The forecasting ability of implicit risk-neutral density functions
- A study of planned economic events in Sweden

Johan Land

MASTER'S THESIS IN FINANCE

STOCKHOLM SCHOOL OF ECONOMICS

Abstract

Financial institutions spend large amounts of money on gaining accurate information. The information they acquire is often kept for themselves and used in order to trade and make profits. When they use the information it is though implicitly put into the prices of the instruments. This because they price the instruments on the information they have. In this thesis option prices are used in order to derive the underlying information that the financial institutions used to price them. Explicitly, the implicit Risk-Neutral Density (RND) is derived from option prices close to planned economic events, and the shape of this function is studied and interpreted. In practise, the Swedish election of 2006 and Swedish Central Bank meetings are studied in order to forecast the outcomes of these events. The result is that the RND changed drastically around the election, but the implications of these changes are uncertain. For the Central bank meetings the result was that it was not possible to forecast the outcomes of any of these meetings.

Keywords: Implicit Volatility, Risk-neutral density, Forecasting

Tutor: Tomas Björk
Presentation: June 12th 2007, 10:15 a.m.
Venue: Room 750
Discussants: Victor Jacobsson and Sebastian Siemiatkowski

Acknowledgements: I would like to thank my tutor for efficient, precise and accessible advice.

Table of contents

1 Introduction

The market in general is affected by many events with uncertain outcome. Some of these events are unexpected, such as earthquakes and wars, while other events are planned but still with uncertain outcome, such as central bank interest meetings or political elections. If we focus on the planned events they are often of the type where there are two or more possible outcomes, often with significantly different economic meaning. In the case of a central bank meeting there are organizations doing different polls among investors in order to forecast the expected decision of the central bank, and in the case of a political election there are different election polls among the citizens also trying to forecast the outcome.

A price on an instrument is formed based on the market's view on the instrument. A future is priced based on the market's view of the price development of the underlying instrument in the future. Hence, futures can tell us about the market's expectation about the future price of the underlying instrument. Options, take this reasoning one step further and carries even more information as there are many options traded on every underlying instrument, both options with different strike prices and options with different strike dates. This way options can tell us much more about the future development of the underlying instrument.

What if we could use the option prices and the information carried in the option prices for extracting the implicit market view on the price developments, and what if we would do this close to a planned economic event with uncertain outcome? Could we then from the option prices derive the implicit market forecast on this economic event and its outcome? More precisely:

Can we calculate which outcome of planned economic events such as elections and interest changes that the market believes to be most probable by only using option prices?

2 Methodology

Option prices are in some sense formed based on investors beliefs about the future. Close to economic events like central bank interest meetings or elections the market forms a view about the outcome of the event and the options are priced based on this view. Hence the information about the market's view on the event is implicitly in the option prices.

The methodology to extract this information chosen in this thesis is to derive something called the Risk-Neutral density function or the RND-function and then interpret the RND-function around economic events. Simple share prices or futures carry much less information about the expectation about the price in the future than options that carry much more information as there are many options for each underlying instrument. The information carried in option prices is best shown by deriving the RND-function.

The RND-function can be derived in many ways and we will use two ways in this thesis. One direct method which will give us a discrete RND-function using elementary claims and butterfly spreads (Bahra 1997). The other method is to fit a mix of two log-normal densities to the actual option prices by a parametric approach and minimizing the deviances. The latter method has been used in other studies (Gemmill and Apostolos 1999) to study economic events as this method very well captures economic events with two possible outcomes. If the aim was to study how certain attributes of the densities, such as skewness, variance etcetera varied over time, a more direct method of deriving the RND would have been better. Following the development of theory around deriving the RND-functions Jackwerth (1999) developed a method for deriving Implied Binomial Trees. This is the corresponding model to the binomial trees used for pricing options. Jackwerths Implied Binomial Trees does not only study the terminal RND-function but more the implied process of the underlying asset price changes.

One problem with the simple methods of deriving the RND-function is that they can give negative probabilities. The Gram-Charlier series expansion method suggested by Rompolis and Tzavalis (2006) is a method that guarantees positive probabilities. The method does however not offer any advantages for studying planned economic events as the two-log-normal method does. Also, Keller and Craig (2005) studied more complex methods and did not find any signs of more complex methods having a better forecasting ability.

Äijö (2006) found that events had larger impacts on the implicit RNDs that had a short maturity than on RNDs derived from options with a long maturity. Therefore, options with an as short maturity as possible have been used throughout this thesis. The reason for events having a larger impact on short maturities is mainly because we will have less noise and less economic factors to take into account. Javiera (1999) confirmed that RNDs derived from options with a short maturity had better forecasting abilities for options on Swedish foreign exchange.

2.1 Why not a questionnaire?

Can we calculate which outcome of planned economic events such as elections and interest rate changes that the market believes to be most probable by only using option prices?

To come back to the original question posed in this thesis ”Can we calculate which outcome of planned economic events such as elections and interest changes that the market believes to be most probable by only using option prices?” one might question whether the method of using option prices is useful at all. Should we not instead just ask investors about what they think about the economic event? Do they think that the central bank will raise the interest rate or not?

The method of using a questionnaire for answering these questions is for sure much more clear in its interpretation and it is certainly easier to explain and it is therefore likely to be considered as more legitimate and more accurate. One does still though need to remember that a questionnaire has certain flaws. Generally the step from the raw data to the conclusion is fairly straight forward, but the raw data and answers from the questionnaires can be strongly questioned. Was the respondent group representative? Did all respondents answer? Were the questions clear? Could the respondents have an incentive of not responding honestly?

In contrast, the raw data using a method of deriving the RND-function is very accurate and reliable. The option prices have been formulated by investors trading real money and they have very large sums at stake. Therefore, the raw data is very reliable and it clearly reflects actual respondent sentiment. With the RND-function method the data is not the questionable part, but the method of calculating the RND-function and even more questionable is the interpretation of the RND-function.

These two totally different methods of the RND-function and the questionnaire posses totally different strengths and I would argue that the two methods complement each other.

3 Data

The data used is end-of-the-day data for both call and put options on the Swedish index OMXS30 for every trading day of 2006. OMXS30 is the largest index in Sweden and the options on this index are the most traded among the options on the Swedish stock market. The data contains end-of-the-day bid and ask prices, last price, highest price, lowest price and volume. The corresponding futures to the options and the index value itself are also contained in the data.

The reason for choosing the OMXS30 is to make sure we get the most traded option data under the restriction of only using data from the Swedish market. But, the OMXS30 is also an index over the 30 largest companies in Sweden. Hence, it very well reflects the overall market sentiment as opposed to if we would have chosen to study the option data of a single large Swedish company. This way we also limit the effect of single company events that otherwise possibly could have affected our results and findings.

The method used in this thesis is highly dependent on the data. For each strike date of the options the data needs to contain as many strike prices as possible as each strike price contains additional information about the market's view on the future and is crucial when deriving the RND-function. But, also each option needs to be accurately priced and for that we need a high volume and a low bid-ask spread.

The option data for 2006 is fairly good for options with strike dates one or two months away. Luckily these are exactly the options that are most useful to us as we only study the short term view of the market around specific events. The tail-options with strike prices far from the expected price do though have a very low volume and often a very large bid-ask spread. Sometimes there might not even be both a bid and an ask.

To cope with the unreliable tail-data some data points have been excluded based on the following criteria:

- Only bid or ask is available
- The last price is clearly off the chart. Examples are call options with a low strike price that have a lower price than options with a higher strike price.

For some options that have a high volume the last price figure can be outside the bid-ask spread. For these options the last price has been changed to the mean of the bid and ask price. For some options the volume was zero and no trades had taken place, but the bid-ask spread was small. For these options the last prices were also replaced by the mean of the bid and ask prices. For all other options the last price has been considered as the actual price of the option, and from now on when we refer to the price of the option we mean this adjusted last price figure.

4 Theory

4.1 Risk-neutrality

Let S_T be a stochastic variable denoting the value of a certain instrument at time T in the future.

The expected value of S_T at time T can be calculated as:

$$E(S_T)$$

What if we wanted to know how much investors today would be willing to pay to get one unit of the asset at the time T. As investors do not know the price of the instrument at time T they have to base their valuation today on some kind of expectations about the value of the instrument at time T. But, investors are risk-averse and need some compensation for the risk and uncertainty, therefore they will not simply pay the discounted expected value.

In order to cope with this risk-aversion a new expected value-measure is defined. The new measure is called the *risk-neutral measure* and is denoted $E^Q()$. This measure is such that the expected value of a stochastic variable S_T takes the risk-aversion into account and hence denotes the safe equivalent to the stochastic pay-off S_T. That is, the price today of getting one unit of an asset of stochastic price S_T will be:

$$P_0 = e^{-rT} E^Q(S_T)$$

Where r is the risk-free rate (Jackwerth 99).

In order to describe the function $E^Q()$ its density function can be derived. The density function of this risk-adjusted measure is called the Risk-Neutral Density function, or simply the RND-function.

RND-functions are very handy when valuing derivatives. E.g. if f(x) is the RND-function of S we can simply calculate the price today of getting an instrument at time T as:

$$P_0 = e^{-rT} E^Q(S_T) = e^{-rT} \int_0^{\infty} x f(x) dx$$

Call-options have the payout max(S-K,0), and the pricing of options given an RND-function is also greatly simplified under the risk-neutral measure as:

$$c(K) = e^{-rT} E^{Q}\left(\max\left(S_T - K, 0\right)\right) = e^{-rT} \int_{0}^{\infty} \max\left(x - K, 0\right) f(x) dx = e^{-rT} \int_{K}^{\infty} \left(x - K\right) f(x) dx$$

In the methods for deriving f(x) that we will present later f(x) will be the Risk-Neutral Density function. In the next section we will show that the common assumption of f(x) being log-normal usually is not true.

4.2 Implied Volatility

For a given day with a given price of the underlying asset and with a given risk-free rate the price of options with a given time to maturity can be expressed as a function of the following[1]:

$$c = c\left(K, \sigma\right)$$

That is, the price of an option depends only on the strike price and the volatility of the underlying asset. Observe here that we have assumed the prices being log-normally distributed. If Assuming we know the observed market prices c*() we can solve the equation for the volatility σ and we end-up with:

$$\sigma = \sigma\left(c^{*}, K\right)$$

The value of sigma should of course be the same no matter what k and c we choose among the many options as the nature of the underlying asset doesn't depend on our choice of option. But, if we plot $\sigma(c^{*},K)$ and c(K, σ) we get the following typical shapes shown in Figure 1.

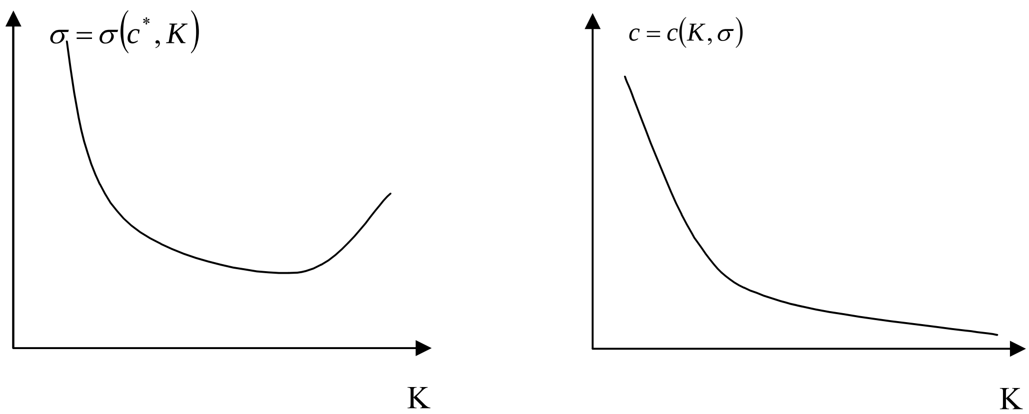

[1] Using e.g. Black and Scholes

Figure 1: A constructed smiley in the left graph. Typical option prices in the graph to the right.

