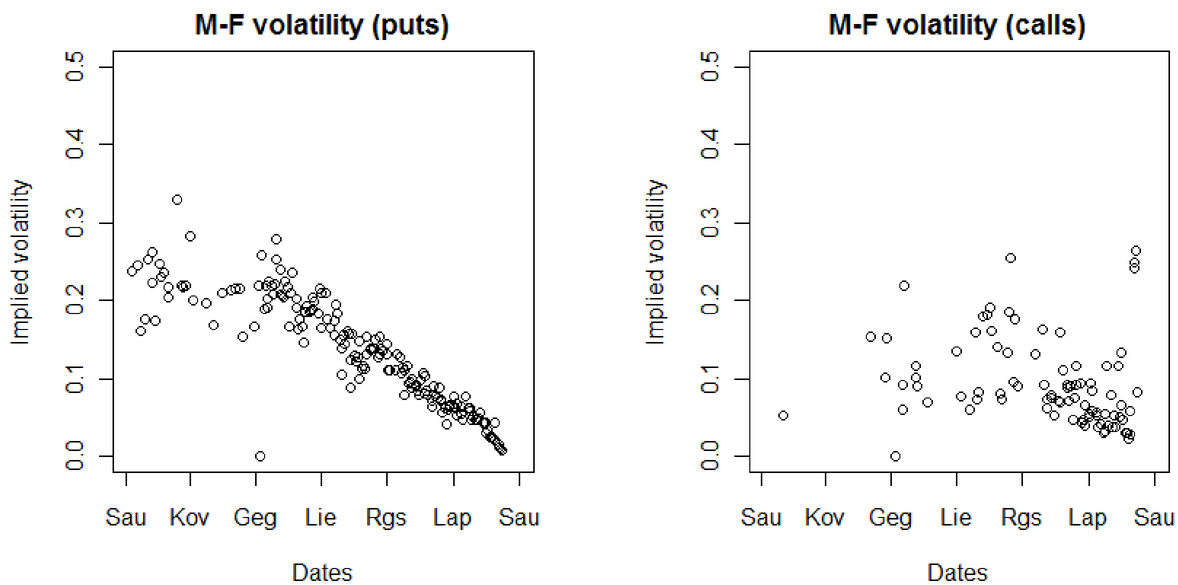

Figure 32

B-S and M-F volatilities

In order to compare like-for-like, Figure 33 shows the previously discussed put options graphs side by side. Visual observation shows that B-S implied volatility is higher than M-F estimates (contrary to Jiang and Tian (2005), however, they make adjustments to the series to account for known biases).

