Stockholm School of Economics
Department of Finance – Master's Thesis
Spring 2011

Risk neutral densities and the September 2008 stock market crash

A study on European data

Misha Wolynski
misha.wolynski@alumni.hhs.se

Martin Theimer
martin.theimer@alumni.hhs.se

Abstract

In this paper, we aim to determine whether the options market predicted the stock market crash of September 15 2008 or reacted to it. In order to do so, we study volatility smiles and RND functions for the EURO STOXX 50 equity index. For our estimated RND functions, retrieved by using the two-lognormal method, we calculate standard deviation, skewness and kurtosis. We find that the options market did not predict the stock market crash. Instead, it reacted to it. Specifically, the reaction consisted of an increase in standard deviation, a decrease in left-skewness and kurtosis and a tendency toward a bimodal shape. Apart from the result regarding the skewness, these findings are consistent with research on earlier stock market crashes. However, earlier studies find that left-skewness increases as a reaction to a stock market crash. Thus, the decreased left-skewness appears to be a finding specific for this particular crash. Lastly, we note that the fact that RNDs seem to lack predictive power does not render them useless, as they can be used to assess market sentiment and how it changes over time, which could be useful for decision-making organs, such as central banks.

Tutor: Professor Tomas Björk

Date and time: June 15 2011, 8.15

Venue: Stockholm School of Economics, Room 348

Discussants: Alexander Pehrsson von Greyerz and Henrik Staaf

Acknowledgements: We would like to thank our tutor Tomas Björk for his guidance and helpful advice throughout the writing of this thesis. We are also grateful to Filip Andersson and Gustaf Linnell for helping us to obtain the necessary data and to Anna Leijon and Victor Salander for valuable comments.

Table of Contents

1. Introduction

Derivative contracts, such as call options and put options, are actively traded in financial markets around the world. Clearly, the price of such a contract reflects the market's view of the likelihood that the contract will yield a positive payoff. Since derivatives are assets whose payoff depends on the state of some underlying asset at some future point in time, it follows that option prices indirectly convey information about the probabilities that the market attaches to the underlying asset being in particular states in the future. By using certain techniques, it is possible to obtain a risk neutral probability density function for the state of the underlying asset at a fixed future point in time from the prices of traded options.[1] This risk neutral density (RND) function can be interpreted as the market's probability distribution for the state of the underlying asset. By studying the RND function, the market's beliefs can be directly examined. For example, it could convey information on whether the market places relatively greater probability on an increase in prices than on a decrease. Furthermore, the evolution of the obtained RND can be used to assess how the market's beliefs change over time. Specifically, it can be used to assess market beliefs about a planned future event, such as an election. It can also be used to look at how market beliefs change around an unplanned event, such as a stock market crash. If market beliefs change prior to such an event, it indicates that the market predicted the event. If, on the other hand, the change in market beliefs comes after the event has occurred, the logical interpretation is that the market did not predict the event, but instead reacted to it.

The late-2000s financial crisis is widely considered the worst financial crisis since the Great Depression. The crisis began in the credit market, particularly the market for mortgage-backed securities based on subprime mortgages.[2] The first indicators of the crisis appeared as early as February and March of 2007, when several subprime lenders declared bankruptcy. However, this did not have an immediate effect on the stock market (see Figure 1 below). Instead, the EURO STOXX 50 index reached a five-year high in June 2007. The decline started in early 2008, when the gravity of the matter became clearer. During 2008, the downturn in the subprime mortgage sector took its toll on major financial institutions heavily invested in subprime mortgage products. On March 16 2008, Bear Stearns was bailed out by the US government in a deal that let J.P. Morgan acquire the bank for less than seven percent of its market value two days prior to the sale. The negative trend continued throughout the spring and summer, but the index level was still comparable to that of 2006. The market did not crash until September 15 2008, when on the same day, Lehman Brothers filed for bankruptcy and Merrill Lynch was sold to Bank of America as a consequence of the bank's subprime

[1] The obtained probability density will be risk neutral, as derivatives are priced under a risk neutral probability measure and not under the real world probability measure. The potentially erroneous conclusions that may arise as a result of this will be elaborated on in later sections. For now, we simply note that since the obtained distributions and probabilities will be risk neutral, one should interpret them with caution.

[2] In this paper, we give a very brief overview of the crisis in order to motivate why we have chosen to look at it. However, the literature on the topic is very extensive and the interested reader should have no problem in finding a book explaining the causes and effects of the crisis in great detail. We would suggest e.g. Authers (2010).

mortgage exposure. After this, stock markets around the world plummeted (as can clearly be seen for the EURO STOXX 50 index in Figure 1 below), bottoming out in March 2009. The fall in levels was accompanied by extreme price volatility of a kind that had not been witnessed since the Great Depression.

Figure 1 – EURO STOXX 50 for the period 2006-01-01 to 2009-12-31

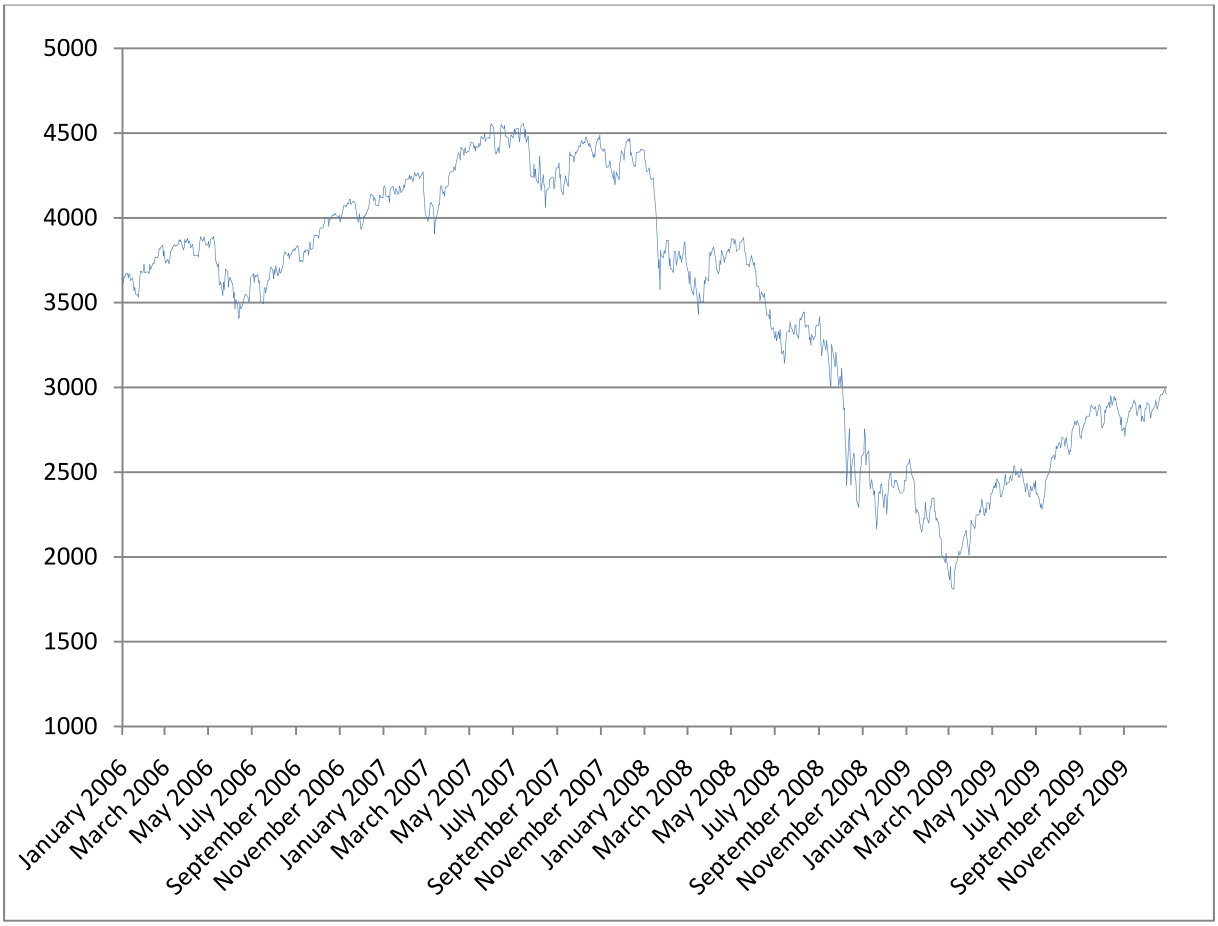

Though the literature on RND functions is extensive, it is mostly focused on methodology rather than application. Furthermore, the articles that do apply the RND framework to data typically look at planned events, such as central bank meetings and elections. Still, there are studies that look at unplanned events, such as various crises and market crashes. However, the bulk of these studies look at earlier stock market crashes, which is not surprising, as we are dealing with a relatively recent event. To our knowledge, no study conducted on the September 2008 stock market crash has been performed on European data. This gives us the opportunity to apply the RND framework in a new setting.