σ appears to be different for different strike prices k. This is not in line with the assumption about the prices being log-normally distributed and we have to abandon that assumption. Instead we will derive the RND-function without assuming it being log-normal. This is described further in the next section.

4.3 Deriving the RND-function

We will now develop the theory for calculating the RND-function using two different methods. The first method called the Elementary claims method is a direct method using the option prices. This method is more used as a reference method as it is too simply to be useful in answering our question, it does though give us the basic knowledge needed to know about the RND-function. The second method called the two-log-normal method uses a parametric form of the RND-function, and determines the parameters by minimizing an error function. This method is better in answering our question about planned economic events as it very well captures different possible outcomes.

4.3.1 Elementary claims method

Let c(K) be the price of a call-option with strike price K. The formula for calculating this price is:

$$c(K) = e^{-rT}\int_K^{\infty}(x-K)f(x)dx$$

Where f(x) is the risk-neutral density function. If we take the second derivative of this expression we end up with:

$$c'(K) = e^{-rT}\cdot 0 - e^{-rT}\int_K^{\infty} f(x)dx = -e^{-rT}\int_K^{\infty} f(x)dx$$
$$c''(K) = -e^{-rT} f(K)$$
$$f(K) = -e^{rT}c''(K)$$

Thus, we can calculate the RND-function if we know the second derivative of the prices of the options (Jackwerth 99). However, we only have option prices for a discrete set of strike prices. In order to cope with this we use the discrete version of the second derivative and end up with the following formula:

$$f(K) = -e^{rT} c''(K) = e^{rT} \cdot \frac{c(K+\Delta K) - 2 \cdot c(K) + c(K-\Delta K)}{(\Delta K)^2}$$

In practice option prices are often in discrete strike price intervals and this formula is therefore fairly easy to use. And as we calculate all possible f(k) we have a discrete version of the RND-function.

4.3.2 Two-log-normal method

A different approach to deriving the RND-function is to first define the RND-function using a certain setting of parameters, then determining these parameters so that the RND-function well reflects the given market option prices. It seems appropriate to let assumed RND-function have certain properties, such as being equally distributed over addition[2] and still having some resemblances with the previous log-normal density such as finite variance in the tails. We will here choose to define the new density function as the weighted sum of two parametric log-normal functions. Keller and Craig (2005) showed that more complex choices of density functions do not give better results and we therefore choose not to go any further in the choice of the density function. The sum of two log-normal functions has only five parameters which is good when the data is limited. Bahra (1997) did a test of five different approaches of calculating the RND-function and he found the two-log-normal function to be the “preferred” method.

The five parameters are calculated by defining a function which measures the error that the option prices generated from this function has compared to the actual option prices. This error-function is minimized over the parameters and the parameters with the lowest error-function value are the parameters that we consider giving us the best density function. Gemmill and Apostolos (1999) showed that this two-log-normal method performs better than the Black and Scholes one-log-normal model.

4.3.2.1 Derivation of minimization function

A variable X is log-normal if it can be written:

---

[2] The log-normal distribution is such that if “daily prices are lognormally distributed then other arbitrary length holding period price distributions must also be lognormal” (Bahra 1999). This is an important attribute as we regard the underlying reasons for price changes as random on not dependent on the time period chosen, hence the density function should not neither be dependent on the time period.

$X = e^Z$

Where Z is normal. Hence, ln(X) is normal. Figure 2 shows a typical log-normal density.

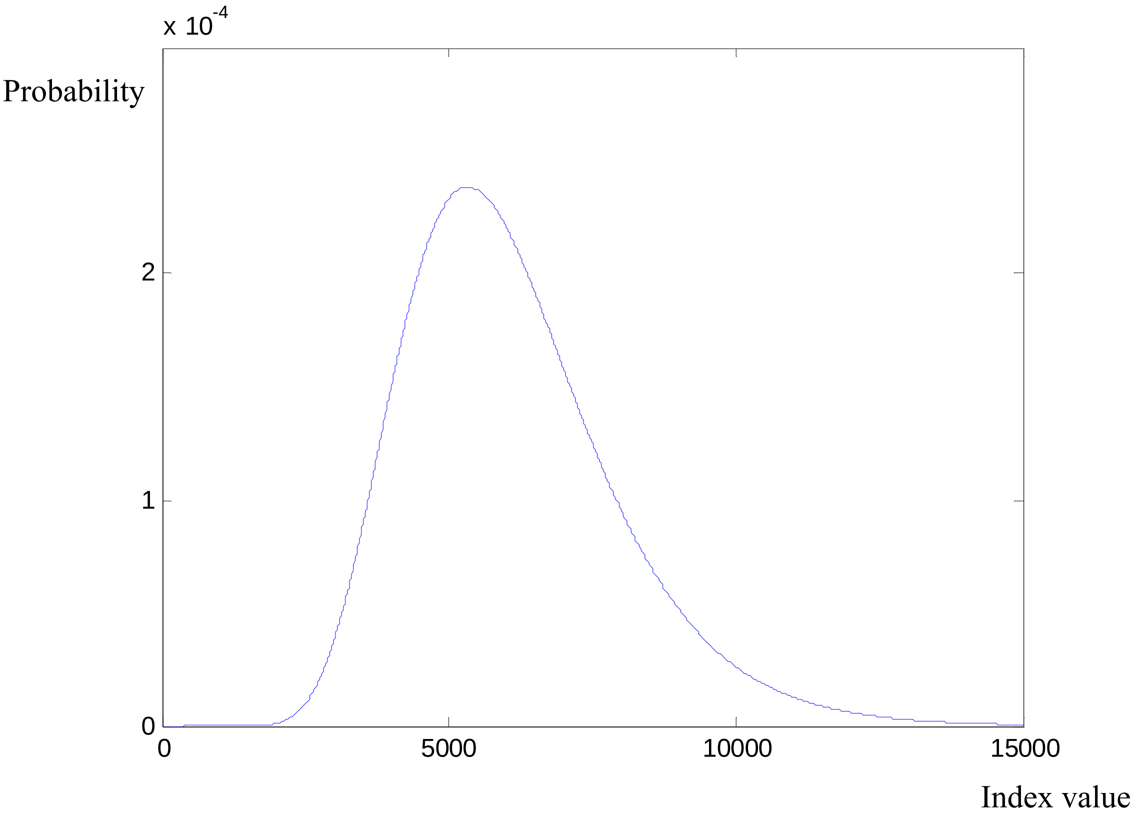

Figure 2: Principle look of a log-normal density

The log-normal density function can be expressed as:

$$f(x,\mu,\sigma) = \frac{e^{-(\ln x-\mu)^2/2\sigma^2}}{x\sigma\sqrt{2\pi}}$$

If we define our function as a weighted sum of two independent log-normal functions we get the following function:

$$\varphi(x,\alpha,\mu_1,\sigma_1,\mu_2,\sigma_2) = \alpha \cdot \frac{e^{-(\ln x-\mu_1)^2/2\sigma_1{}^2}}{x\sigma_1\sqrt{2\pi}} + (1-\alpha)\cdot \frac{e^{-(\ln x-\mu_2)^2/2\sigma_2{}^2}}{x\sigma_2\sqrt{2\pi}}$$

Our aim now is to calculate the five parameters that best fit φ() to the given option prices. Assuming that φ () describes the actual probability density at strike date we can calculate the theoretical price of put and call options. For a call option with strike price K we have the following theoretical price for that option:

$$c(K) = e^{-r\tau} \int_K^{\infty} \varphi(s,...)(s-K)ds$$

And for put options we have:

$$p(K) = e^{-r\tau} \int_0^{K} \varphi(s,...)(K-s)ds$$

Observe here that the formula simply is the present value of the sum of all possible pay-outs weighted by their risk-neutral probabilities.

To determine our five parameters we need to define certain criteria that φ () needs to fulfil. First we want all theoretical prices on both call options and put options to be as close to the given option prices as possible. But from φ () we can also calculate the expected value of the underlying asset. The present value of this expected value should theoretically be equal to the current value of the underlying asset.

Let c_i and p_i be the given option prices with corresponding strike prices K_i. The error-value for a call or put option is then:

$$e_{c,i} = c(K_i) - c_i$$
$$e_{p,i} = p(K_i) - p_i$$

We showed earlier that the price today of an asset can be calculated as:

$$P_0 = e^{-rT} E^Q(S_T)$$

The expected value of a variable X that is log-normal is:

$$E(X) = e^{\mu + \sigma^2/2}$$

If S_T has the distribution φ() we get:

$$P_0 = e^{-rT} E^Q(S_T) = e^{-rT}\left(\alpha \cdot e^{\mu_1 + \sigma_1^2/2} + (1-\alpha)\cdot e^{\mu_2 + \sigma_2^2/2}\right)$$

But the price today can be observed, that is we know P_0^*. But P_0^* should be equal to P_0 and we have the following difference:

$$e_{\text{expected}} = e^{-r\tau}\left(\alpha \cdot e^{\mu_1 + \sigma_1^2/2} + (1-\alpha)\cdot e^{\mu_2 + \sigma_2^2/2}\right) - P_0^*$$

To determine our parameters we define a sum of all errors square and minimize this function in order to determine the parameters that best fit our assumed density function to the given parameters. We define the minimization function as follows:

$$E(\alpha,\mu_1,\sigma_1,\mu_2,\sigma_2)=\sum_i e_{c,i}^2+\sum_i e_{p,i}^2+e_{\text{expected}}^2=\sum_i\left(c(K_i)-c_i\right)^2+\sum_i\left(P(K_i)-P_i\right)^2+e_{\text{expected}}^2$$

$$=\sum_i\left(e^{-r\tau}\int_{K_i}^{\infty}\varphi(s,\ldots)(s-K_I)ds-c_i\right)^2+\sum_i\left(e^{-r\tau}\int_0^{K_i}\varphi(s,\ldots)(K_i-s)ds-p_i\right)^2+$$

$$\left(e^{-r\tau}\left(\alpha\cdot e^{\mu_1+\sigma_1{}^2/2}+(1-\alpha)\cdot e^{\mu_2+\sigma_2{}^2/2}\right)-P_0^*\right)^2$$

5 Previous research

Many early studies are focused on determining whether RND-functions outperformed other historically based methods for forecasting the future such as GARCH. Most of these studies have found that the RND-method greatly outperforms historical methods. Examples of studies on this is Gemmill and Apostolos (1999) and Wilkens and Röder (2006).

The main study that is referred to in most articles on implicit volatilities and RND functions is the study by Bhupinder Bahra (1997). In this study Bahra lays the ground for five different methods of calculating the implicit RND-functions and he roughly tests them on empirical data. He concludes that the best method is the parametric method of a weighted sum of two independent log-normal density functions.

Most studies are done on either options on interest rates or options on equities or equity indexes, but there also exists studies on options on foreign exchanges. Also, most studies are done on the main markets of the world, but there are also studies on smaller exchanges. One example is the study by Arild Syrdal (2002) on the Norwegian option market where he finds the Norwegian market to be too illiquid to derive RND-functions that make any sense. There is also a study on the Swedish foreign exchange market by Javiera (1999) where he concludes that options on Swedish exchange rates outperform other historically based forecasting methods such as GARCH, but RNDs on Swedish exchange rates are far worse than RNDs on international exchange rates. A Finnish study by Nikkinen (2003) finds that RNDs are usable on the Finnish market, but that there exists large differences compared to larger markets.

Bahra (1997) originally highlighted the occurrences of instable solutions with some models including the two-log-normal method and he concluded that this was due to bad data. Syrdal (2002) confirms this by concluding that the data on Norwegian options is too illiquid to at all find any stable solutions.