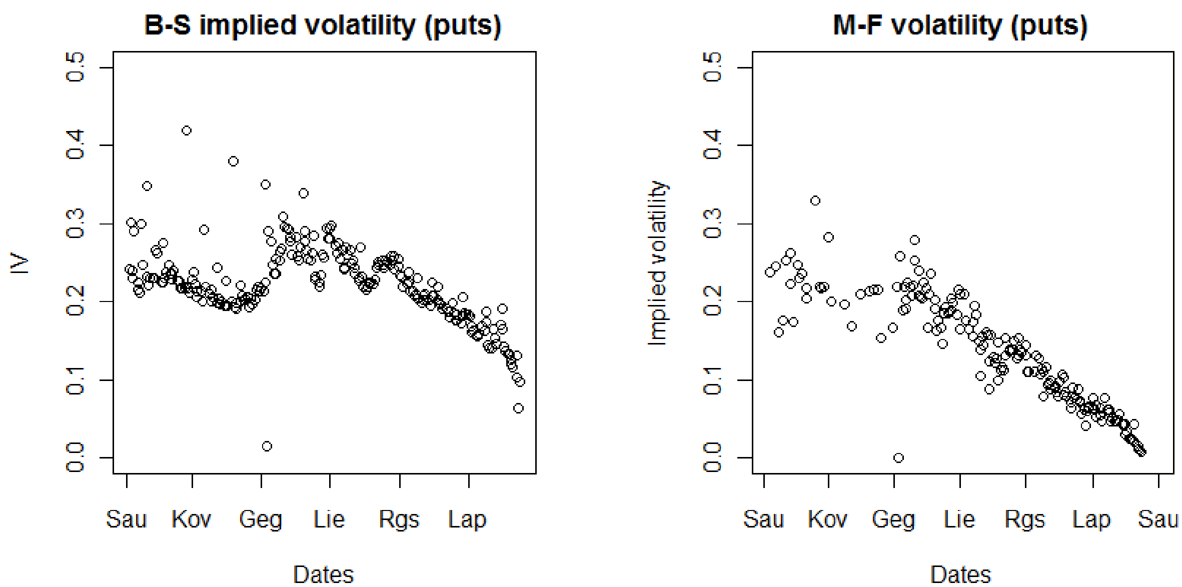

Figure 33

The same information is presented in a single graph (Figure 35) – the relationship between M-F and B-S volatilities is even more apparent. Excluding the period until June, 2010 (not enough data for M-F volatility calculations, because at least 4 distinct traded options are required to calculate M-F volatility – this is required cubic spline interpolation method), both series follow very similar trajectories, but B-S implied volatility is higher.

Realised volatility is also shown in the graph – it converges with the M-F implied volatility at the end of the series. Furthermore, and more interestingly, implied volatilities and realised volatility consistently converge through the period – this could indicate higher risk premiums for events further in the future.

Figure 34 shows a scatterplot diagram of M-F implied volatilities plotted against B-S implied volatilities (even though the series exhibit exponential behaviour, taking logarithms decreases the fit of the model) and also the best-fit-line for the data. Adjusted R^2 for this model is 57% indicating quite high but not perfect co-movement of implied volatilities, correlation coefficient is 76%.