In this paper, we intend to use the RND framework to study the stock market crash of September 15 2008. Specifically, we will look at the evolution of the RND function of the EURO STOXX 50 equity index before and after September 15 2008 to try to determine whether the options market predicted the stock market crash or reacted to it. We will also look at implied volatilities (i.e. the volatilities implied for the Black and Scholes (1973) model by market prices of options), as these are closely linked to the shape of the RND function.

We proceed by presenting the necessary theoretical framework in the next section, before providing a brief overview of the previous research on the matter in section 3. In section 4, we introduce our data and explain the procedure that is used to extract reliable observations from the initial data set. Section 5 describes our methodology and explains how the theoretical framework is applied to the data. In section 6, we present and discuss our results, before summing them up and presenting more general conclusions in section 7.

2. Theoretical framework

In this section, we present the theoretical framework necessary to retrieve RND functions and implied volatilities from the data. Our aim is to present the framework in a way that is intuitively accessible rather than mathematically rigorous. However, given the nature of the subject, a rather extensive use of mathematics is necessary.

2.1. Risk neutral valuation

Risk neutral valuation was first derived by Cox and Ross (1976). The authors show that if it is possible to find an analytical expression in the form of a differential or difference-differential equation that every contingent claim must satisfy and in which one of the original model parameters does not appear, this parameter can be chosen so that the underlying asset earns the risk free rate. The value of the contingent claim can then be obtained by calculating the expected value, using the modified parameter, and then discounting at the risk free rate. Harrison and Kreps (1979) formalize this approach and make it more rigorous by introducing the theory of equivalent martingale measures. They show that the method proposed by Cox and Ross is equivalent to changing the probability measure from the real world probability measure $\mathbb{P}$ to the equivalent[3] martingale[4] measure $\mathbb{Q}$. For obvious reasons, this measure is also commonly referred to as the risk neutral measure. The value at time t of a contingent claim X maturing at time T can be obtained by using risk neutral valuation as:

$$\Pi_t = \mathbb{E}^{\mathbb{Q}}\left[e^{-\int_t^T r_u du} \cdot X_T | \mathcal{F}_t^S\right] \tag{1}$$

In the equation above, Π_t denotes the time t price of the contingent claim X, $\mathbb{E}^{\mathbb{Q}}[\cdot]$ denotes the expected value taken under the probability measure $\mathbb{Q}$, r denotes the risk free rate and $\mathcal{F}_t^S$ denotes the

[3] Two measures are said to be equivalent if for the two measures (here denoted by $\mathbb{P}$ and $\mathbb{Q}$) on the measureable space $(\Omega, \mathcal{F})$, it holds that:

$$\mathbb{P}(\mathrm{A}) = 0 \leftrightarrow \mathbb{Q}(\mathrm{A}) = 0 \;\forall\; \mathrm{A} \in \mathcal{F}$$

In words, this means that the two measures agree on all impossible events. This implies that the two measures also agree on all certain events, as a certain event is the compliment of an impossible event. Hence, two equivalent measures agree on all impossible and on all certain events. Equivalence between two measures $\mathbb{P}$ and $\mathbb{Q}$ is denoted by $\mathbb{P} \sim \mathbb{Q}$.

[4] A stochastic process M_t is said to be martingale if $\mathbb{E}[|\mathrm{M_t}|] < \infty \;\forall\; t \in [0, \infty)$ and $\mathbb{E}[\mathrm{M_t}|\mathcal{F}_\mathrm{s}] = \mathrm{M_s}$ for every pair s, t, such that $s < t$. The latter condition is commonly referred to as the martingale property. In words, it means that the best prediction of the value of the process at any future point in time, given all available information, is the current value of the process. The term "equivalent martingale measure" arises because the discounted price process is a martingale under $\mathbb{Q}$.

filtration generated by the price process of the underlying asset S over the period $[0, t]$. Thus, the expression entails computing a conditional expectation under $\mathbb{Q}$ at time t. Using more compact notation, this can be rewritten as:

$$\Pi_t = \mathbb{E}_t^{\mathbb{Q}}\left[e^{-\int_t^T r_u du} \cdot X_T\right] \tag{2}$$

It can be shown that if an equivalent martingale measure $\mathbb{Q}$ exists, the market is free from arbitrage. If the measure is unique, the market is referred to as complete, meaning that all contingent claims can be replicated using other assets. This also implies that the arbitrage free price is unique.[5]

For a derivative, the payoff of the contingent claim X at maturity (i.e. at time T) can be expressed as a function $h(\cdot)$ of the value of the underlying asset S at time T, i.e. $X_T = h(S_T)$. Expression (2) then becomes:

$$\Pi_t = \mathbb{E}_t^{\mathbb{Q}}\left[e^{-\int_t^T r_u du} \cdot h(S_T)\right] \tag{3}$$

One should note that since the expected value of a product does not generally equal the product of the expected values (i.e. $\mathbb{E}[\mathrm{A} \cdot \mathrm{B}] \neq \mathbb{E}[\mathrm{A}] \cdot \mathbb{E}[\mathrm{B}]$), the conditional expectation above can be rather difficult to compute. Therefore, the simplifying assumption of a constant risk free interest rate over the time period $[t, T]$ is usually made, i.e. $r_u = r \; \forall \, u \in [t, T]$. Given this assumption, it holds that $e^{-\int_t^T r_u du} = e^{-r(T-t)}$. Since $e^{-r(T-t)}$ is a constant, it can be taken out of the conditional expectation operator. The resulting expression is:

$$\Pi_t = e^{-r(T-t)} \cdot \mathbb{E}_t^{\mathbb{Q}}[h(S_T)] \tag{4}$$

Thus, it is clear that in order to obtain Π_t, all that is needed is the probability density function of S_T at time t under the equivalent martingale measure $\mathbb{Q}$. This is the previously mentioned RND function, denoted by $q_t(S_T)$. Assuming that the RND function $q_t(S_T)$ is known, the conditional expectation in expression (4) can be computed as[6]:

$$\mathbb{E}_t^{\mathbb{Q}}[h(S_T)] = \int_0^{\infty} q_t(S_T) h(S_T) dS_T \tag{5}$$

Consequently, the price of the contingent claim at time t can be obtained as:

$$\Pi_t = e^{-r(T-t)} \int_0^{\infty} q_t(S_T) h(S_T) dS_T \tag{6}$$

Expression (6) is typically used in one of two ways. The focus is either on computing the price Π_t, in which case certain assumptions regarding the price process of the underlying asset in order to obtain

[5] For more on the connection between equivalent martingale measures and arbitrage, see e.g. Björk (2004).

[6] The observant reader may note that the integral in expression (5) is taken over the interval $[0, \infty)$ rather than $(-\infty, \infty)$, which is the correct integration interval when computing an expected value. This is a result of the fact that S_T is only defined on the interval $[0, \infty)$, as a price cannot take negative values. Hence, it is assumed that $q_t(S_T) = 0 \; \forall \, S_T \in (-\infty, 0)$. Consequently, integrating over $[0, \infty)$ will yield the same result as integrating over $(-\infty, \infty)$ in this case.