There are some studies on specific events but most studies have focused on shifts in the RND-functions around specific planned and non-planned events whereas others have studied the shapes of the RND-functions in detail around events. Gemmill and Apostolos (1999) studied UK-elections and found RND-functions with two maximums. They could not though relate the two maximums to any political outcomes and it did not correlate with election polls. Mandler (2002) studied European Central Bank meetings and the RND-functions around these meetings. He even developed a new method for doing this using mean values over many days. Still he could not find any evidence of the central bank meetings in the RND-functions.

His conclusion was that there is too much noise and that the meetings did not have enough economic effect on the option prices and RND-functions.

6 Analysis

6.1 Elementary claims method

6.1.1 Test on a random day – January 2nd 2006

In the theory section a formula for calculating the implicit RND function from option prices via an elementary claim was shown. This method had option prices for equally spaced strike prices as input data. As a start we will calculate the RND-function for just a random day, say the first trading day of the time period, that is January 2nd 2006. We choose options with strike date of January 31st.

After adjustments of the data as described in the data section, the call option prices shown in Table 1 were observed.

Strike prices	Option prices
930	35
940	27
950	19,75
960	13,25
970	8
980	4,75
990	2,6
1000	1,1

Table 1: Call option prices January 2nd 2006 with maturity January 31st

The corresponding put option prices are shown in Table 2.

Strike prices	Option prices
880	0,4
890	0,55
900	1
910	1,1
920	1,5
930	2,25
940	3,6

950	5,5
960	8,75
970	14
980	20,125
990	26,75

Table 2: Put option prices January 2nd 2006 with maturity January 31st

It is interesting to note that we have more data for put options than for call options. This is generally the case which is explained by put options on OMXS30 being more traded than call options.

First, all put-options are recalculated as call-options using the put-call parity. In case of overlapping values the average price is taken. Secondly, the formula described in the theory section is used and we simply calculate the discrete risk-neutral probability density. As we are only interested in the shape of the density we choose to normalize the probabilities over the available data. The resulting RND-function is shown in Figure 3.

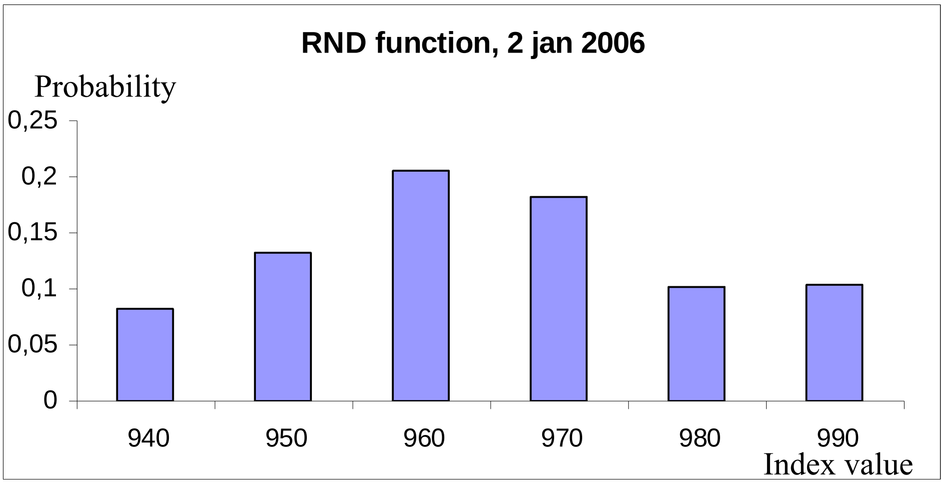

Figure 3: RND function as of January 2nd 2006 calculated on options with maturity January 31st 2006

From the figure above we can note that the RND function resembles what we would have expected – a graph centred on a specific strike price with tails diminishing as we move away from the centre. Worth noting is that we seem to have only one mode centred on 960. Futures traded at 964,5 at January 2nd 2006 so the mode of the graph is in line with the future price

The RND-function for the non-overlapping values in the tail is shown in Figure 4.

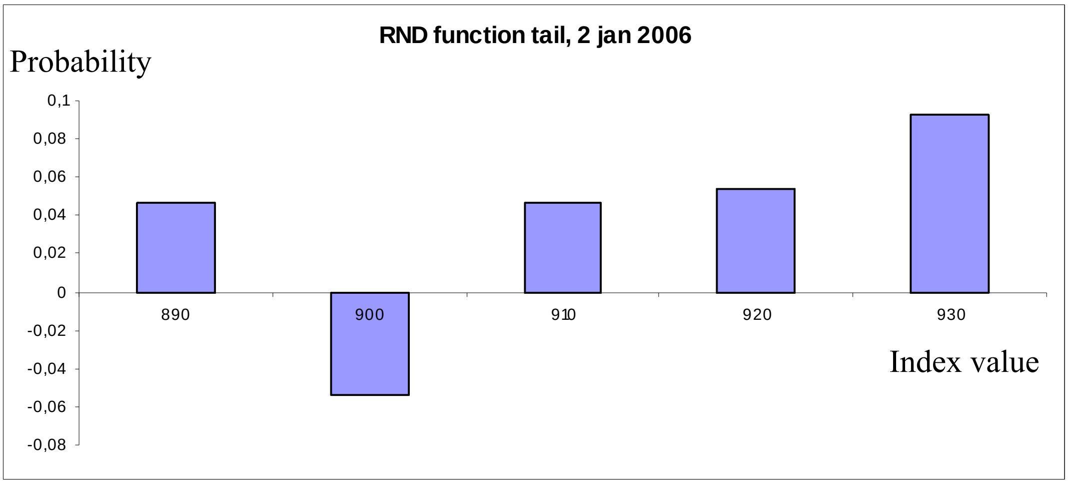

Figure 4: Tail of the RND function as of January 2nd 2006 calculated on options with maturity January 31st 2006

One thing worth noting about the graph of the RND-tail is that we have a negative probability for index value 900. If we go back to the formula in the theory this means:

$$p_i = e^{rT} \cdot \frac{c(910) - 2 \cdot c(900) + c(890)}{\Delta S} < 0$$

$$c(910) - 2 \cdot c(900) + c(890) < 0$$

$$c(910) + c(890) < 2 \cdot c(900)$$

What it means is that we can sell two options with strike price 900 and buy an option with strike 910 and one option with strike 890 and have a positive amount of money left and a positive pay-back. The conclusion is that we get money today and have a certain positive pay-back in the future. This is an arbitrage. So, in theory negative probabilities is the same thing as an arbitrage. Rompolis and Tzavalis (2006) presented a method for calculating RND-functions that assured only positive probabilities. As the negative probabilities only occur in the tails it does not really affect our results and we will stick to the method of using elementary claims as suggested by Bahra (1997).

In practise though these negative probabilities only occur in the tails of the probability densities, and in practise there is no arbitrage as in practise it is impossible to buy and sell options at the same price and when we calculate "the price" of the option we do it as the last traded price, but it might not be possible to actually both buy and sell options at this price. Therefore, the elementary method seems to work well for the central part of the density but it seems to be highly unreliable for the tails of the density.

6.1.2 Test the day before the election

In order to answer the original question of this thesis we should examine the implicit RND functions before and after certain planned economic events. Sweden's probably most important planned economic event during 2006 was the election in September. Figure 5 shows the implicit RND-function calculated using the elementary claims method Friday September 15th – the trading day before the election.

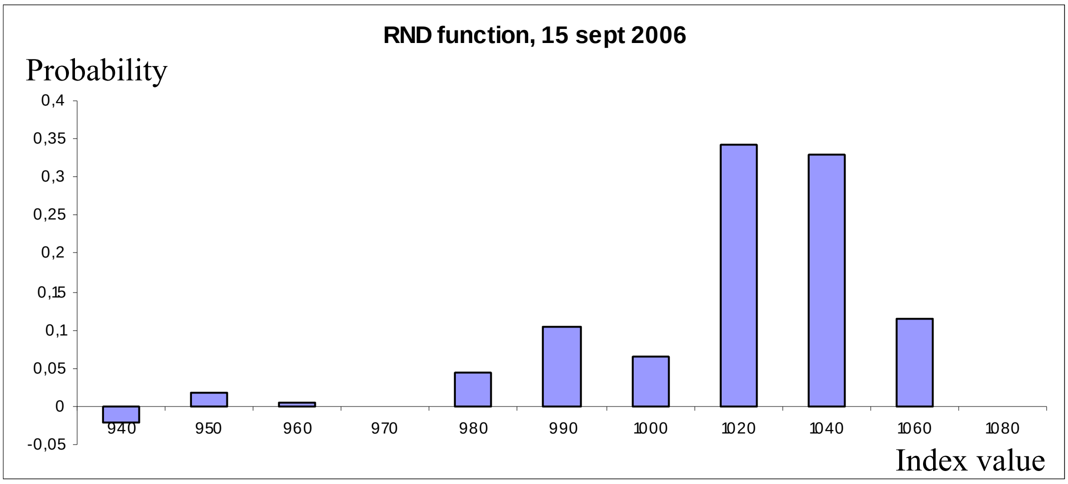

Figure 5: RND function as of September 15th 2006 calculated on options with maturity September 29th 2006

If the market would put different economic meaning to the two possible outcomes of the election, and if the market believed that these two outcomes were roughly equally probable, we should expect to see two maximums of the function – one for each outcome. This would of course require the two outcomes to have sufficiently different economic meanings, and it would also require none of the outcomes to be dominantly probable over the other.

In the graph above we can though, with a great deal of goodwill, see a large maximum around 1030 and a smaller maximum at 990 somewhat to the left. This is shown in Figure 6.

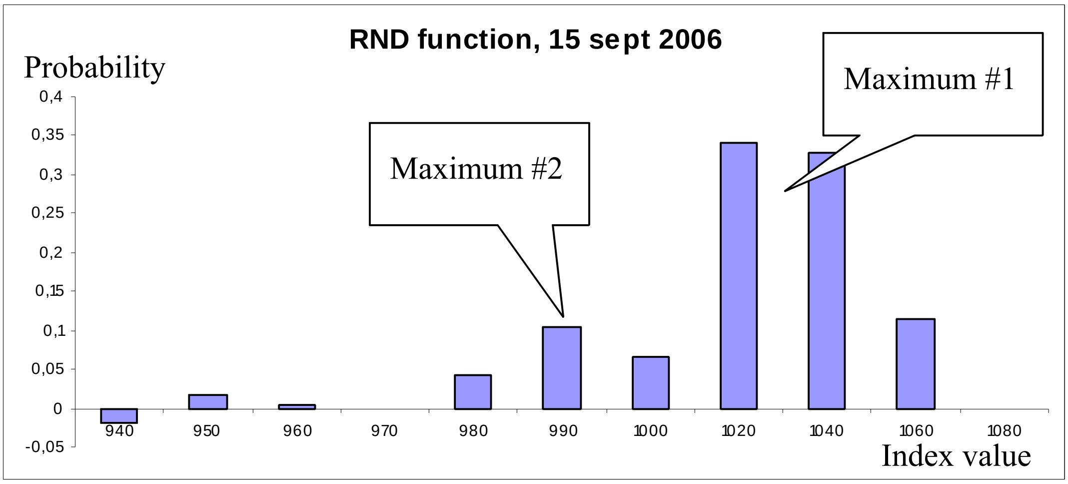

Figure 6: RND function as of September 15th 2006 calculated on options with maturity September 29th 2006 with pointers at two possible maximums

This argument is however very uncertain and requires a great deal of positive interpretation and it certainly could be criticized and claimed to be totally unreliable due to the graphs varying so much and that we have negative probabilities in some cases. I will not take it any further, but only conclude that the method of elementary claims does not fulfil our needs as it requires a lot of human interpretation which is easily criticized.

Instead, the interpretation will be left to the other more mathematical approach described in the theory section as the two-log-normal method.