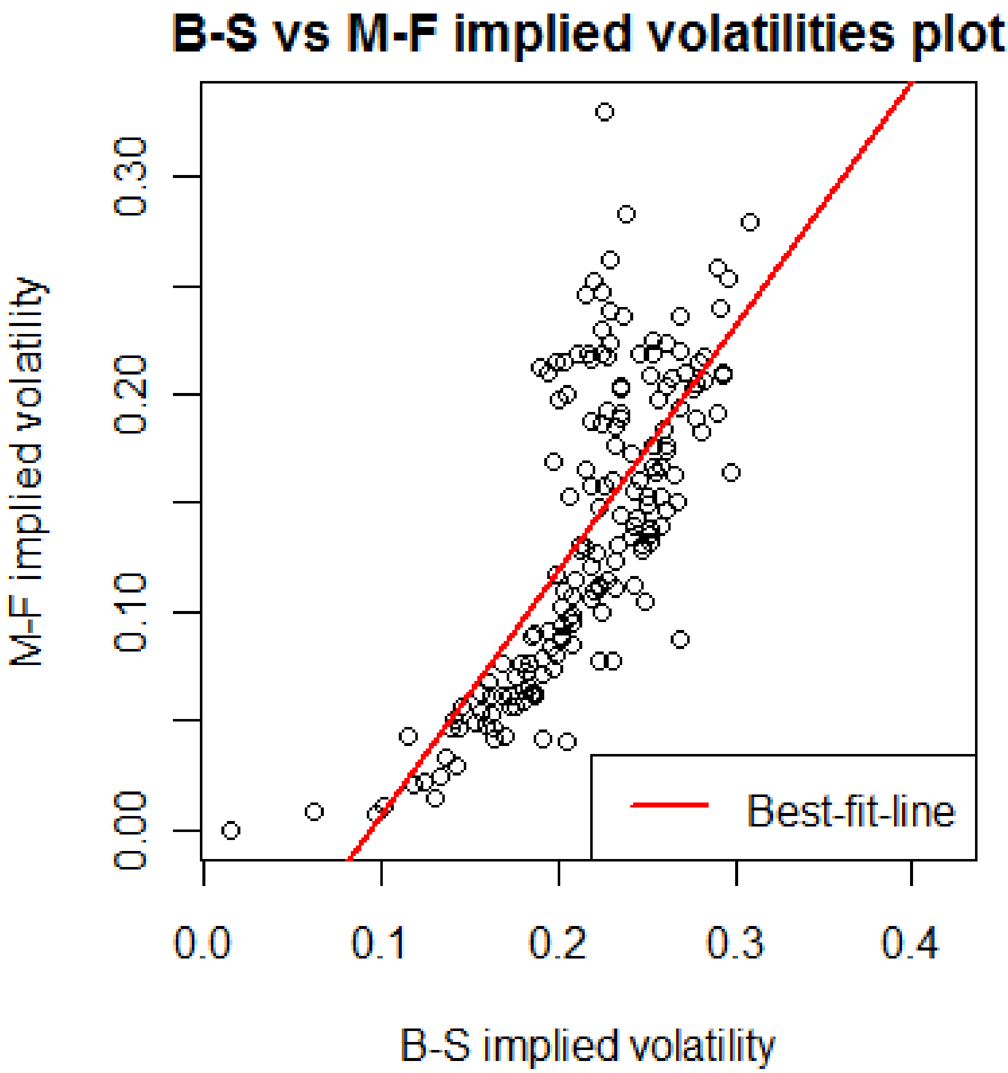

Figure 34

Analysing the data more formally, I find that realised volatility and B-S volatility correlation coefficient is 50%, whereas realised and M-F volatilities are correlated at 90% - this suggests that M-F volatility should be a better indicator of volatility.