$q_t(S_T)$ are made, or on using the available prices of traded derivatives to estimate the RND function $q_t(S_T)$ implied by those prices. The focus of this paper is on the latter.

It should be pointed out that the usage of risk neutral valuation in no way implies the (obviously incorrect) assumption that investors are risk neutral. Instead, usage of the equivalent martingale measure can be viewed as a different approach to modeling risk. Instead of compensating for higher risk by using a higher discount rate, the probabilities for "good" outcomes are adjusted down (and hence, the probabilities for "bad" outcomes are adjusted up, as the total probability has to sum to one). Hence, the expected value under $\mathbb{Q}$ will be lower than under $\mathbb{P}$, thus eliminating the need for a higher discount rate to obtain the correct price. Consequently, the expected rate of return under the equivalent martingale measure $\mathbb{Q}$ is equal to the risk free rate r for all assets.

2.2. The Black-Scholes model

In their seminal paper, Black and Scholes (1973) developed the model that has since become the benchmark in option pricing. The model, known simply as the Black-Scholes model, postulates that the price process for the underlying asset follows a geometric Brownian motion (GBM), i.e.:

$$\frac{dS_t}{S_t} = \mu dt + \sigma dW_t^{\mathbb{P}} \tag{7}$$

In the equation above, μ represents the drift term and σ represents the diffusion term for the return process of the underlying asset.[7] $W_t^{\mathbb{P}}$ denotes a Wiener process under the real world probability measure $\mathbb{P}$. Recall that for a Wiener process, the increments are normally distributed with mean 0 and variance dt, i.e. $dW_t^{\mathbb{P}} \epsilon N(0, dt)$.[8] Therefore, it is clear that the return process for the underlying asset under the real world probability measure $\mathbb{P}$ has normally distributed increments. Hence, the price process has lognormally distributed increments. Thus, the Black-Scholes dynamics for the price of the underlying asset imply that it is lognormally distributed. As the present price of the underlying asset is known, the assumption of a stochastic process for the price of the underlying asset makes it possible to derive the distribution of the price of the underlying at some future point in time.

Black and Scholes show that the price of a derivative is given by $\Pi_t = f(t, S_t)$, where the pricing function $f(\cdot)$ satisfies the partial differential equation (PDE) below, commonly referred to as the Black-Scholes PDE:

$$\frac{df}{dt}(t, S_t) + rS_t \frac{\partial f}{\partial S_t}(t, S_t) + \frac{1}{2}\sigma^2 S_t^2 \frac{\partial^2 f}{\partial S_t^2}(t, S_t) - rf(t, S_t) = 0 \tag{8}$$

[7] It is important to note that μ and σ denote the drift and diffusion terms for the return process and not for the price process. The drift and diffusion terms for the price process at time t are μS_t and σS_t respectively, and thus vary with t as S_t varies with t, whereas μ and σ are constants and thus time-invariant.

[8] For more on Wiener processes and their applications in finance, see e.g. Kijima (2002) or Björk (2004).

The reader familiar with PDEs will notice that the expression above is insufficient in order to obtain a specific solution. In order to do so, a boundary condition is also necessary. The boundary condition is given by:

$$f(T, S_T) = h(S_T) \tag{9}$$

Recall that $h(S_T)$ is the payoff function of the derivative at maturity. Now, there is a unique solution for this PDE, so in a Black-Scholes economy[9], there is a unique price for every derivative. Notice that all derivatives in the economy have to satisfy the PDE in expression (8). The only difference between derivatives lies in the boundary condition, i.e. expression (9).

Black and Scholes also derive explicit formulas for the pricing of European call and put options. Recall that for a European call, the payoff function is $h(S_T) = \max(S_T - K, 0)$, where S_T is the price of the underlying asset at maturity and K is the exercise price. Similarly, for a European put, the payoff function is $h(S_T) = \max(K - S_T, 0)$. Thus, the boundary condition in expression (9) is set to the respective payoff function.

The Black-Scholes formula for European calls and puts respectively, is:

$$\begin{aligned} c_t &= S_t \boldsymbol{\Phi}(\mathrm{d}_1) - \mathrm{e}^{-\mathrm{r}(\mathrm{T}-\mathrm{t})} \mathrm{K} \boldsymbol{\Phi}(\mathrm{d}_2) \\ p_t &= e^{-r(T-t)} K \boldsymbol{\Phi}(-d_2) - S_t \boldsymbol{\Phi}(-d_1) \end{aligned} \tag{10}$$

The parameters d_1 and d_2 are given by:

$$\begin{aligned} d_1 &= \frac{\log\left(\frac{S_t}{K}\right) + \left(r + \frac{\sigma^2}{2}\right)(T-t)}{\sigma\sqrt{T-t}} \\ d_2 &= \frac{\log\left(\frac{S_t}{K}\right) + \left(r - \frac{\sigma^2}{2}\right)(T-t)}{\sigma\sqrt{T-t}} \end{aligned} \tag{11}$$

In expressions (10) and (11) above, $\boldsymbol{\Phi}$ denotes the cumulative distribution function of the standard normal distribution[10] and $\log(\cdot)$ denotes the natural logarithm. The parameters involved have already been defined, though we will return to the parameter σ shortly. Also, one should note that the expected return of the underlying asset, μ, is not included in the valuation formulas. This is to be expected, as μ is the expected return of the underlying asset under the real world probability measure $\mathbb{P}$. However, as has already been explained, risk neutral valuation is carried out under the equivalent martingale measure $\mathbb{Q}$, where the expected return on all assets is the risk free rate r.

[9] See Black and Scholes (1973) for all of the assumptions that make up a Black-Scholes economy.

[10] Recall that a cumulative distribution function for a random variable X is given by $F(x) = P(X \leq x)$, where P is an arbitrary probability measure (not necessarily the real world probability measure $\mathbb{P}$). This can be defined in terms of the probability density function $f(x)$ as $F(x) = \int_{-\infty}^{x} f(u)du$. For the standard normal distribution, the probability density function is given by $\phi(x) = \frac{e^{-\frac{x^2}{2}}}{\sqrt{2\pi}}$. Hence, the cumulative distribution function is given by $\boldsymbol{\Phi}(x) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{x} e^{-\frac{u^2}{2}} du$. Unfortunately, there is no way to express this integral analytically, so it has to be evaluated numerically.

Garman and Kolhagen (1983) extend the Black-Scholes model, enabling it to cope with the presence of two interest rates. This is done for foreign exchange (FX) options, where both the domestic risk free rate r_d and the foreign risk free rate r_f must be taken into account. The resulting difference is that while only K is discounted in the original Black-Scholes model, the Garman-Kolhagen model also discounts the price of the underlying asset S_t, at the foreign risk free rate r_f, while K is obviously still discounted at the domestic risk free rate r_d. The reason for discounting the underlying asset is that the investor forgoes the foreign interest rate by owning the option rather than the underlying asset directly. Though Garman and Kolhagen do their derivation for FX options, it is clear that the same framework can be applied to any type of underlying asset where there is a continuous return that the investor relinquishes by owning the option rather than the underlying asset. With the commonly made simplifying assumption that equity indices pay a continuous dividend yield rather than discrete dividends, this is clearly the case for index options. Denoting the dividend yield by q, the extended Black-Scholes formula for European index options becomes:

$$\begin{aligned} c_t &= e^{-q(T-t)} S_t \boldsymbol{\Phi}(d_1) - e^{-r(T-t)} K \boldsymbol{\Phi}(d_2) \\ p_t &= e^{-r(T-t)} K \boldsymbol{\Phi}(-d_2) - e^{-q(T-t)} S_t \boldsymbol{\Phi}(-d_1) \end{aligned} \tag{12}$$

The parameters d_1 and d_2 are now given by:

$$\begin{aligned} d_1 &= \frac{\log\left(\frac{S_t}{K}\right) + \left(r - q + \frac{\sigma^2}{2}\right)(T-t)}{\sigma\sqrt{T-t}} \\ d_2 &= \frac{\log\left(\frac{S_t}{K}\right) + \left(r - q - \frac{\sigma^2}{2}\right)(T-t)}{\sigma\sqrt{T-t}} \end{aligned} \tag{13}$$

One should note that while options on single stocks are typically American, index options are typically European. Hence, the Black-Scholes formula is particularly suitable for working with index options. As we are dealing with index options in this paper, the model presented in expressions (12) and (13) will be used.