6.2 Two-log-normal method

Let us now move on to the approach of using two log-normal functions to approximate the density function. In order to use this method we need to undertake some further theory to make the method work in practice and for our data and purpose. After that we will have a look at some empirical examples.

6.2.1 Further theory to apply the two-log-normal method

6.2.1.1 Discretization of the minimization function

The minimization function from the theory section contained among other things two expressions for call and option values. These expressions had integrals in them over very large intervals. In practice we can only handle these integrals numerically and hence we need to find discrete versions of them. Below is the expression from the theory section:

$$E(\alpha,\mu_1,\sigma_1,\mu_2,\sigma_2)=\sum_i\left(e^{-r\tau}\int_S^{\infty}L(s,\ldots)(s-K_I)ds-c_i\right)^2+\sum_i\left(e^{-r\tau}\int_0^{S}L(s,\ldots)(K_i-s)ds-p_i\right)^2+$$
$$\left(e^{-r\tau}\left(\alpha\cdot e^{\mu_1+\sigma_1^2/2}+(1-\alpha)\cdot e^{\mu_2+\sigma_2^2/2}\right)-I\right)^2$$

In order to calculate the integrals we choose the trapezium rule where the function that we are integrating over is replaced by a linear approximation in many discrete intervals. In practice this is done by using the built-in method trapz in Matlab. Trapz reforms the formula as follows[3]:

$$E(\alpha,\mu_1,\sigma_1,\mu_2,\sigma_2)=$$
$$\sum_i\left(e^{-r\tau}\frac{S}{n}\left(\frac{L(S,\ldots)(S-K_I)+L(2S,\ldots)(2S-K_I)}{2}+\sum_{k=1}^{n-1}L\left(S+k\frac{S}{n},\ldots\right)\left(S+k\frac{S}{n}-K_I\right)\right)-c_i\right)^2$$
$$+\sum_i\left(e^{-r\tau}\frac{S}{2n}\left(\frac{L(S/2,\ldots)(K_I-S/2)+L(S,\ldots)(K_I-S)}{2}+\sum_{k=1}^{n-1}L\left(S/2+k\frac{S}{2n},\ldots\right)\left(K_I-S/2+k\frac{S}{sn}\right)\right)-p_i\right)^2+$$
$$\left(e^{-r\tau}\left(\alpha\cdot e^{\mu_1+\sigma_1^2/2}+(1-\alpha)\cdot e^{\mu_2+\sigma_2^2/2}\right)-P_0^*\right)^2$$

6.2.1.2 Algorithm to minimize the function

Most minimization methods are based on some start values from which small changes are made in each dimension and new values are calculated, then the smallest value is chosen as the new start value and the method is iterated. The theory around minimization and optimization is a science of its own and I have chosen not to include any further discussion around it in this thesis.

The minimization method is the standard method in Matlab called fminsearch. It is using a simplex search method[4] and is not dependent on any gradients which is good for our discrete and empirical data.

---

[3] We here replace infinity by the double start value and zero with half the end value. In practice this is good enough as φ() is very small for these values for strike dates of only a few months.

[4] The method is described futher in Lagarias, Reeds, Wright, SIAM Journal of Optimization, 1998

6.2.1.3 Start values

The minimization algorithm requires some start values as input and we need to supply initial α, μ_1, σ_1, μ_2 and σ_2. The only requirement on the start values is that they need to give an initial density function that resembles the solution that we expect to get.

It makes sense to set α=0.5 as this gives equal weight to the two log-normal functions. As for σ_1 and σ_2 they are chosen to give a density function of normal variance. It proves that a value of 0.01 gives a normal variance.

The solution will be centred on the futures price with same strike date and this value is chosen for μ_1. We can though not choose $\mu_2 = \mu_1$ as this can make the minimization algorithm behave improperly as the two log-normal functions are equal. Therefore we choose $\mu_2 = \mu_1 - \varepsilon$ where ε is a very small value. This way there is a small difference in the two log-normal functions.

6.2.2 Test on a random day – January 2nd 2006

As with the elementary claims method we will test the two-log-normal method on just any random day in order to get a grasp of what kind of results we can expect. We choose to examine the same day as with the elementary claims method – the day January 2006 2nd – with the same strike date at end of January.

Futures were trading at 964 and the start values were chosen around this value. The actual start density shown in Figure 7 is hence the weighted sum of these two graphs with weight 0.5.

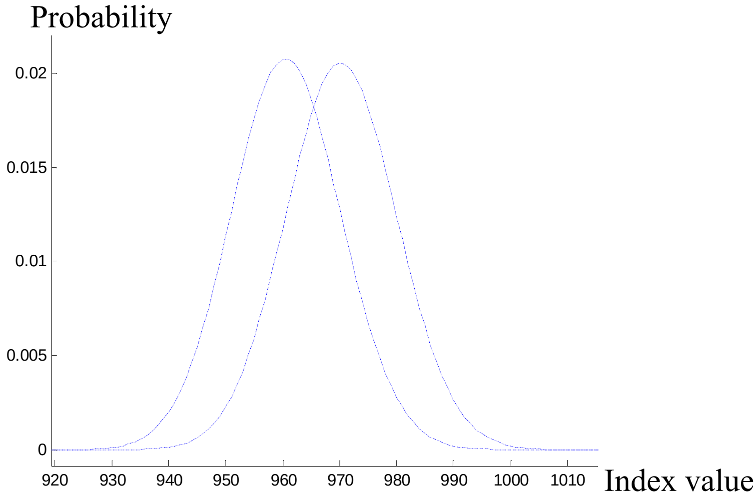

Figure 7: Start values for the optimization algorithm shown as the corresponding two log-normal functions

The solution we get after using fminsearch with start values as above and with the discrete formula developed in the section before is shown as non-dashed in Figure 8.

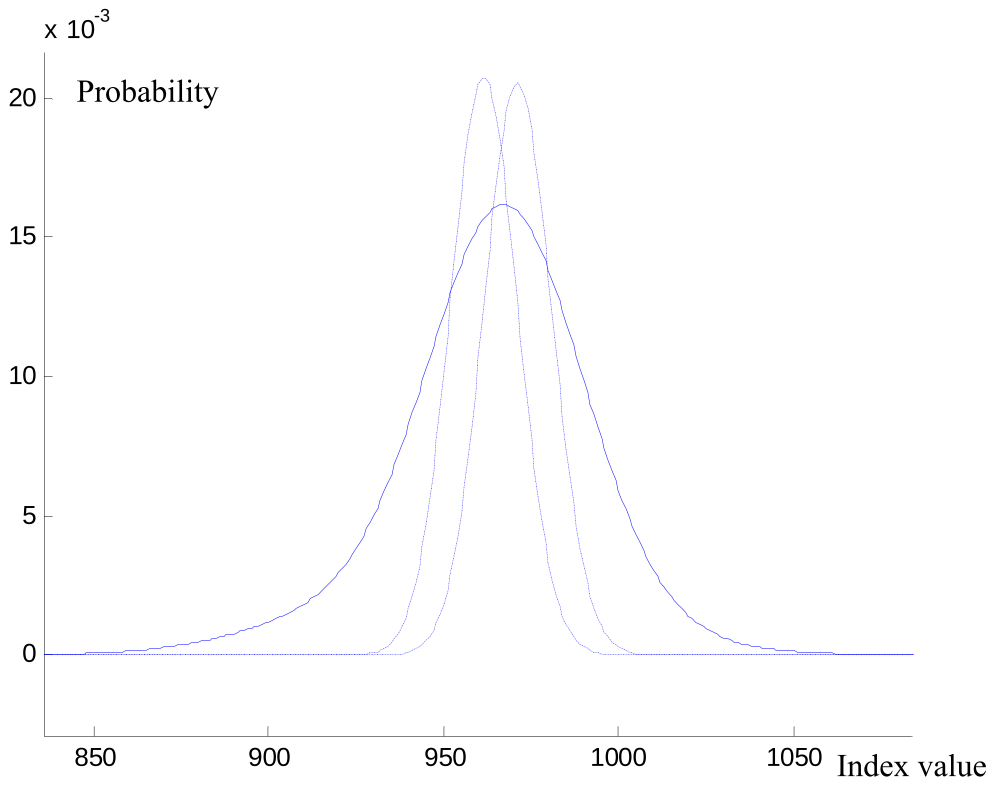

Figure 8: RND function as of January 2nd 2006 calculated on options with maturity January 31st 2006. Start-values are shown in dashed.

Hence we see that the main change from the start values is an increase in the volatility. We can also conclude that we seem to have only one maximum which is in line with what we found using the elementary claims method. As expected we do also see that the tails are handled much better by this method than by the elementary claims method.

It could be interesting to examine what the two underlying log-normal functions look like and what end-parameters we actually end up with. Figure 9 shows the build-up of the two-log-normal solution RND-function.

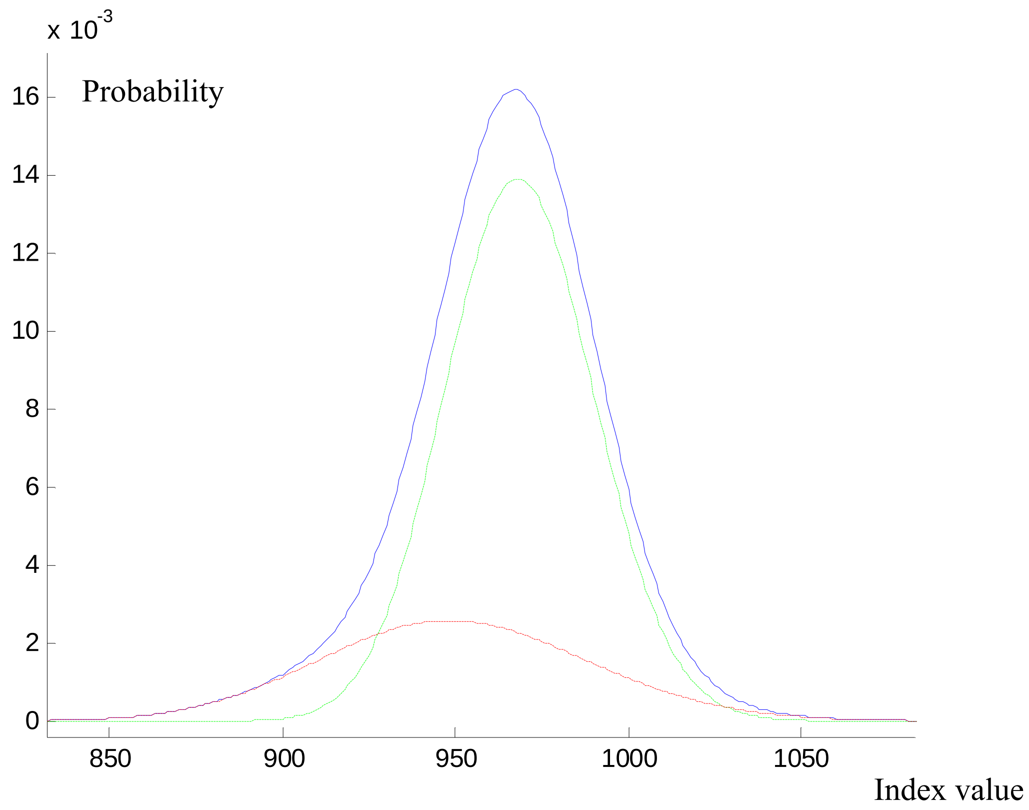

Figure 9: RND function as of January 2nd 2006 calculated on options with maturity January 31st 2006. The two underlying log-normal functions are shown as dashed.

It seems like one of the log-normal functions is more dominant and the second log-normal function seems to add additional variance. Hence, we end up with an altered log-normal function with thicker tails.

For some days the start values can be misleading and we might end-up with odd solutions. These solutions generally converge to stable solutions with slight adjustments of the start values. Bahra (1997) also found these odd solutions and explained them with poor data. One odd solution is presented in Figure 10.