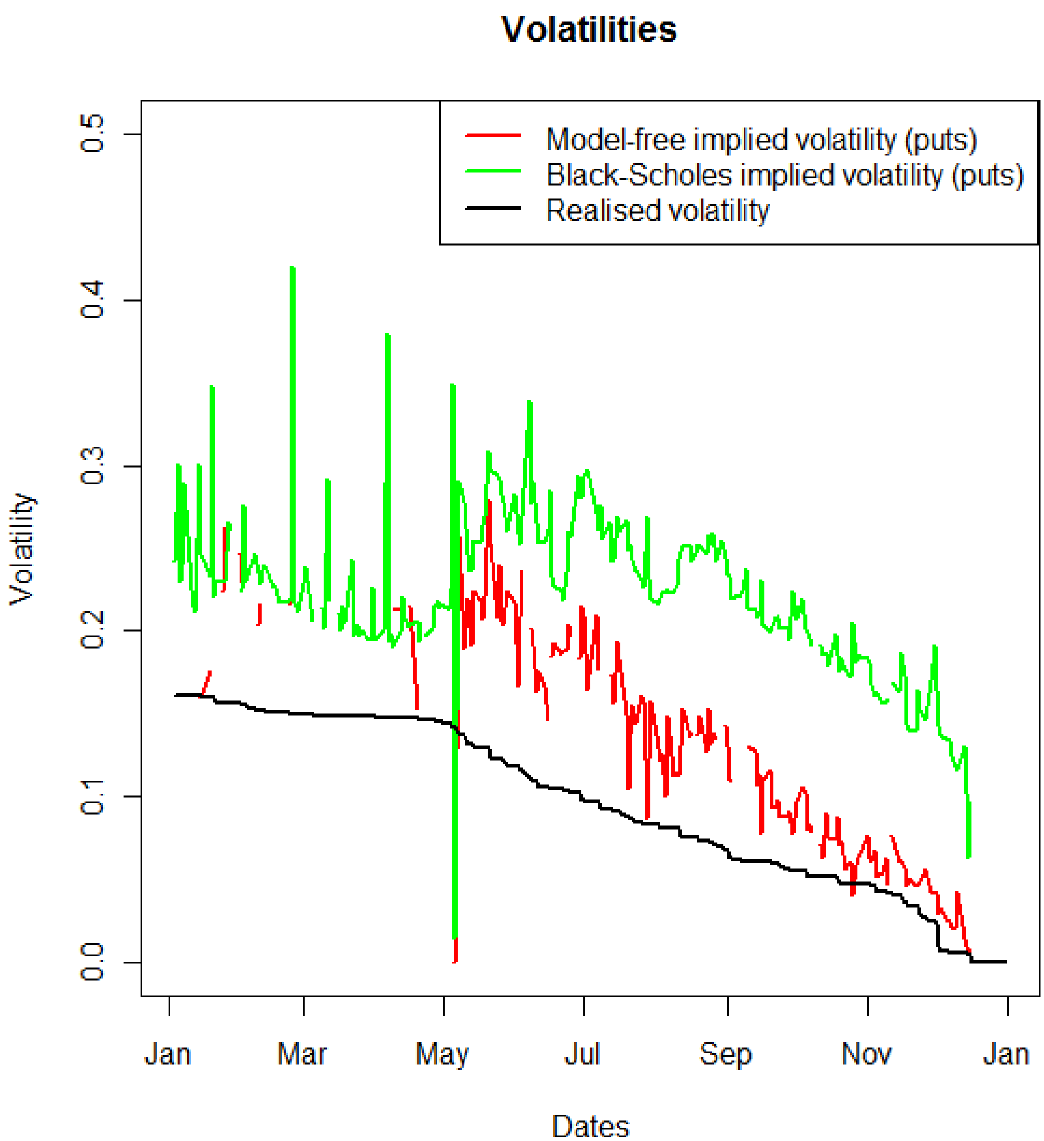

Figure 35

Regressions of volatilities

Regressing realised volatility on various combinations of implied volatilities I find some interesting results. The main characteristics of the models are presented in Table 3.

Table 3

Model name		Degrees of freedom	Intercept	M-F IV (puts)	B-S IV (puts)	M-F IV (calls)	B-S IV (calls)	Adjusted R^2
F	Estimate	78	0.0548	0.6911	0.2640	-0.0099	-0.6917	79.9%
	Standard error		0.0097	0.0642	0.1302	0.0322	0.1574	
	p-value		0.0000	0.0000	0.0459	**0.7587**	0.0000	
F-	Estimate	147	0.0443	0.6200	0.0774		-0.3592	84.6%
	Standard error		0.0070	0.0336	0.0807		0.0808	
	p-value		0.0000	0.0000	**0.3391**		0.0000	
P	Estimate	171	0.0351	0.6442	-0.1770			81.5%
	Standard error		0.0075	0.0310	0.0462			
	p-value		0.0000	0.0000	0.0002			
P.M-F	Estimate	172	0.0091	0.5542				80.0%
	Standard error		0.0032	0.0211				
	p-value		0.0053	0.0000				
C	Estimate	85	-0.0207			0.0056	0.5038	24.9%
	Standard error		0.0152			0.0650	0.0957	
	p-value		**0.1757**			**0.9321**	0.0000	
C.B-S	Estimate	190	0.0631				0.1713	1.2%
	Standard error		0.0160				0.0938	
	p-value		0.0001				**0.0693**	
P.M-F.C.B-S	Estimate	148	0.0470	0.6421			-0.2947	84.6%
	Standard error		0.0064	0.0244			0.0448	
	p-value		0.0000	0.0000			0.0000	

Table 3 explanation

Table 3 shows linear model fitting results where realised volatility is a dependent variable and various combinations of implied volatilities (IV in the table) are explanatory variables. Reported in this table are estimates of the regression coefficients, their standard errors, p-values (probability that true regression coefficient is actually 0, given observed values), degrees of freedom of the model (number of observations minus the number of explanatory variables and intercept), and the adjusted R^2 – how much of variability in dependent variable is explained by the explanatory variables, adjusted for the number of regressors.

The models are named in order to facilitate easier exposition: model F is the "full" model – realised volatility regressed on all implied volatility measures, model F- – same as model F, but M-F implied volatility from call options is excluded, model P – implied volatilities from put options are used as regressors, model P.M-F – single explanatory variable – M-F implied volatility calculated from put options, model C – regressors are implied volatilities calculated from call options, finally, model C.B-S is using B-S implied volatility as an explanatory variable.

It is interesting to note that most of the p-values are near zero, consequently, the only statistically not significant regression coefficients were marked in red (5% significance level was used according to convention). Most other coefficients are significant even at 1% significance (except B-S implied volatility coefficient in model F).

Discussion of regression results

Degrees of freedom of the regressions show that there is a sufficient amount of data points for the regressions (it is usually accepted that at least 30 degrees of freedom are required for hypothesis testing – central limit theorem to hold (Smith & Wells, 2006)) – minimum degrees of freedom is 78. The "bottleneck" in this case is M-F implied volatility series from call options (as explained in part *Implied volatility*). Moreover, out of the two reported models, which include M-F implied volatility from calls (F and C models), the volatility is not statistically significant. This in part justifies the decision to use put options for implied volatility estimations. Another consequence of omitting data points where at least one regressor has no value is that it becomes difficult to compare regression results. For example, models C and C.B-S give different regression coefficients for B-S implied volatility, because model C uses 88 observations, whereas C.B-S uses 192; furthermore, for the same reason model C gives R^2 of 25%, whereas C.B-S – 1% (correcting for the change in number of observations I find that C.B-S R^2 would be 26% for the same 88 observations as in model C).