2.2.1. Implied volatility and the volatility smile

Given the framework presented above, the price of a European index option is a function of six parameters[11], namely the current level of the index (S_t), the exercise price (K), the time to maturity ($T-t$), the risk free interest rate (r), the dividend yield (q) and the volatility of the index return (σ). The values of the first five parameters at time t are readily observable, so there is generally little

[11] A word on notation might be appropriate at this point. Since the price of an equity index option under the Black-Scholes model is a function of six parameters, the most general way to denote the time t option price function is $c_t(S_t, K, T-t, r, q, \sigma)$ and $p_t(S_t, K, T-t, r, q, \sigma)$ for European calls and puts respectively. If a more complex model than the Black-Scholes model is used, even more parameters become involved. Clearly, writing them all out every time is highly impractical. Therefore, we will typically use more compact notation and only explicitly write the most relevant variables for the particular application. Thus, when we write e.g. $c_t(K)$, it does not mean that the exercise price is the only variable that the call price depends on, but rather that it is the one most relevant for the task at hand.

disagreement about them. The value of the parameter σ, however, is unobservable. One should note that since the values of the other five parameters are known, the option price at time t can be considered a function of σ only. Hence, it is possible to obtain an estimate of σ from the prices of traded options by choosing σ so that the Black-Scholes price corresponds to the market price. This type of estimate of σ is known as implied volatility. One should note that the Black-Scholes price of an option (call or put) is a monotonically increasing function of σ. Therefore, a higher implied volatility, ceteris paribus, means that an option is trading at a higher price.

Under the Black-Scholes assumptions, the price of the underlying asset evolves according to a GBM. In this context, σ should be constant, as can clearly be seen in expression (7). Thus, the implied volatility should not vary with either exercise price or time to maturity. Rubinstein (1994) points out that the Black-Scholes framework can be easily adjusted to allow for time-dependent implied volatility. Still though, the implied volatility should be constant for different exercise prices, given a fixed maturity. However, implied volatility is usually observed in the market as a convex function of exercise price. Because of this, implied volatility as a function of exercise price, $\sigma(K)$, is typically referred to as the "volatility smile".[12]

Rubinstein (1994) studies options on the S&P 500 index and finds that the assumption of a constant implied volatility for different exercise prices, given a fixed maturity, held fairly well until the stock market crash of 1987. Since then, the implied volatility as a function of the exercise price has exhibited the reverse skew (or "smirk") shape that can be seen in Figure 2 below and that is characteristic for equity index options today. Rubinstein suggests that one possible explanation for this is "crash-o-phobia", i.e. that the market prices out-of-the-money (OTM) puts (and hence in-the-money (ITM) calls as a result of the put-call parity) relatively higher than options with higher exercise prices in order to provide insurance against stock market crashes. Another possible explanation is the "leverage effect", proposed by Black (1976), though Figlewski and Wang (2000) convincingly argue against this explanation.

[12] Though typically convex, the shape of the function $\sigma(K)$ is not always a regular smile. Depending on the underlying asset, the shape can range from a reverse skew to a forward skew, with the regular smile somewhere in between.

Figure 2 – Volatility smile on December 22 2008, three months

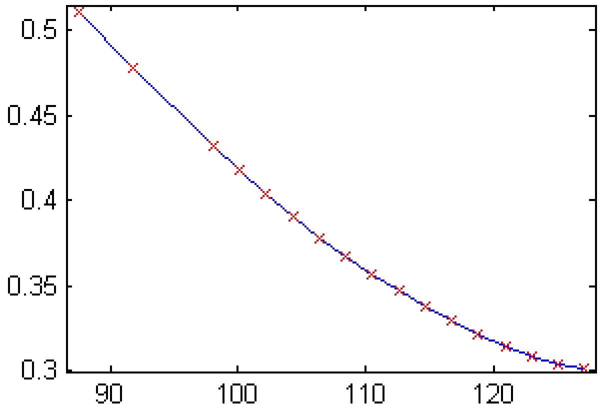

Regardless what explanation for it one chooses to believe, it is clear that the existence of the implied volatility smile indicates that market participants make more complex assumptions than a GBM about the path of the underlying asset price. As a result, they attach different probabilities to the possible values of the underlying asset at maturity than those that are consistent with a lognormal distribution. Bahra (1997) points out that the extent of the convexity of the smile curve indicates the degree to which the market RND function differs from a lognormal (Black-Scholes) RND. Specifically, a more convex volatility smile function indicates that greater probability is attached to extreme outcomes of S_T. As a result, the market RND will have fatter tails than those associated with a lognormal distribution. Bahra further notes that the slope of the volatility smile function is related to the skewness of the market RND function. A positive slope implies an RND function that is more right-skewed than a lognormal RND function, whereas a negative slope implies that the market RND function is more left-skewed than a lognormal RND function. Thus, we would expect an equity index to exhibit RND functions that are more left-skewed than a lognormal RND function, as the volatility smirk has a negative slope. Overall, it is clear that there is a close connection between the volatility smile and the market RND function. This will become apparent when we look at different methods for estimating the RND function in the next section.

2.3. The RND function

Before going into the various ways of recovering RND functions, it is useful to review the concept of elementary claims. An elementary claim is the most fundamental state-contingent claim[13] and was introduced by Arrow (1964), based on the time-state preference model of Arrow and Debreu (1954). For this reason, it is commonly referred to as an Arrow-Debreu security. An Arrow-Debreu security is an asset that pays one unit of currency at a future time T if the underlying asset S is in a particular state at that time, and zero otherwise. The price of an Arrow-Debreu security for a certain state is simply the risk neutral probability of that state occurring, multiplied by the discounted value of one unit of

[13] A state-contingent claim is a claim whose value depends on the future state of some variable. Hence, it should be clear that any derivative constitutes a state-contingent claim.

currency. Hence, if Arrow-Debreu securities were traded, recovering the risk neutral probability $\mathbb{Q}_t(S_T = K)$ would simply entail observing the price for the Arrow-Debreu security corresponding to the future state $S_T = K$ and compounding it by the risk free rate. Doing this across all states would yield all risk neutral probabilities, thus making it trivial to obtain the RND function $q_t(S_T)$. However, the securities are not traded and have to be replicated. This can be achieved by taking a long position in a so-called "butterfly spread".