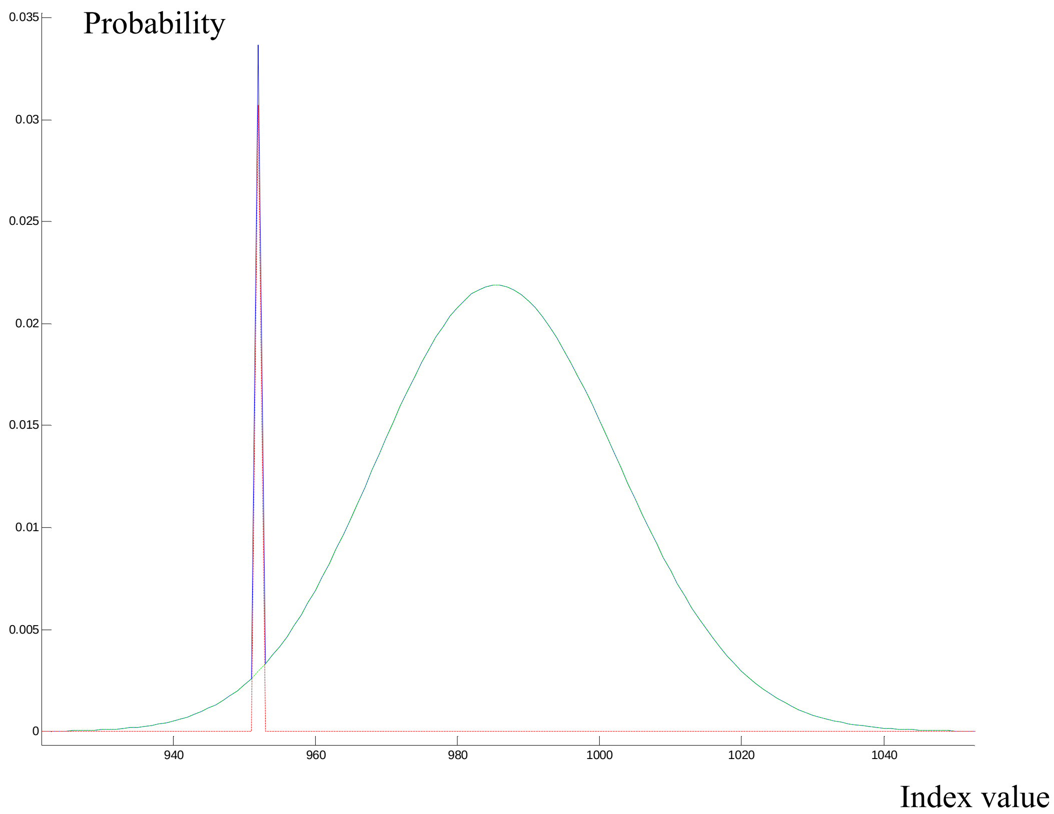

Figure 10: RND function as of February 17th 2006 calculated on options with maturity February 28th 2006.

6.2.3 Test of the election – the day before and the day after

For the option prices the day before the election we had a somewhat peculiar result with the elementary claims method and it was very hard to interpret. The nature of the two-log-normal method is totally different and as it is the sum of two log-normal functions it very well captures situations where there are two possible outcomes as in the case of elections.

The result from using the two-log-normal method the Friday before the election is shown in Figure 11.

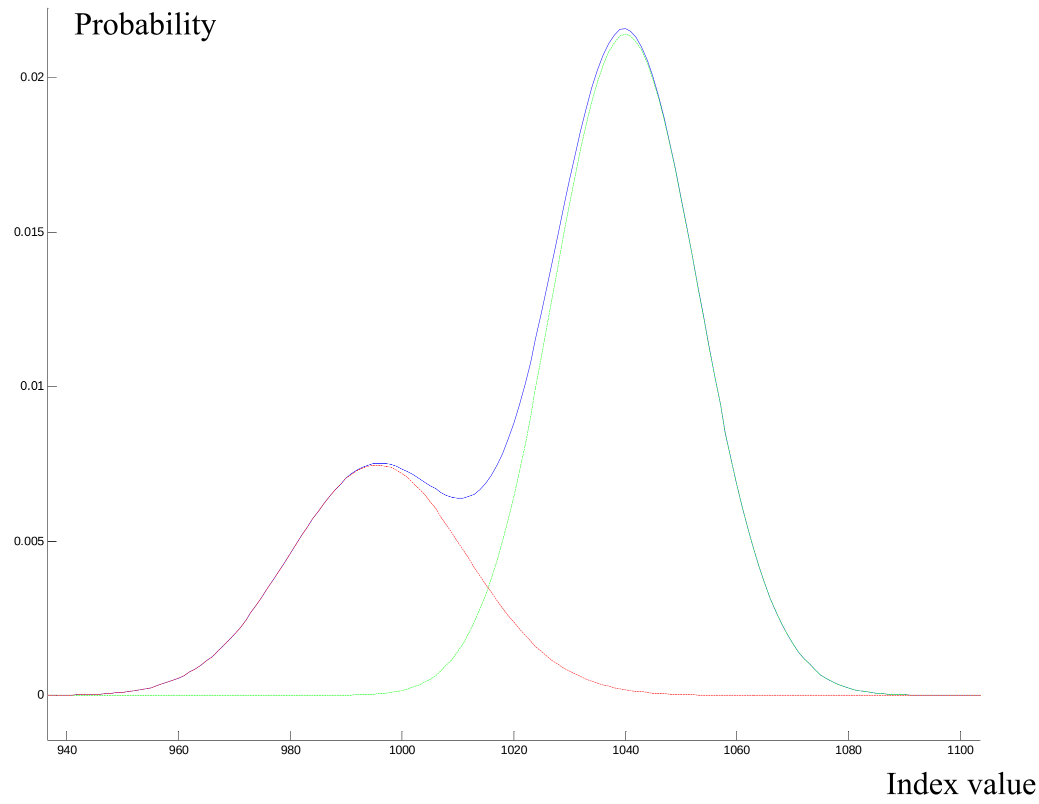

Figure 11: RND function as of September 15th 2006 calculated on options with maturity September 29th 2006

We clearly see that we have two maximums centred on very different prices. One of the maximums seems to be of larger importance than the other, but the variances seem to be roughly the same. This graph very much resembles the graph that we got using the elementary claims method. This further strengthens the two-log-normal method as it obviously catches the same shapes but in a much more clear way.

One natural question to ask how the graph of the RND-function changes after the election results have been announced. Figure 12 show th RND-function the day after the election.

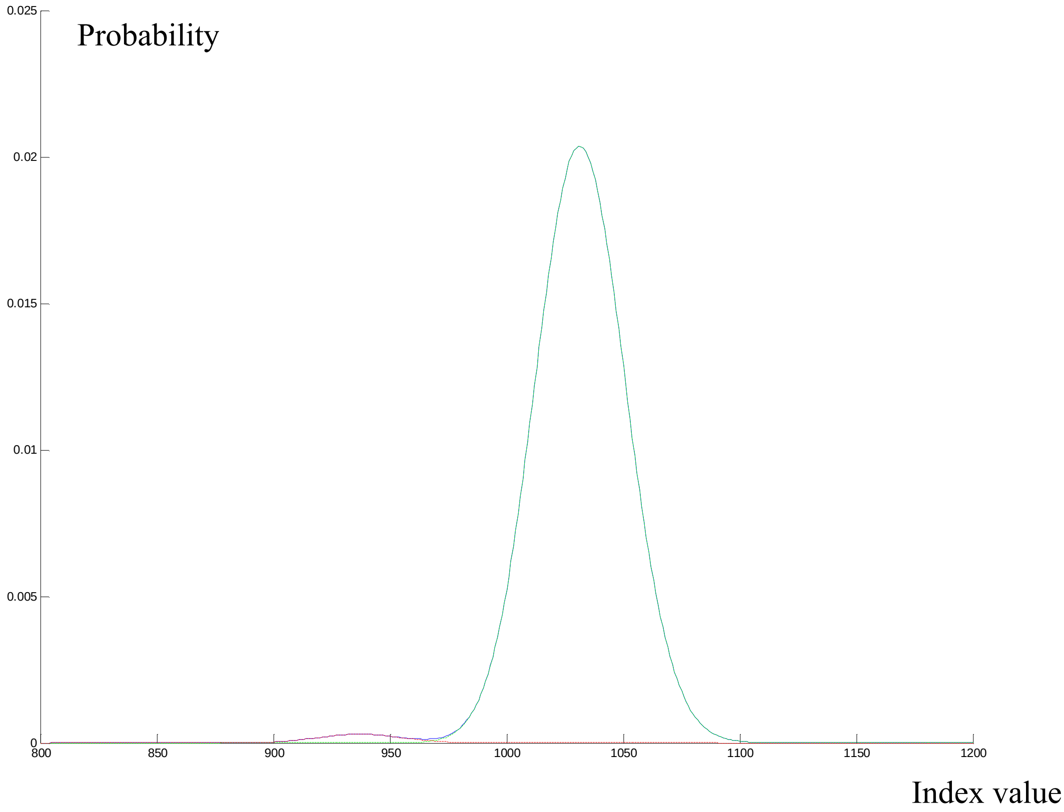

Figure 12: RND function as of September 18th 2006 calculated on options with maturity September 29th 2006

After the election the appearance of the two very distinct maximums vanishes and we end up with a RND-function that is log-normal.

6.2.4 Test of the election – the week before

In the last section we saw how the implicit RND-function changed over the election and more importantly how it had two maximums the Friday before the election. This could of course have been a coincidence. To tackle this we calculate the RND-functions for all the days of the election week.

Figure 13 shows the RND-functions for the five days before the election, and the day after the election.