Model F achieves a high adjusted R^2 value – 80% of variation in realised volatility is explained by the different measures of volatility, however, this model includes a statistically insignificant parameter (M-F implied volatility from calls) – it should be made more parsimonious. Dropping the parameter gives model F- (R^2 increases to 85%) which also has almost twice as many degrees of freedom. The increase in the number of observations also increases the fit of the model (R^2 after adjusting for change in the number of observations is actually the same for F and F-), meaning that M-F implied volatility from calls is actually a useless explanatory variable, because it is not significant and reduces the number of available observations.

Comparing models P and P.M-F is very illuminating, because I find that dropping B-S implied volatility reduces the fit of the model by a small amount (R^2 change - 1%), but the model becomes more parsimonious. This indicates that B-S implied volatility from puts could be dropped from the regression without much loss – information from put options is almost fully subsumed by M-F implied volatility – B-S implied volatility adds only a little.

The best fit is provided by P.M-F.C.B-S model (realised volatility regressed on M-F implied volatility from puts and B-S implied volatility from calls). However, this model presents a theoretical problem – explaining negative regression coefficients of implied volatilities is theoretically difficult – there are no apparent reasons why higher B-S implied volatility should reduce realised volatility when the other independent variable is M-F implied volatility.

Following Jiang *et al.* (2005), if any calculated implied volatility estimate was an unbiased estimator of realised volatility, then intercept would be equal to zero and the regression coefficient of that regressor would be equal to 1. The formal Chi-square hypothesis tests results are not presented here because looking at Table 3 it is obvious that none of the regressors even comes close to being equal to 1, therefore, it is impractical to present test results on two parameters.

None of the regressions yield an unbiased estimator of realised volatility; however, the objective of this part is to investigate information available on the options market – not find the best method to forecast realised volatility. I conclude that the best source of information about future volatility is the model-free implied volatility because (a) it achieves a high degree of fit (80%), (b) the fitting model is parsimonious, and (c) most importantly – it does not rely on model assumptions.

Overview of available information

As presented above, two distinct and useful forecasts of the stock market emerge when prediction markets and option markets are examined. First, it is the standard deviation of DJIA expected values – a unique straightforward measure of dispersion of future values of DJIA. Second, option markets allow the extraction of DJIA returns volatility. Both estimates are nonparametric (B-S implied volatility is used as a benchmark for M-F performance), which allows testing various hypotheses using this information without testing the models used to calculate these estimates.

Discussion

In summary, I find that all analysed markets – stock, prediction, and options – are "well-behaved", i.e. asset prices follow random walk. Building on that foundation I use prediction market data to calculate expected DJIA values and standard deviation of those expected values. Analysis shows that expected values differ from current DJIA values, but that is probably a result of transaction costs and time value of money. I also find that expected DJIA values do not follow normal or lognormal distributions. Finally, standard deviation of the expected values is quite stable for most of the sample period, but drops abruptly as forecast horizon approaches.

Using techniques developed by Jiang *et al.* (2005) I also calculate Black-Scholes and model-free implied volatilities and test them for explanatory power of realised volatility. Even though none of the tested models provides an unbiased estimate of the realised volatility, I find that they have quite good explanatory power – up to 85%. I conclude that nonparametric implied volatility calculation is the best estimator for future volatility. On the other hand, Jiang *et al.* (2005) results from M-F implied volatilities and my findings are somewhat different, for instance, the authors find much higher regression coefficients for regressions of realised volatility on implied volatility (M-F and B-F). Of course, the authors use much shorter forecasting horizons, different realised volatility measures and more data. Also, their sample is pre-financial crisis which could have influenced investors' behaviour, consequently influencing my results.

This research is the first to use prediction market data to estimate market sentiment on future stock market realisations. It should add to the theoretical studies of available information on prediction markets (Wolfers & Zitzewitz, 2004) and bridge the gap to studies focusing on information in options markets (Jiang & Tian, 2005).

A further avenue of study could be to use Ait-Sahalia *et al.* (1998) approach to estimate DJIA returns PDFs and, assuming constant returns distribution, estimate the same DJIA expected values PDF as calculated from prediction markets. These two distributions could then be compared to gain further insights into asset prices and processes driving them.