A butterfly spread is a portfolio, denoted by P_t, of European call options[14], formed by taking a short position in two European call options with exercise price K, a long position in one European call option with exercise price $K + \mathbf{\Delta}K$ and a long position in one European call option with exercise price $K - \mathbf{\Delta}K$, where $\mathbf{\Delta}K$ represents the constant step size between adjacent exercise prices. Notice that if $S_T = K$, the payoff of a butterfly spread is equal to $\mathbf{\Delta}K$, and that if $S_T = n \cdot \mathbf{\Delta}K,\ n \epsilon \mathbb{Z}$, the payoff is zero. Thus, by investing $\frac{1}{\mathbf{\Delta}K}$ in a butterfly spread, the payoff is one when $S_T = K$ and zero elsewhere. Hence, a discrete approximation of an elementary claim for a given future state $S_T = K$ is given by:

$$\left.\frac{P_t}{\mathbf{\Delta}K}\right|_{K=S_T} = \left.\frac{c_t(K + \mathbf{\Delta}K) - 2c_t(K) + c_t(K - \mathbf{\Delta}K)}{\mathbf{\Delta}K}\right|_{K=S_T} \tag{14}$$

In the expression above, $c_t(K)$ denotes the current (time t) price of a European call option with exercise price K and expiry date T. As this expression replicates an elementary claim, it is clear that the risk neutral probability for the future state $S_T = K$ is given by:

$$\mathbb{Q}_t(S_T = K) = e^{r(T-t)} \left.\frac{c_t(K + \mathbf{\Delta}K) - 2c_t(K) + c_t(K - \mathbf{\Delta}K)}{\mathbf{\Delta}K}\right|_{K=S_T} \tag{15}$$

Hence, the risk neutral probability density for the state will be given by:

$$q_t(S_T = K) = e^{r(T-t)} \left.\frac{c_t(K + \mathbf{\Delta}K) - 2c_t(K) + c_t(K - \mathbf{\Delta}K)}{(\mathbf{\Delta}K)^2}\right|_{K=S_T} \tag{16}$$

Clearly, this framework is not ideal, as it only allows us to replicate discrete states of S_T, spaced by the discrete distance $\mathbf{\Delta}K$. However, one should note that the fraction in expression (16) is the second order central finite difference approximation, i.e. an approximation of the second order derivative of $c_t(K)$ with respect to K. Thus, it is clear that:

$$\lim_{\mathbf{\Delta}K \to 0} q_t(S_T = K) = e^{r(T-t)} \left.\frac{\partial^2 c_t(K)}{\partial K^2}\right|_{K=S_T} \tag{17}$$

[14] As a consequence of the put-call parity, a butterfly spread can also be formed by using European puts. However, following the approach of Breeden and Litzenberger (1978), we use calls throughout. To see how to construct a butterfly spread with puts, see e.g. Hull (2006).

In words, this means that if European call options for all possible exercise prices were traded (i.e. as $\Delta K \to 0$), the probability density for all possible future states of S_T could be obtained. Applying expression (17) across the continuum of all possible states, the RND function $q_t(S_T)$ is obtained as:

$$q_t(S_T) = e^{r(T-t)} \frac{\partial^2 c_t(K)}{\partial K^2} \quad (18)$$

This is the famous result arrived at in the seminal paper by Breeden and Litzenberger (1978).[15] It is important to note that since the derivation of expression (18) does not make any assumptions about the dynamics of the underlying price process, it can be used to obtain the implied RND function irrespective of what the underlying price process looks like.

2.3.1. Techniques for estimating the RND function

The simplest way to estimate the RND function is to derive a risk neutral histogram for it (an example can be seen in Figure 3 below). This is done by using expression (15) for every exercise price K. By applying this technique to all available exercise prices K for a certain maturity $T - t$, discrete approximations of the implied risk neutral probabilities for that maturity is obtained.

Figure 3 – Risk neutral histogram for December 22 2008, three months

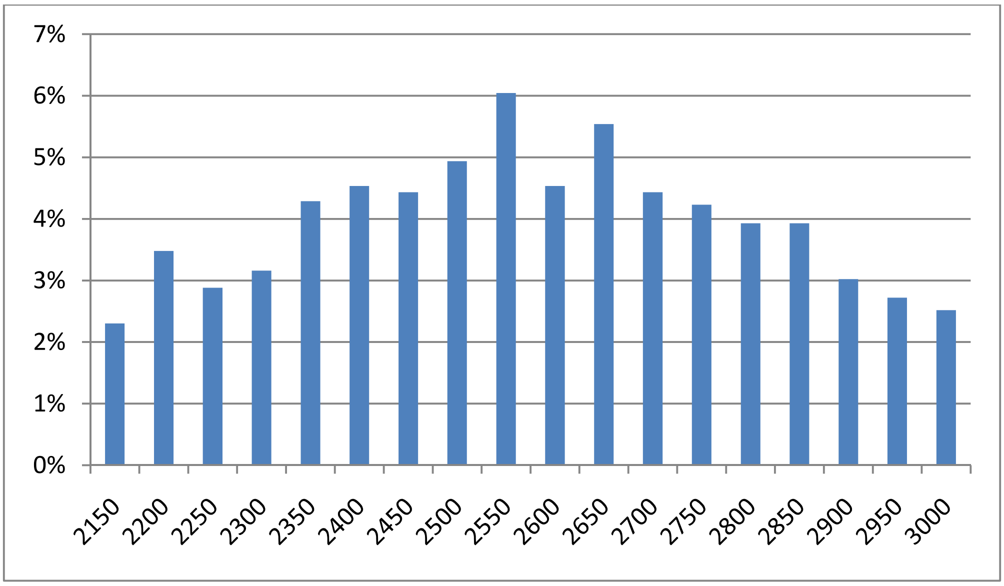

Though simple, the risk neutral histogram method has a number of notable weaknesses. One such weakness is that it requires large amounts of data. In order to obtain estimates for n state probabilities, $n + 2$ option prices are needed. Furthermore, all of the $n + 2$ option prices need to correspond to evenly spaced exercise prices, with the distance between adjacent exercise prices given by ΔK. In practice, this is a big limitation, because reliable price estimates for options are not necessarily evenly spaced (in the data section, we elaborate on the criteria used to determine what a “reliable” price

[15] It should be pointed out that replication is not necessary to obtain expression (18). Differentiating the call option price given in expression (21) twice with respect to the exercise price will obviously yield the same result, but it is harder to do and does not provide the same intuitive explanation as to why this result is to be expected.

estimate is in this context). Also, it is clear that this approximation will always result in a truncated distribution (i.e. $\int_{-\infty}^{\infty} q_t(S_T) dS_T < 1$), as options for very high and very low exercise prices are not traded. Additionally, Bahra (1997) points out that this procedure is highly sensitive to badly behaved call prices. Observed prices sometimes exhibit small but sudden changes in convexity across exercise prices, as well as small degrees of concavity in exercise price. These irregularities result in large variations in probabilities over adjacent exercise prices and negative probabilities respectively. Where bid-ask spreads are observed rather than actual traded prices, these irregularities can arise due to measurement errors arising from using mid prices. Problems of this kind are present in our data and sometimes lead to histograms looking precisely like explained above (see Figure 4 below). Hence, it is clear that more sophisticated methods for retrieving the RND function are needed. The proposed methods for estimating the implied RND function can be broken down into three main categories.[16]

Figure 4 – Risk neutral histogram for December 21 2007, three months

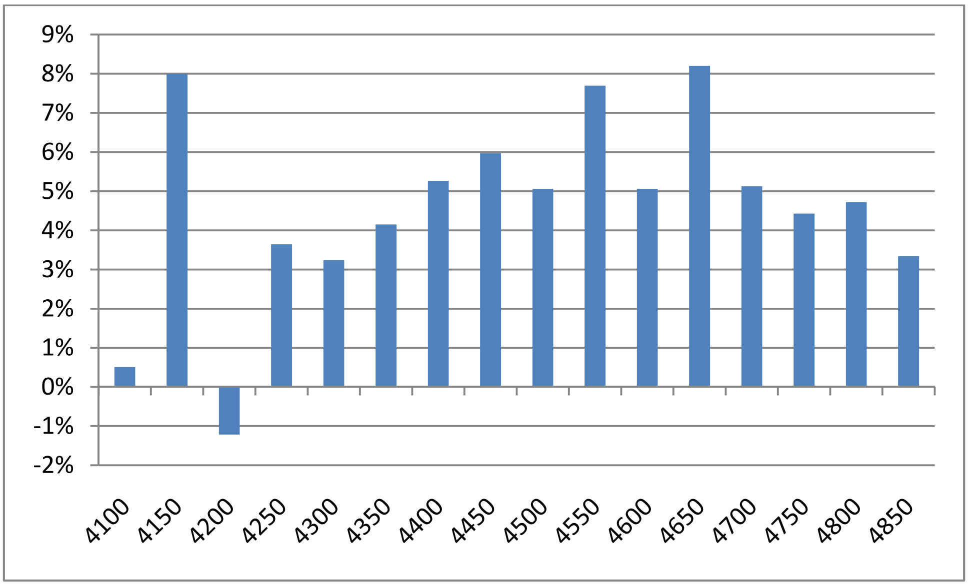

The first category is made up of strictly non-parametric methods. A notable example is Aït-Sahalia and Lo (1998), who apply the Nadaraya-Watson kernel estimator[17] to estimate the entire call pricing function. Strictly non-parametric methods have the advantage of not making any assumptions at all about the underlying RND function, thus allowing for more general RND functions. However, they are particularly data-intensive and thus require a large amount of available option prices to work well.