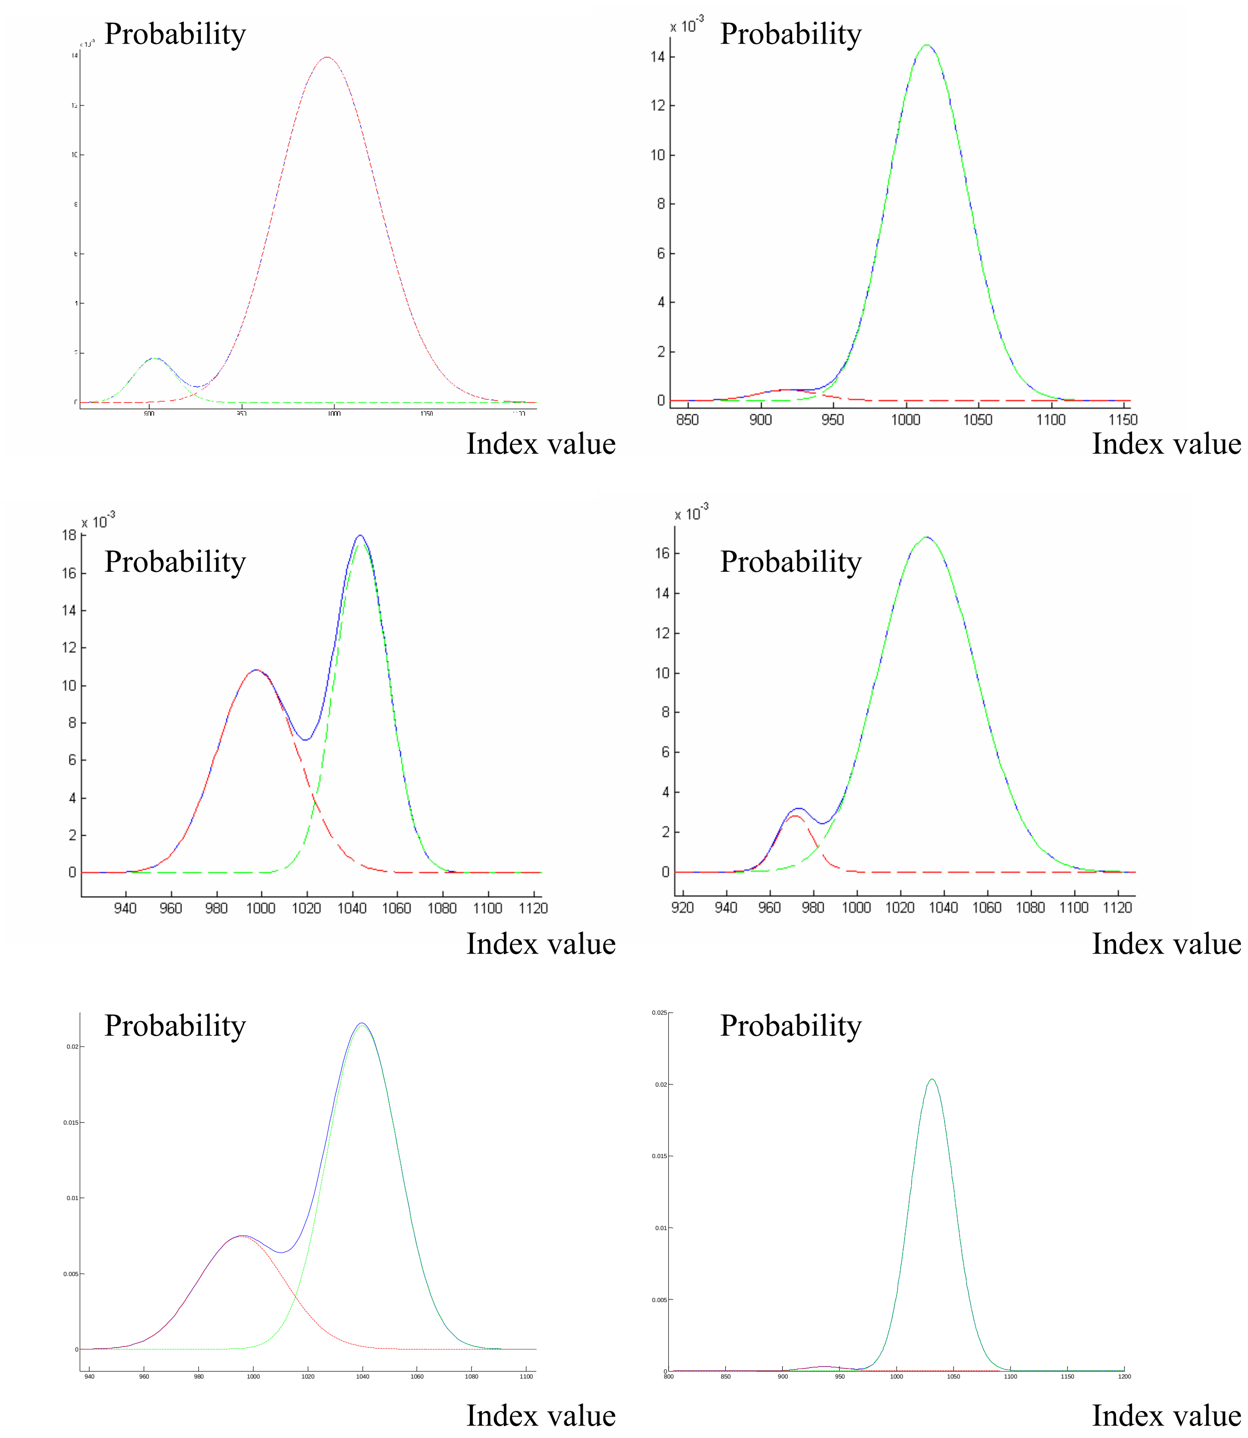

Figure 13: RND functions as of September 11th (upper left), September 12th (upper right), September 13th (middle left), September 14th (middle right), September 15th (lower left), September 18th (lower right), calculated on options with maturity September 29th 2006

We can see that all graphs show signs of two maximums, and in all graphs the left maximum is much smaller than the right maximum. The difference in sizes and shapes are though huge between the days. Considering that the view on the election outcome should reflect the shapes of these graphs, and that the view on the outcome of the election is very unlikely to have changed that much during the election week, the graphs above should be considered as somewhat unstable and the sizes of the different maximums should not be regarded as reliable.

6.2.5 Test of central bank meetings

For the election tests we saw clear implications on two maximums corresponding to the two possible outcomes of the planned economic event. To test for the same appearances for central bank meetings we will study the first three central bank meetings of 2006. We will derive the RND-function the day before the election and the day after the election for all of these meetings.

The RND-functions before and after the first central bank meeting are shown in Figure 14 and Figure 15.

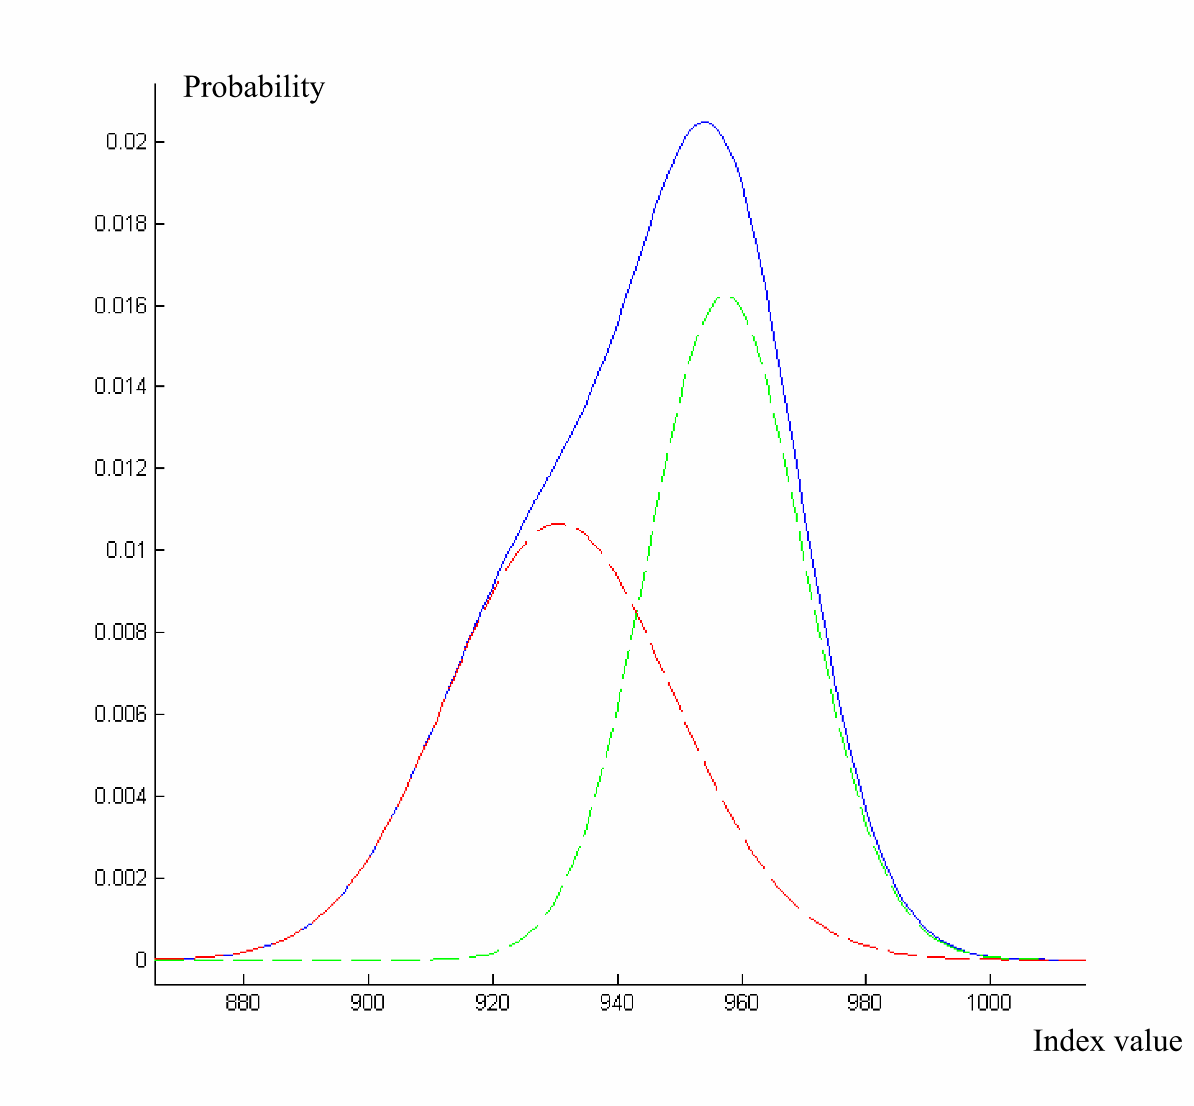

Figure 14: RND function as of 19th January (before the meeting) for options with maturity January 31st 2006

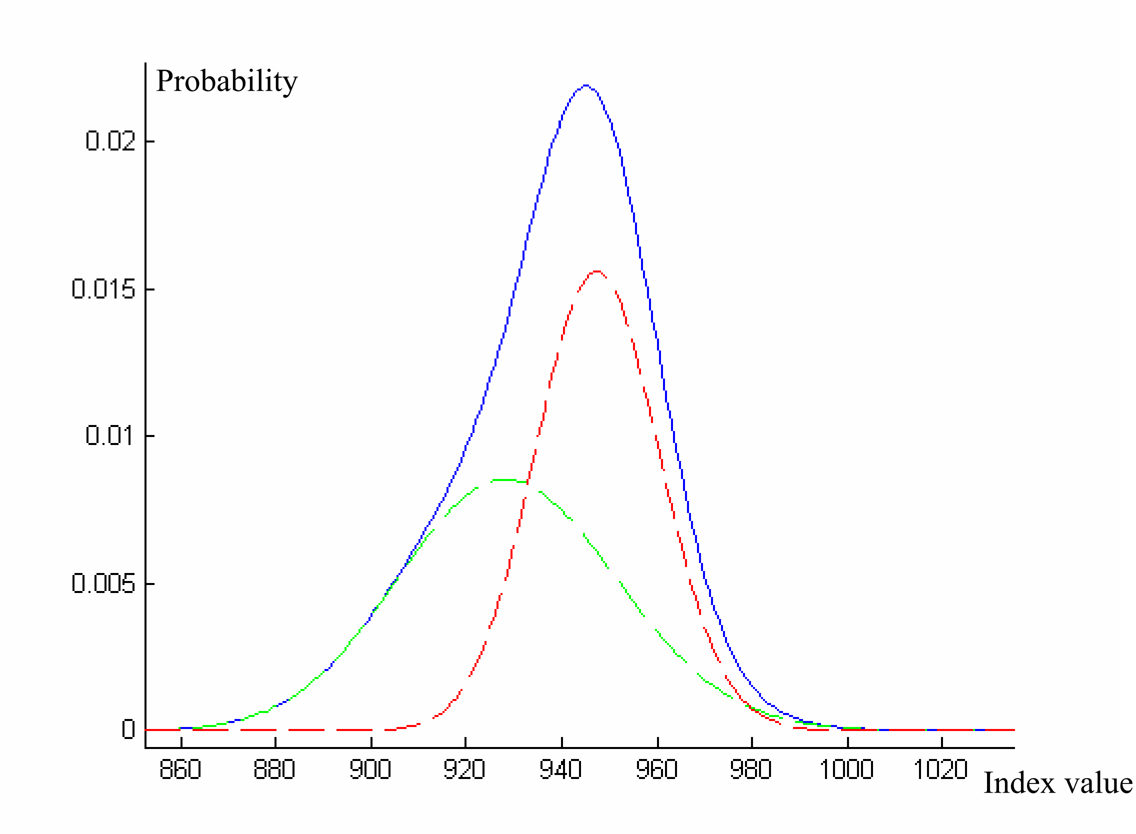

Figure 15: RND function as of 20th January (after the meeting) for options with maturity January 31st 2006

We can see a slight change where the RND-function the day before the publication of the meeting outcomes is somewhat more skewed and the two underlying log-normal functions are somewhat more distinct than the day after the publication of the outcome.