A simpler investigation would be to test if financial outcomes are predicted as well as election results ((Berg, Nelson, & Rietz, 2007); (Rhode & Strumpf, 2004)), however, this investigation needs much more data and, preferably, increased prediction market liquidity. Furthermore, this investigation could be expanded to use more prediction market data – markets are also available for monthly forecast horizons and annual forecasts for previous years.

The shortfalls of this thesis are numerous, but they do not substantially affect the proof-of-concept nature of this paper. First, dividend payments should be included in the analysis – this can be done by calculating implied index values, which take into account market belief about future dividends (this approach is based on Ait-Sahalia *et al.* (1998)), by observing paid out dividends and integrating them ex post into the analysis, or by using historical dividend yield to adjust. Furthermore, higher quality data would greatly improve potential avenues for analysis – quote-level data for options would allow incorporation of more information into model-free volatility calculations (as per Jiang *et al.* (2005)), whereas quote-level data (or intraday observations) of stock prices would greatly improve realised volatility estimate (Andersen, Bollerslev, Diebold, & Ebens, 2001).

A further improvement that could be implemented with more data is the method – proposed by Christensen *et al.* (2001) – to avoid the problem of telescoping data which could cause some problems with regressions.

Conclusion

The aim of the thesis was to improve the understanding of available information on different markets with regard to equities, therefore, research problem was formulated – determine market predictions of future financial events from options and predictions market data. The problem was solved through calculation of DJIA expected values probability density functions from prediction markets and Black-Scholes and model-free implied volatilities from options markets.

Reiterating the previously tested hypotheses results:

1. All three market prices follow random walk;
2. Implied volatility is a good if biased estimator of realised volatility;
3. DJIA expected values do not appear to follow Gaussian distribution;

The main result of the empirical work was to show that it is possible to use prediction market data to construct the DJIA expected values PDF, also, following Jiang *et al.* (2005) model-free implied volatility was calculated for the DJIA. Even though the model-free implied volatility did not fully correspond with the Jiang *et al.* (2005), it is still a step forward in directly comparing information extracted from the two markets.

The goal of the thesis was achieved to some extent, however, this proof-of-concept can be greatly improved if previously discussed shortcomings are corrected and investigation expanded to take into account suggested improvements.

Bibliography

Ait-Sahalia, Y., & Lo, A. W. (1998, April). Nonparametric Estimation of State-Price Densities Implicit in Financial Asset Prices. *Journal of Finance, LIII*(2), 499-547.

Andersen, T. G., & Bollerslev, T. (1998, November). Answering the Skeptics: Yes, Standard Volatility Models Do Provide Accurate Forecasts. *International Economic Review, 39*(4), 885-905.

Andersen, T. G., Bollerslev, T., Diebold, F. X., & Ebens, H. (2001). The distribution of realized stock return volatility. *Journal of Financial Economics, 61*(1), 43-76.

Bakshi, G., Cao, C., & Chen, Z. (1997, December). Empirical Performance of Alternative Option Pricing Models. *Journal of Finance, 52*(5), 2003-2049.

Berg, J., Nelson, F., & Rietz, T. (2007). Prediction Market Accuracy in the Long Run. *Working Draft*.

Black, F., & Scholes, M. (1973). The Pricing of Options and Corporate Liabilities. *Journal of Political Economy, 81*(3), 637-654.

Blair, B. J., Poon, S.-H., & Taylor, S. J. (2000, November). Forecasting S&P 100 Volatility: The Incremental Information Content of Implied Volatilities and High Frequency Index Returns. *Journal of Econometrics, 105*(1), 5-26.

Breeden, D. T., & Litzenberger, R. H. (1978). Prices of State-contingent Claims Implicit in Option Prices. *Journal of Business, 51*(4), 621-651.

Brigham, E. F., & Houston, J. F. (2008). *Fundamentals of financial management* (6 ed.). Mason: Cengage Learning.

Britten-Jones, M., & Neuberger, A. (2000, April). Option Prices, Implied Price Processes, and Stochastic Volatility. *The Journal of Finance, LV*(2), 839.

Canina, L., & Figlewski, S. (1993). The Informational Content of Implied Volatility. *The Review of Financial Studies, 6*(3), 659-681.