The second category encompasses curve-fitting methods. These are methods where the RND function is derived directly from some parametric specification of either the call pricing function or of the implied volatility smile curve. A notable example is Shimko (1993), who fits a quadratic polynomial

[16] A wide variety of different methods for estimating the implied RND function have been proposed. Here, we only intend to give a very brief overview. For a thorough review of the literature on the matter, we refer the interested reader to Jackwerth (1999) or Figlewski (2009).
[17] Going into the specifics of kernel regression is beyond the scope of this paper. The interested reader is referred to e.g. Härdle (1992).

to the implied volatility smile and then uses the Black-Scholes formula to obtain the call price as a continuous function of the exercise price.[18] The rationale behind interpolating in the implied volatility domain rather than in the call price domain directly is that implied volatilities are typically smoother than option prices themselves. The resulting call price function is then twice differentiated with respect to the exercise price in order to obtain the RND function between the lowest and the highest exercise prices. Clearly, the resulting distribution will be truncated. To cope with this, Shimko grafts lognormal tails onto each of the endpoints of the obtained density in order to get the resulting RND to integrate to one. Methods of this kind are non-parametric in the sense that the RND function is never explicitly parameterized, but they cannot be called strictly non-parametric, as they demand the estimation of certain parameters in the process of deriving the RND.

The third category comprises fully parametric methods, where assumptions are made about either the price process of the underlying asset or about the functional form of the RND directly. Examples include Bates (1991), Aparicio and Hodges (1998), Ritchey (1990) and Bahra (1997). Bates assumes that the price process of the underlying asset evolves according to an asymmetric jump-diffusion process and derives the RND based on this assumption. Aparicio and Hodges use the generalized beta distribution of the second kind, a four-parameter distributions first described by Bookstaber and McDonald (1987). The generalized beta distribution of the second kind encompasses many commonly used distributions, such as the lognormal distribution, the gamma distribution, the exponential distribution and several Burr type distributions (to mention a few) as special cases. The rationale for using such an advanced distribution is that one does not want to impose an overly restrictive functional form on the RND. Another way to achieve this is to use mixtures of simpler distributions. Richey proposes a method where the RND is expressed as a weighted sum of k lognormal distributions. Specified in this way, the RND is able to capture the main contributions to the implied volatility smile curve, namely the skewness and the kurtosis of the distribution of the underlying asset. The drawback of this method is that it requires the estimation of a large number of parameters as k increases. Two parameters are used for each lognormal distribution and $k-1$ mixing parameters are also needed. Hence, the total number of parameters to be estimated when k lognormal distributions are mixed is $3k-1$. However, Bahra finds that even when using $k=2$, the model is able to capture the skewness and the kurtosis of the underlying distribution, whilst only requiring five input parameters. Because of its flexibility and the relatively small number of required parameters to be estimated, Bahra finds the two-lognormal approach to be the preferred method to estimate the RND function. He also derives explicit formulas for European calls and puts for the two-lognormal method.

Interestingly, Jackwerth (1999) finds that unless there are very few available option prices, the various methods presented above tend to give rather similar estimates of the implied RND function. Hence,

[18] Note that the use of the Black-Scholes formula in this context does not require it to be true. It is merely used as a translation device between implied volatilities and option prices.

Jackwerth concludes that just about any reasonable method can be used without affecting the results too much. Consequently, we choose to use the two-lognormal method, as it is relatively simple, while allowing for a wide variety of possible RND shapes.

2.3.2. The two-lognormal method

When using a method where a functional form for the RND is assumed, the parameters are recovered by minimizing the distance between the observed option prices and those that are generated by the assumed parametric form. Melick and Thomas (1997) point out that this is a more general approach than assuming a stochastic process for the underlying price process, as a stochastic process implies a unique RND function, whereas any given RND function is consistent with many different stochastic processes.

A random variable is lognormal if its natural logarithm is normally distributed. Thus, if the random variable Z is normal with parameters μ and σ, e^Z is lognormal with parameters μ and σ, i.e. $Z \,\epsilon\, N(\mu, \sigma) \leftrightarrow e^Z \,\epsilon\, L(\mu, \sigma)$.[19] The probability density function for a lognormal random variable is given by:

$$\ell(x; \mu, \sigma) = \frac{e^{-\frac{(\log(x)-\mu)^2}{2\sigma^2}}}{x\sigma\sqrt{2\pi}}\ , \ \ x > 0 \tag{19}$$

Hence, if the RND is assumed to be a weighted sum of two lognormal random variables, it will have the following functional form:

$$q_t(S_T; \mu_1, \sigma_1, \mu_2, \sigma_2, \theta) = \theta\,\frac{e^{-\frac{(\log(S_T)-\mu_1)^2}{2\sigma_1^2}}}{S_T\sigma_1\sqrt{2\pi}} + (1-\theta)\,\frac{e^{-\frac{(\log(S_T)-\mu_2)^2}{2\sigma_2^2}}}{S_T\sigma_2\sqrt{2\pi}}\ , \ \ S_T > 0 \tag{20}$$

Since expression (20) above is a weighted sum, the weights for the respective lognormal densities must sum to one, i.e. $\theta \,\epsilon\, [0, 1]$.

Recall from expression (6) that the time t price of any contingent claim maturing at time T can be calculated as $\Pi_t = e^{-r(T-t)} \int_0^\infty q_t(S_T) h(S_T) dS_T$. Also recall that the payoff functions for European calls and puts respectively are $h(S_T) = \max(S_T - K, 0)$ and $h(S_T) = \max(K - S_T, 0)$. Thus, the price of a European call and a European put respectively can be computed as:

$$\begin{aligned} c_t &= \mathrm{e}^{-\mathrm{r(T-t)}} \int_K^\infty q_t(S_T)(S_T - K) dS_T \\ \mathrm{p_t} &= \mathrm{e}^{-\mathrm{r(T-t)}} \int_0^K q_t(S_T)(K - S_T) dS_T \end{aligned} \tag{21}$$

[19] Note that the parameters μ and σ here have nothing to do with the Black-Scholes parameters that were denoted in the same way earlier.

Given the functional form for the RND function $q_t(S_T)$ presented in expression (20) above, Bahra (1997) derives closed-form solutions for pricing European calls and puts:

$$c_t = \mathrm{e}^{-\mathrm{r}(\mathrm{T}-\mathrm{t})}\left(\theta\left(e^{\mu_1+\frac{1}{2}\sigma_1^2}\mathbf{\Phi}(d_1) - K\mathbf{\Phi}(\mathrm{d}_2)\right) + (1-\theta)\left(e^{\mu_2+\frac{1}{2}\sigma_2^2}\mathbf{\Phi}(d_3) - K\mathbf{\Phi}(\mathrm{d}_4)\right)\right)$$
$$\mathrm{p_t} = \mathrm{e}^{-\mathrm{r}(\mathrm{T}-\mathrm{t})}\left(\theta\left(K\mathbf{\Phi}(-\mathrm{d}_2) - e^{\mu_1+\frac{1}{2}\sigma_1^2}\mathbf{\Phi}(-d_1)\right) + (1-\theta)\left(K\mathbf{\Phi}(-\mathrm{d}_4) - e^{\mu_2+\frac{1}{2}\sigma_2^2}\mathbf{\Phi}(-d_3)\right)\right) \tag{22}$$

The parameters d_1, d_2, d_3 and d_4 are given by:

$$\begin{aligned} d_1 &= \frac{-\log(K) + \mu_1 + \sigma_1^2}{\sigma_1} \\ d_2 &= d_1 - \sigma_1 \\ d_3 &= \frac{-\log(K) + \mu_2 + \sigma_2^2}{\sigma_2} \\ d_4 &= d_3 - \sigma_2 \end{aligned} \tag{23}$$

It is interesting to note that unlike the Black-Scholes formula for index options, the dividend yield q is not explicitly considered in the closed-form solution presented above. The reason for this is that the derivation of the Black-Scholes formula starts by assuming a price process for the underlying asset under the real world probability measure $\mathbb{P}$ and then transforms it to the equivalent martingale measure $\mathbb{Q}$, whereas the approach taken here is to find the model parameters directly under $\mathbb{Q}$. Hence, the dividend yield does not need to be considered explicitly, as its presence will affect the values of the other parameters, thus giving it an implicit effect. Also note that the time to maturity, $T-t$, is not present, other than in the discount factor. This is because the parameters are estimated for a specific maturity, so it too will be implicitly included in them.