Figure 16 and Figure 17 show the corresponding graphs for the second central bank meeting.

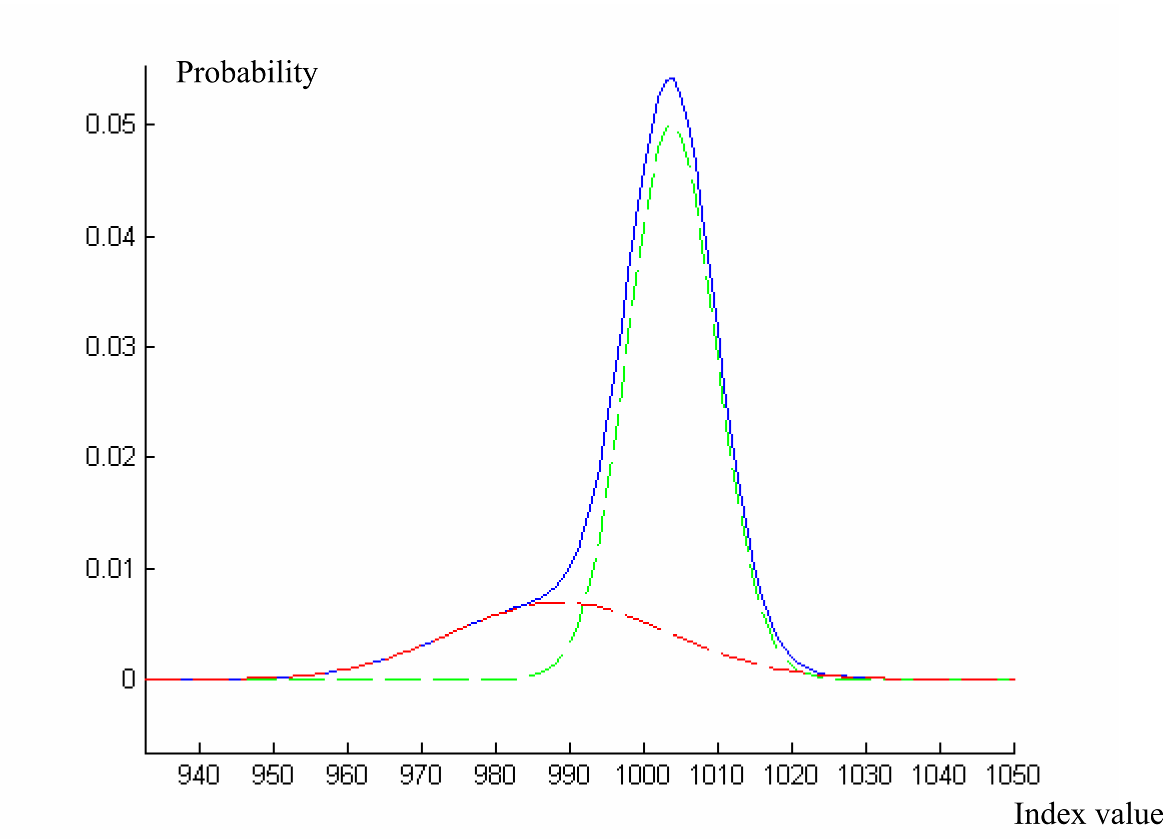

Figure 16: RND function as of 22nd February (before the meeting) for options with maturity February 28th 2006

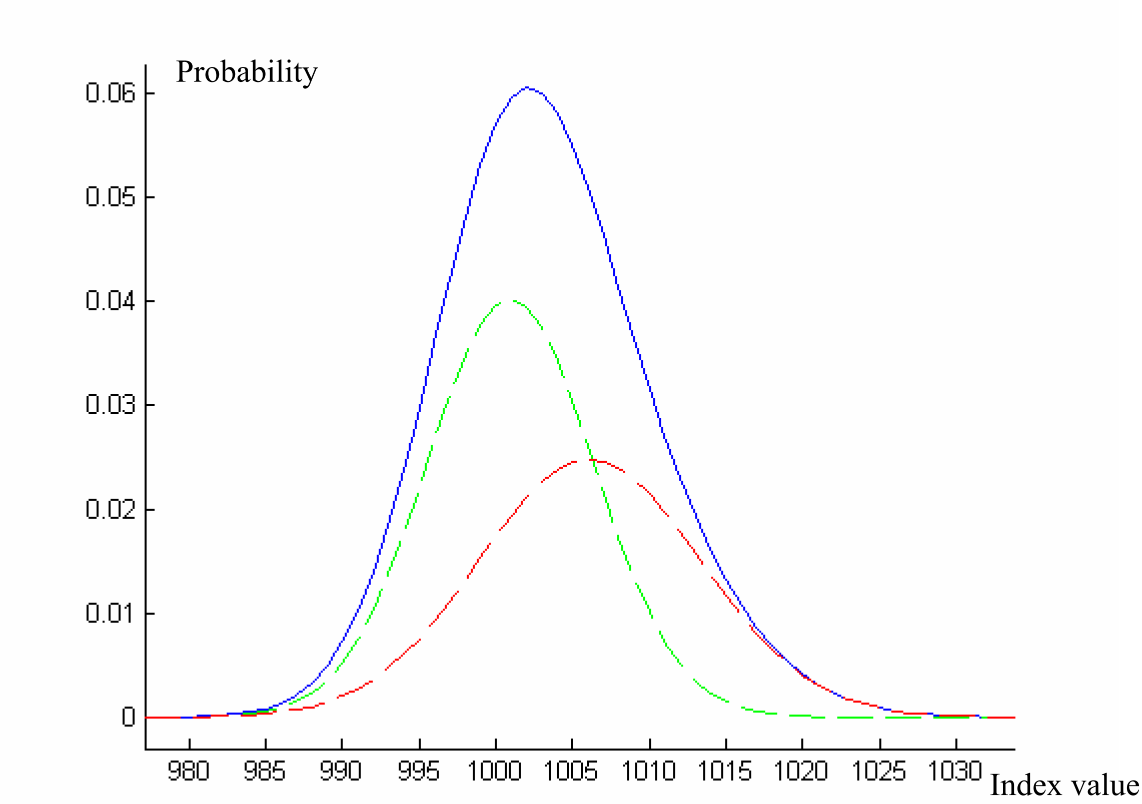

Figure 17: RND function as of February 23rd (after the meeting) for options with maturity February 28th 2006

Again, it is difficult to identify any maximums. With positive interpretation we could possibly identify maximums and even interpret them as having some economic underlying reason, but it is highly spectacular and unreliable and subject to a high degree of subjectivity and data snooping.

Figure 18 and Figure 19 lastly show the RND-function for the third central bank meeting

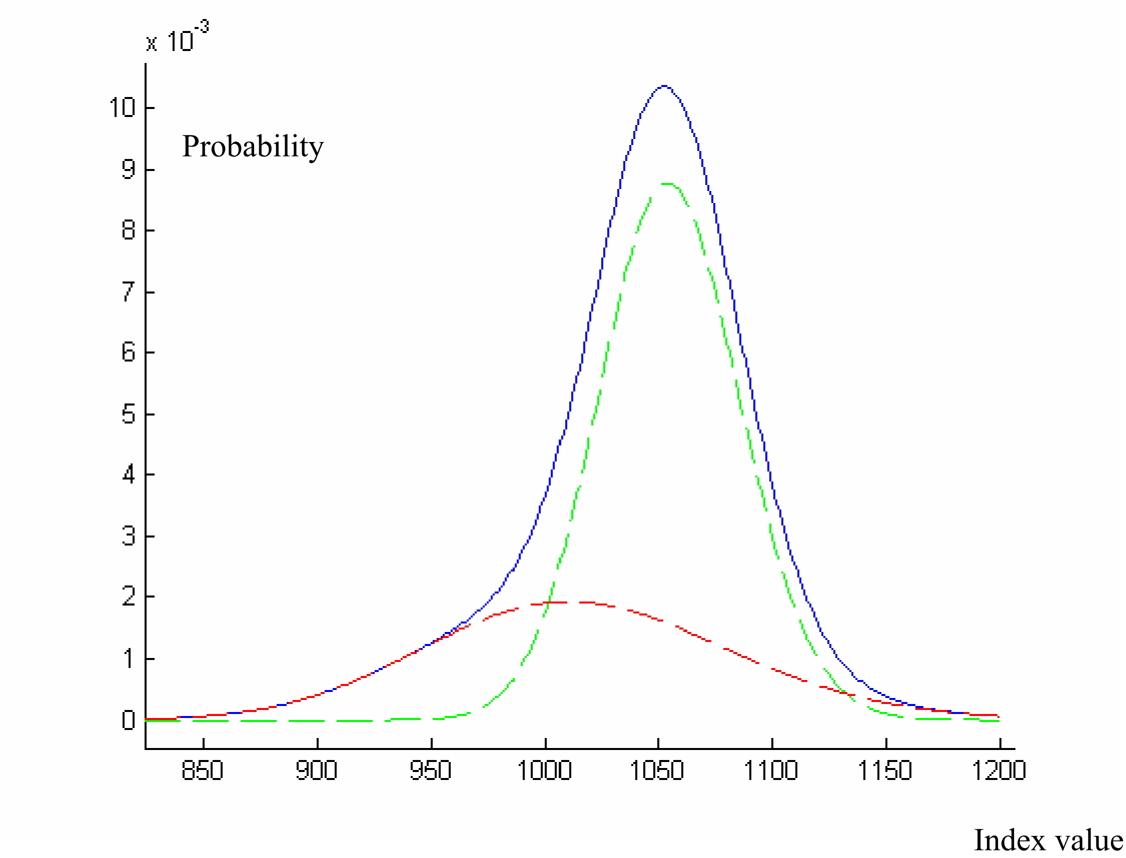

Figure 18: RND function as of April 27th (before the meeting) for options with maturity May 31st 2006

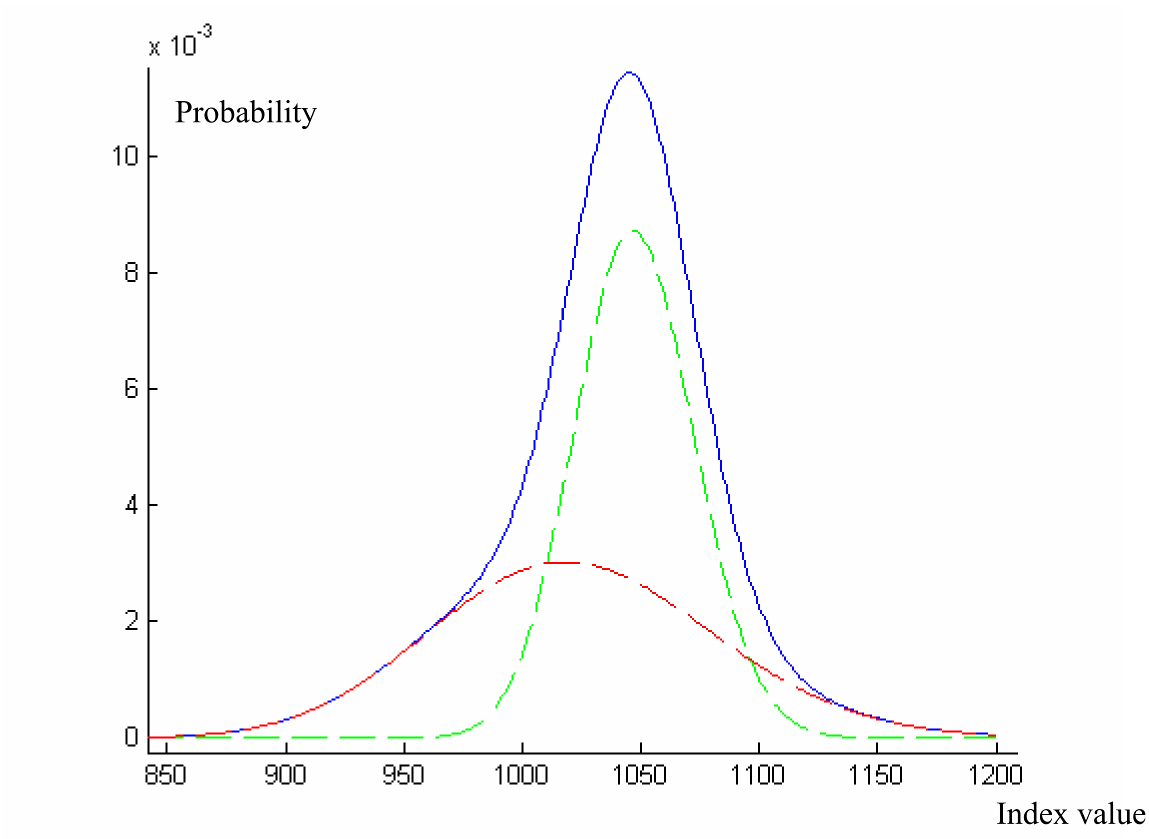

Figure 19: RND function as of 28th April (after the meeting) for options with maturity May 31st 2006

Also here it seems like the difference is too small to draw any real conclusions.