Chicago Board Options Exchange, Inc. (2011). *Micro Site - DJX*. Retrieved March 29, 2011, from CBOE: http://www.cboe.com/micro/djx/introduction.aspx

Chordia, T., Roll, R., & Subrahmanyam, A. (2001). Market Liquidity and Trading Activity. *Journal of Finance*, 20.

CME Group Index Services, LLC. (2010, July). *Dow Jones Industrial Average Historical Divisor Changes.* Retrieved March 29, 2011, from CME Group: http://www.cmegroup.com/trading/equity-index/files/djia-history-divisor.pdf

CME Group Index Services, LLC. (2011). *Calculating the Averages*. Retrieved March 29, 2011, from Dow Jones Averages: http://www.djaverages.com/?go=industrial-calculation

Dumas, B., Fleming, J., & Whaley, R. E. (1998, December). Implied Volatility Functions: Empirical Tests. *The Journal of Finance, LIII*(6), 2059-2106.

Duong, T. (2001, May 24). *An introduction to kernel density estimation*. Retrieved March 26, 2011, from http://school.maths.uwa.edu.au/~duongt/seminars/intro2kde/

Fama, E. F. (1970). Efficient Capital Markets: A Review of Theory and Empirical Work. *Journal of Finance, 25*(2), 383-417.

Figlewski, S. (1997, February). Forecasting Volatility. *Financial Markets, Institutions, and Instruments, 6*(1), 1-88.

Fleming, J. (1998). The quality of market volatility forecasts implied by S&P 100 index option prices. *Journal of Empirical Finance, 5*, 317-345.

Fornberg, B., & Zuev, J. (2006). The Runge phenomeno and spatially variable shape parameters in RBF interpolation. *Comput. Math. Appl*, 379-398.

Galai, D. (1978). Empirical tests of boundary conditions for CBOE options. *Journal of Financial Economics, 6*(2-3), 187-211.

Hafner, W., & Zimmermann, H. (2004). Amazing Discovery: Professor Bronzin's Option Pricing Models (1908). *Working Paper*, 1-66.

Hansen, C. S., Prabhala, N. R., & Christensen, B. J. (2001). The Telescoping Overlap Problem in Options Data. *AFA 2002 Atlanta Meetings.* Atlanta.

Hollywood Stock Exchange. (2011, February 17). *Hollywood Stock Exchange*. Retrieved February 17, 2011, from Hollywood Stock Exchange: www.hsx.com

Intrade. (2011, February 17). *Intrade Prediction Markets*. Retrieved February 17, 2011, from Intrade: www.intrade.com

Jackwerth, J. C. (2000). Recovering Risk Aversion from Option Prices and Realized Returns. *The Review of Financial Studies*, 433-451.

Jiang, G. J., & Tian, Y. S. (2005). The Model-Free Implied Volatility and Its Content. *The Review of Financial Studies, 18*(4).

Kou, S. (2008). Lévy Processes in Asset Pricing. *Wiley Online Library*.

Leigh, A., Wolfers, J., & Zitzewitz, E. (2003). What Do Financial Markets Think of War in Iraq? *NBER Working Paper*.

Mauldin, J. (2009, April 4). *Deep Inside the Dow*. Retrieved March 29, 2011, from Safe Haven: http://www.safehaven.com/article/13005/deep-inside-the-dow

Mayhew, S. (1995). Implied Volatility. *Financial Analysts Journal, 51*(4), 8-20.

McKinley, S., & Levine, M. (1998). *Cubic Spline Interpolation.* Retrieved March 31, 2011, from Student Projects in Linear Algebra: http://online.redwoods.cc.ca.us/instruct/darnold/laproj/Fall98/SkyMeg/Proj.PDF

Meyer, M. C. (2010, October 21). *R Code for Constrained Penalized Spline Estimation*. Retrieved April 25, 2011, from http://www.stat.colostate.edu/~meyer/penspl.htm

Nau, B. (2005, May 16). *Identifying the numbers of AR or MA terms*. Retrieved April 12, 2011, from Decision 411 Forecasting: http://www.duke.edu/~rnau/411arim3.htm#plots

Nau, B. (2005, May 16). *Stationarity and differencing*. Retrieved April 13, 2011, from Desicision 411 Forecasting: http://www.duke.edu/~rnau/411diff.htm

Nelson, F., Neumann, G., & Polgreen, P. (2006). Operating With Doctors: Results from the 2004 and 2005 Influenza Markets. *University of Iowa*.