The expected value of a lognormally distributed random variable e^Z with parameters μ and σ is given by $\mathbb{E}[e^Z] = \mathrm{e}^{\mu+\frac{1}{2}\sigma^2}$. Thus, by the linearity of the expected value, it is clear that the time t expected value of the RND function under $\mathbb{Q}$ will be given by:

$$\mathbb{E}_t^{\mathbb{Q}}[S_T] = \theta e^{\mu_1+\frac{1}{2}\sigma_1^2} + (1-\theta)e^{\mu_2+\frac{1}{2}\sigma_2^2} \tag{24}$$

This expected value should equal the time t price of a futures contract maturing at time T, denoted by $F_{t,T}$. Hence, it should hold that:

$$F_{t,T} = \theta e^{\mu_1+\frac{1}{2}\sigma_1^2} + (1-\theta)e^{\mu_2+\frac{1}{2}\sigma_2^2} \tag{25}$$

Thus, in order to fit a two-lognormal RND function to the data, the task is to solve the following minimization problem, where observed call and put prices for an exercise price K are denoted by $\hat{c}_t(K)$ and $\hat{p}_t(K)$ respectively:

$$\min_{\mu_1,\sigma_1,\mu_2,\sigma_2,\theta} \sum_{i=1}^{n}\left(c_t(K_i) - \hat{c}_t(K_i)\right)^2 + \sum_{i=1}^{n}\left(p_t(K_i) - \hat{p}_t(K_i)\right)^2 + \left(\theta e^{\mu_1+\frac{1}{2}\sigma_1^2} + (1-\theta)e^{\mu_2+\frac{1}{2}\sigma_2^2} - F_{t,T}\right)^2$$

$$\begin{gathered}\text{subject to}\\ \mu_1,\mu_2 \,\epsilon\, \mathbb{R}\\ \sigma_1,\sigma_2 \geq 0\\ \theta \,\epsilon\, [0,1]\end{gathered} \tag{26}$$

The time to maturity for all options is obviously fixed to $T - t$, as the aim is to derive $q_t(S_T)$, i.e. the RND function for time T at time t. At this point, all the tools necessary to carry out our analysis have been presented. Before doing so, however, we will give an overview of the previous research conducted on the forecasting ability of RND functions, focusing on studies on market crashes, as well as present the data that has been used.

3. Previous research

The literature on implied RND functions is extensive. However, much of it focuses on exploring methods to extract RND functions from option prices and identifying the best ones. A brief overview of literature of this kind was presented in the previous section. The literature that focuses on using the RND to look at the market's probability beliefs about specific events is more sparse. Studies of this kind can be divided into two categories, namely those that look at planned events, such as elections or central bank meetings, and those that look at unplanned events, i.e. various crises. Here, we intend to give a summary of the research conducted in this area, focusing on research on unplanned events.

Äijö (2006) finds that "good" news cause implied volatility to decrease and make the RND function less left-skewed, while increasing its kurtosis. Conversely, "bad" news increase implied volatility, make the RND more left-skewed and decrease its kurtosis. These are general findings and should apply irrespective of whether planned or unplanned events are studied. Another general finding, presented by Ederington and Lee (1996), is that there is an inverse relationship between the time to maturity of the options studied and the effect of new information on the implied volatility. Thus, we would expect RND functions for shorter maturities to more accurately reflect the market's probability beliefs about an event.

Mandler (2002) studies RND functions around European Central Bank (ECB) meetings. He uses a curve-fitting method in the implied volatility domain to estimate the RND and finds that ECB meetings do not have a clear effect on the estimated implied RND functions. He concludes that ECB meetings have too small an impact on the market in order for their effect to rise above the "noise" present in the RND function. Thus, it is clear that in order for an event to have an effect on the RND, the event has to be of great importance to the market. In their broad study of the usefulness of implied RNDs, Gemmill and Saflekos (2000) look at (among other things) British elections. They extract the RNDs from FTSE options using the two-lognormal method and find that it does help to reveal market sentiment during elections, but that it lacks forecasting ability.

Obviously, stock market crashes are vastly significant, and hence, they are expected to be important enough to the market to affect the RND. Unlike the planned events, the time of occurrence of events of this kind is unknown ex ante. Therefore, events like these can potentially test the predictive power of RNDs to a greater degree than planned events, since it is possible to study whether RNDs before the event predicted its occurrence at all, rather than just its outcome. In the study already mentioned above, Gemmill and Saflekos also look at the effects of the crash of October 1987, the mini crash of October 1989 and the market turmoil of October 1997 on the British stock market. They find that the implied RND did not predict any of these events and conclude that the index options market reacts to rather than predicts crashes. Specifically, they find that the RND becomes more left-skewed after the event and not before it. Still, the authors point out that the RND is useful for revealing market sentiment after an event has occurred. Bates (1991) conducted one of the first studies on market crashes. He looks at RND functions implied by S&P 500 futures options for the period leading up to October 1987 and finds that the subsequent crash was anticipated as much as two months in advance. Fung (2007) looks at implied volatility on the Hong Kong stock exchange and finds that it gave early warning signs of the 1997 Hong Kong stock market crash. However, Bhabra et al. (2001) arrive at the opposite conclusion when studying the 1997 Korean financial crisis. They study the implied volatility of KOSPI200 index options and, much like Gemmill and Saflekos, conclude that option prices react to rather than predict crashes. Hence, it is obvious that the literature is not clear and points in different directions when it comes to the predictive power of implied volatilities and RNDs. However, Lynch and Panigirtzoglou (2008), who summarize the literature on the matter, find that the conclusion that option prices (and hence RND functions) react to rather than predict crashes is supported by most studies. Birru and Figlewski (2010) study the market crash of September 2008, i.e. the same crash that is studied in this paper. However, their methodology differs greatly from most other studies, as they look at intra-day changes of the RND implied by S&P 500 index options and the effects that news have on it rather than on inter-day RNDs. Hence, the question of prediction is barely touched upon, though the authors do find that the RND is highly responsive to changes in the level of the stock index, indicating that RNDs react to rather than predict movements in the underlying asset price.

4. Data

The bulk of the data used for the analysis consists of European options on the EURO STOXX 50 index for all of the trading days during the period December 1 2006 to December 31 2008. The time period is chosen so as to cover the entire period of the financial crisis that led up to the stock market crash, from the first indicators of it in early 2007 to the actual crash in September 2008. In addition, we include the end of 2006 so as not to miss the normal market conditions prior to the crash, as well as the end of 2008, when the crisis was in full force. Thus, the data at hand covers a time period of varying market conditions, making the chosen time period interesting to study. The data is also interesting because it consists of information on the Euro zone, whereas the other studies in this field have been focused on American (typically S&P 500), Asian or British data. The reason why the EURO STOXX 50 index specifically is chosen is that it is a very large index with a liquid derivatives market, which is essential to obtain reliable data.