7 Results

7.1 Elementary claims method

The very direct and simple method using elementary claims proves to work very well in practice. The implicit RND functions are centred around the forward price which is expected as the forward price is the expected value of the RND-function. They are fairly smooth around the centre and the variance is about as expected. Also the probabilities diminish from the centre in a way that could suggest a log-normality.

The tails are however highly unreliable and we sometimes see negative values. This suggests the presence of arbitrages, but in practice it is due to illiquid options in the tails and a large bid-ask spread. The implication for us is that we can only rely on very few data points around the centre. As the method calculates a value on the implicit RND-function using three surrounding option prices, we get even fewer data points on the RND function than we have available option prices. Hence, the tail problem and the nature of the method makes the actually usable data points on the RND function very few.

Besides having very few usable data points the method only gives us discrete values and considering our original question about extracting the market's view on planned economic events I find it very hard to draw any conclusions on the market's view on these events. There is a tendency to bimodality in the data around the election, but the interpretation needed to draw that conclusion is highly spectacular, unreliable and unscientific.

To sum up, the elementary claims method is a very direct and easy method and it works very well even on the Swedish somewhat illiquid data. But, it does not help us in answering our question and it is discarded.

7.2 Two-log-normal method

The two-log-normal method also proves to work as expected. We end up with solution centred around the forward price. The variances of the RNDs are reasonable and in line with expectations, and the functions are smooth. We end up with somewhat thicker tails than normal single-log-normal functions. Thicker tails is actually one modification to Black and Scholes used in more sophisticated models and the market generally believes that the tails of the log-normal density are too light, hence the result of thicker tails is in line with expectations.

The minimization is working fine for most start values and days. There do though exist solutions that lack economic meaning as shown in the analysis section. This does not necessarily mean any large problems as the solutions are clearly wrong and they can easily be avoided by changing the start values.

To summarize, the two-log-normal method is working very well, the solutions are stable and reliable. Also, the solutions are smooth and have properties that can answer our question. This is further discussed in the next section.

7.3 Election RND-functions

The RND-function the day before the election shows two clear distinct maximums and the RND-function the day after the election shows no signs of two maximums. This tells us that there was a fundamental difference in the view on the future price development of OMXS30 the day before the election and the day after the election. As the election was more or less the only economic event happening the week-end between these two dates it is natural to assume that it is the election that was the reason for the difference in investors view on the future.

Hence, the day before the election investors believed that there were two possible outcomes and that each of these outcomes had different implications on the future development of the underlying asset OMXS30. Apparently one of the outcomes would have more positive effect on the OMXS30 than the other outcome. Attached to each of these outcomes was a most-probable value of the OMXS30 represented by the two maximums in the implicit RND-function. But, the very value of the OMXS30 given the outcome of the election was still uncertain and therefore we see a log-normality centred around the most probable outcome given the election results.

In the RND-function the day before the election the two maximums seem to be of different sizes both measured as there marginal probabilities and there cumulated probabilities. It is though unclear whether these sizes in the RND-function can give us any implications about whether the market believes one of the outcomes to be more probable than the other. If we accept the RND-function the day before the election to be the very truth we could draw the conclusion that the market believes one of the outcomes to be more probable, but to do that we need to know how stable the solution we have is. Nevertheless, assuming the RND-function is stable and exact we can conclude that the outcome corresponding to a less positive development of the OMXS30 is less probable than the outcome with a more positive implication on the OMXS30.

The Swedish election of 2006 had two sides – the Alliance and the Socialists – of which the polls the week before the election showed an outcome of an Alliance win as the most probable. Without going to deep into a political discussion we could bravely argue that the maximum in the RND with the highest expected value of OMXS30, that is the maximum with the most positive view on the future, corresponds to an Alliance win. With this interpretation, we could conclude that the market regarded an Alliance win as the most probable as this maximum in the RND-function was the biggest. We should though remember that the interpretation of economic events impact on prices is highly uncertain as highlighted by McQueen and Roley (1993) where they argued that an economic event can be seen as having different effects on prices by different investors.

To examine the stability of the RND-functions we calculated them for all of the election week – the week before the election. As OMXS30 changed a lot during this week new option prices were formed every day which can also be seen in the RND-functions being centred around different prices. Hence, we have new prices formed every day and we have a new RND-function for every day. All of the RND-functions showed two maximums. The maximums were though of varying distinctions and of highly varying sizes. The only underlying difference between the days should theoretically be the view on the outcome of the election. Despite the election week being very hectic politically the view upon the election outcome should and did not vary as much as the implicit RND-function implicate. Therefore, the difference in the shapes of the RND-functions seems to be of a much larger quantity than the underlying reason could explain. Because of this we observe a low stability and there is a great uncertainty in any conclusions about the market view on the election.

7.4 Central bank meetings

From the RND-functions on the central bank meetings we see that the difference in the functions before and after the publication of the meeting outcomes is very small. In the cases where there is a difference the difference is very small and it is difficult to tell whether the difference is an adjustment to the log-normal curve to better fit economic practice or if its underlying cause actually is the upcoming meeting. In none of the three meetings we have two maximums.

The reason for choosing to study these meetings is that the interest rate that is decided upon on these meetings has very high economic impact. The question is though if these very meetings have any impact or is the outcome of the meetings more or less already known by the market? Or could it be that the outcome of the meetings is more or less based on the market's expectations, hence do we have a causality problem? No matter what underlying reason, the RND-functions do not give us any information about the market's expectations about the outcome of the central bank meetings.

Mandler (2002) studied European Central Bank meetings using RND-functions and did not find any signs of the meetings having any effect on the RND-functions. He concluded the reason to be too much noise and the meetings having a too small impact on the market. He did even take this one step further and developed a method to calculate the average of RND-functions over many days to see if the original findings were due to lack of data, but also with the new method he did not find anything. Castrén (2005) studied the impact of monetary policy changes in the new member states and found shifts in the RND-functions but still no signs of bimodality. Also Bahra (1997) found these shifts on LIFFE-options, but still he did not find any bimodality.

8 Conclusions

To come back to the original question of this thesis of whether option prices can be used to understand the market's view on planned economic events the conclusion reached is that to a limited extent option prices can be used for that purpose.

First we can conclude that Sweden is a large and liquid enough market for using RND-functions. We had good results from both methods used, and the two-log-normal method had stable solutions. We can though probably only use the two-log-normal method on the most liquid options in Sweden. This result is in line with previous studies on Swedish options on foreign exchange (Javiera 1999).

Secondly, the events that we want to study need to be of a very high economic importance in order for us to be able of drawing any conclusions from the shapes of the implicit RND-functions. Empirically we found that central bank meetings on interest rate changes were not important enough. This is in line with the results found by Mandlers (2002) studies of the European Central Bank meetings. The election of 2006 was though clearly an economic event important enough.

Third, we conclude that we can see the effect and market view on the planned economic events and we can identify the different outcomes of the events in the RND-functions before the outcomes are known by the market. But, we found it difficult to quantify the probabilities of the different outcomes from the RND-functions as the stability of the solutions is low. This is though subject to interpretation and a brave interpreter would conclude that it was possible to forecast the Alliance win in the 2006 election. Our findings of two maximums in the RND-functions are in line with Gemmill's and Apostolos' (1999) study on UK-elections. Our results do show slightly stronger evidence of bimodal densities though we can still not find any correlation to election polls.

To summarise, implicit RNDs could have been used during the week of the election as a complement to the election polls, but only as a complement and the results from it needs to be interpreted with great caution. That is though the case also for classical election polls which are also very unreliable. Hence, the method of using RND-functions might be just as accurate as any of the other methods, and it might fit very well into the quasi scientific method of Election polls.

9 References

9.1 Articles

Arild Syrdal S., 2002, A Study of Implied Risk-Neutral Density Functions in the Norwegian Option Market, *Securities Markets and International Finance Department Working Paper*

Bahra B., 1997, Implied risk-neutral probability density function from option prices: theory and application, *Bank of England*

Castrén O., 2005, Estimating and analysing currency options implied risk-neutral density functions for the largest new EU member states, *European Central Bank Working Paper Series*

Craig B. and Keller J., 2005, The forecast ability of risk-neutral densities of foreign exchange, *Banking and Financial Studies No 05/2005*

Gemmill G. and Apostolos S., 1999, How Useful Are Implied Distributions? Evidence From Stock-Index Options, *Bank for International Settlements*

Jackwerth J.C., 1999, Option-Implied Risk-Neutral Distributions and Implied Binomial Trees: A Literature Review, *Journal of Derivatives*

Javiera A., 1999, GARCH, Implied Volatilities and Implied Distribution: An Evaluation for Forecasting, *Sveriges Riksbank Working Paper Series*

Mandler M., 2002, Comparing risk-neutral probability density function implied by option prices – market uncertainty and ECB-council meetings, *Spring Meeting of Young Economists in Paris*

McQueen G. and Roley V., 1993, Stock Prices, news and business conditions, *Review of Financial Studies*

Nikkinen J., 2003, Normality tests of option-implied risk-neutral densities: evidence from the small Finnish market, *International Review of Financial* Analysi

Rompolis L. and Tzavalis E., 2006, The C-type Gram-Charlier Series Expansion of the Risk Neutral Density Implied by Option Prices, *Mathematical and Computer Modelling*

Wilkens S and Röder K, 2006, The information content of option-implied distributions: Evidence from the Eurex index and interest rate futures options market, *Global Finance Journal*

Äijö J., 2006, Impact of US and UK macroeconomic news announcements on the return distribution implied by FTSE-100 index options, *International Review of Financial Analysis*

9.2 Books

Brockwell P. and Davis R., 2002, Introduction to Time Series and Forecasting, *Springer*

Campbell J., Lo A. and MacKinlay A, 1997, The Econometrics of Financial Markets, *Princeton University Press*

Demmel J., 1997, Applied Numerical Linear Algebra, *Siam*

Heath M., 2002, Scientific Computing, *McGraw-Hill*

Hull J., 2003, Options, Futures and Other Derivatives, *Prentice Hall*

Luenberger D., 1998, Investment Science, *Oxford University Press*

McNeil A., Frey R. and Embrechts P, 2005, Quantitative Risk Management, *Princeton University Press*

Råde L and Westergren B, 1998, Mathematics Handbook for Science and Engineering, *Studentlitteratur*