NYSE Euronext. (2011). *NYSE Number of Trading Days 2010.* Retrieved March 18, 2011, from NYSE, New York Stock Exchange: http://www.nyse.com/pdfs/tradeday_10.pdf

Phillips, M. (2010, June 18). *Yep, It's the Dreaded 'Quadruple Witching'*. Retrieved March 29, 2011, from The Wall Street Journal: http://blogs.wsj.com/marketbeat/2010/06/18/yep-its-the-dreaded-quadruple-witching/

R Development Core Team. (2011). *Introduction to R*. Retrieved April 27, 2011, from R-project: http://cran.us.r-project.org/doc/manuals/R-intro.html#Examining-the-distribution-of-a-set-of-data

R Development Core Team. (2011). R: A Language and Environment for Statistical Computing. Vienna, Austria.

Rhode, P. W., & Strumpf, K. S. (2004). Historical Presidential Betting Markets. *Journal of Economic Perspectives, 18*(2), 127.

Rubinstein, M. (1994, July). Implied Binomial Trees. *Journal of Finance, LXIX*(3), 771-818.

Schneider-Joseph, D. (2004, July 13). *Idea Futures Exchanges*. Retrieved February 27, 2011, from David Schneider-Joseph: http://www.davidsj.com/post.php?id=103_0_1_0_C5

Shiller, R. J., & Perron, P. (1985). Testing the Random Walk Hypothesis. *Economics Letters, 18*, 381-386.

Shoven, J., & Sialm, C. (2000). The Dow Jones Average: The Impact of Fixing Its Flaws. *SIEPR Discussion Paper, 99-016*.

Smith, Z. R., & Wells, C. S. (2006). Central Limit Theorem and Sample Size. *Northeastern Educational Research Association*, (pp. 2-22). Kerhonkson.

Stoll, H. R. (1969). The Relationship between Put and Call Option Prices. *Journal of Finance, 24*(5), 801-824.

The Federal Reserve Board. (2011, February 21). *Selected Interest Rates.* Retrieved February 21, 2011, from FRB: Federal Reserve Statistical Release H.15: http://www.federalreserve.gov/releases/h15/data.htm

Weisstein, E. W. (2011, May 3). *Distribution Function*. Retrieved May 4, 2011, from From MathWorld--A Wolfram Web Resource: http://mathworld.wolfram.com/DistributionFunction.html

Wolfers, J., & Zitzewitz, E. (2004). Prediction Markets. *Journal of Economic Perspectives, 18*(2), 107-126.

Wolfers, J., & Zitzewitz, E. (2006). Interpreting Prediction Market Prices as Probabilities. *NBER Working Paper*.

Wolfers, J., & Zitzewitz, E. (2006). Prediction Markets in Theory and Practice. *NBER Working Paper*.

Yahoo! Inc. (2011). *^DJI Historical Prices*. Retrieved February 5, 2011, from Yahoo! Finance: http://finance.yahoo.com/q/hp?s=%5EDJI+Historical+Prices

Appendix 1 – DJIA component stocks

Ticker	Company Name
MMM	3M Company
AA	Alcoa Incorporated
AXP	American Express Company
T	AT&T Incorporation
BAC	Bank of America Corporation
BA	Boeing Company
CAT	Caterpillar Incorporated
CVX	Chevron Corporation
CSCO	Cisco Systems, Inc.
KO	Coca-Cola Company
DD	E.I. DuPont de Nemours & Company
XOM	Exxon Mobil Corporation
GE	General Electric Company
HPQ	Hewlett-Packard Company
HD	Home Depot Incorporated
INTC	Intel Corporation
IBM	International Business Machines Corporation
JNJ	Johnson & Johnson
JPM	JPMorgan Chase & Company
KFT	Kraft Foods Incorporated Cl A
MCD	McDonald's Corporation
MRK	Merck & Co. Incorporated
MSFT	Microsoft Corporation
PFE	Pfizer Incorporated
PG	Procter & Gamble Company
TRV	Travelers Companies, Inc.
UTX	United Technologies Corporation
VZ	Verizon Communications Incorporated
WMT	Wal-Mart Stores Incorporated
DIS	Walt Disney Company