The initial data set, obtained from iVolatility.com[20], consists of all quoted calls and puts during the mentioned period for a total of 490508 options, divided equally between puts and calls. For each option, the data gives information about maturity ($T - t$), exercise price (K), current index level (S_t), traded volume, open interest, and bid and ask quotes. We use the mid price, i.e. the simple average of the bid and ask for an option, as our option price estimate ($\hat{c}_t$ and $\hat{p}_t$ for calls and puts respectively). To this data set, we apply a cleaning procedure along the lines of Bakshi, Cao and Chen (1997).

In order to exclude observations that may distort the analysis, we apply a cleaning procedure consisting of nine filters. Specifically, we remove: options with no traded volume and/or open interest (1), options with less than six days to maturity (2), options with negative bid and/or ask (3), options where the bid price is greater than the ask price (4), options where bid and/or ask is greater than the current level of the index (5), options for which bid and/or ask $< \max(S_t - K, 0)$ (6), options where the ratio of ask price to bid price is greater than 1.2 (7), options with bid and/or ask smaller than 0.1 (8) and finally, options that are puts (9).

The reason for removing options with no traded volume and/or open interest (filter 1) is that these are options that are illiquid and hence, the information contained in their prices is unreliable. Options with less than six days to maturity (filter 2) are removed, since they may suffer from liquidity biases, caused by traders having to buy or sell large quantities to close out existing positions, as pointed out by Bakshi, Cao and Chen. The filters applied in steps 3 to 6 remove options that violate obvious no-arbitrage conditions, such as negative price, negative bid-ask spread and negative time value. The rationale behind filter 7 is that we want to use options with as narrow bid-ask spreads as possible so as to obtain reliable estimates of option prices. However, if the requirement on the ask to bid ratio being close to one is too strict, we are left with very few option prices, making further analysis difficult or

[20] www.ivolatility.com

even impossible. After having tried different values, we find that 1.2 is a satisfactory cutoff point. The reason for removing options with prices of less than ten cents (filter 8) is that these are options were price changes will always have a large percentage effect, as the minimal increment that a price can change by is one cent. In order to mitigate this effect of discrete prices, these options are removed. At this point, we are left with 71660 options, divided between 45114 calls (63%) and 26546 puts (37%). Thus, calls make up roughly two thirds of the option prices that we deem reliable. One possible approach at this point would have been to convert all puts to calls using the put-call parity and to use the average between the call mid price and the mid price implied by the put (i.e. the call price obtained after converting the mid put price into a call price by using the put-call parity) as our call price estimate. However, given that the remaining options have made it through a rigorous cleaning procedure, they should already give reliable price estimates. Hence, we feel that this procedure adds unnecessary complexity without significantly improving reliability. Moreover, we would not be able to do this for all options, as we have more calls than puts, thus leading to option prices being estimated in an inconsistent way. What could still be done, though, is to remove the puts that have corresponding calls (i.e. calls for the same exercise price and maturity), but to keep the unique puts so as to obtain option price estimates for a larger number of exercise prices. However, Birru and Figlewski (2010) point out that equity index puts typically trade at different implied volatilities than corresponding calls.[21] Thus, this approach will create artificial jumps in the implied volatility curve wherever a put price rather than a call price is used, which is precisely the result that we obtained when we tried this method (see Figure 5 below). Birru and Figlewski also point out that this is likely to result in badly behaved RND functions. Hence, we choose to exclude puts altogether in our final filter (9). Thus, our final data set consists of 45114 call options. A summary of the cleaning procedure with the number of options removed in each filter is presented in Appendix C.

Figure 5 – Volatility smile with five unique puts on December 18 2006, three months

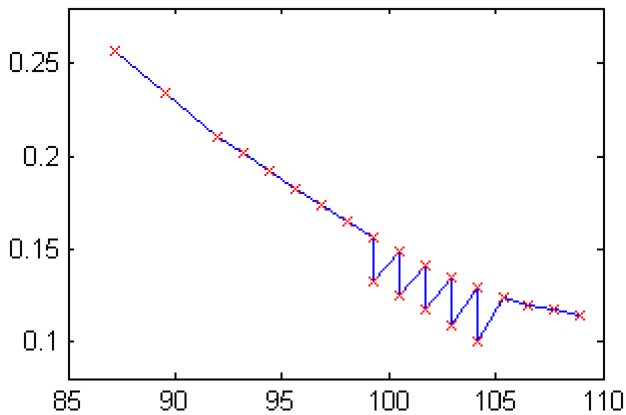

[21] In theory, where trading is assumed to be costless, the put-call parity implies that the implied volatilities for a put and a call for the same exercise price and time to maturity should be equal in order for there not to be any arbitrage. In practice, however, there is cost associated with putting on a trade, which is why these implied volatilities can differ. How much they can differ is still limited by arbitrage, and hence depends on the trading cost. Birru and Figlewski (2010) find that for S&P 500 index options, puts can trade at implied volatilities of one to two percentage points higher than calls at the money.

In addition to option price data, data on prices for futures on the EURO STOXX 50 index, the risk free interest rate and the dividend yield for the index is also needed. The EURIBOR rate is used as proxy for the risk free rate. For all of the trading days during the period December 1 2006 to December 31 2008, data for EURIBOR rates for three, six and nine months, as well as data for the dividend yield on the EURO STOXX 50 index, was obtained from Thomson Reuters Datastream. Futures prices for maturities of three, six and nine months for the last month of every quarter were obtained from Bloomberg Terminal. The reason why these particular dates and maturities were chosen will be elaborated on in the methodology section below.

5. Methodology

As was stated in the introduction, the aim of this paper is to study whether the options market predicted the stock market crash of September 2008 or reacted to it. The simplest way to do this is to look at the implied volatility before and after the crash. In our analysis, we look at the level of the at-the-money (ATM) implied volatility as well as the overall shape of the implied volatility curve. The next step is to retrieve the implied RND function and examine it. Birru and Figlewski (2010) point out that the RND has a significant advantage over the implied volatility in that it is model-independent[22], whereas the implied volatility is extracted by using the Black-Scholes model. Looking at the RND function allows us to study how the skewness and the kurtosis of the risk neutral distribution of the underlying asset changes over time, thus contributing information that is not explicitly present when looking only at the implied volatility.

For the last month of every quarter (i.e. March, June, September and December) of our sample period, we generate implied volatility curves for maturities roughly equal to three, six and nine months.[23] We choose to look at different maturities in order to see how far in advance, if at all, the options market was able to predict the stock market crash. The reason for going as far back in time as to the end of 2006 is twofold. First, we want a clear picture of what the implied volatility curves and RND functions looked like before the crash and second, early indicators of the stock market crash appeared as early as the first quarter of 2007, so the starting point had to be set prior to this. The reason for choosing the last month of every quarter is practical. The expiry dates of options on the EURO STOXX 50 index do not span all months. However, options expiring in March, June, September and December are always available. Thus, carrying out the analysis during these months ensures us of the availability of options with maturities of three, six and nine months. This is particularly useful, since the crash that we are studying occurred during the month of September.

[22] While the RND is model-dependent in the sense that it relies on the model chosen to retrieve it, recall that Jackwerth (1999) finds that different methods tend to give similar RND functions if there is enough data available.
[23] For each of these months, a date that gives roughly the desired maturities is chosen and all of the analysis is carried out on this date. In order to stress test this approach, we have varied this date back and forth by about a week. We find that this does not significantly alter the results.

The methodology described above differs from the one most commonly observed in studies of this kind. Typically, the evolution of the implied volatility and/or the RND function is studied every day around the specific event. Therefore, we also generate implied volatility curves and RNDs for every trading day of September 2008, looking at contracts with expiry in December 2008 (i.e. a time to maturity of roughly three months). This allows us to see if the options market was able to predict the crash on a horizon much shorter than three months.

5.1. Implied volatility

In order to calculate the implied volatilities, we use the Newton-Rhapson method, an iterative method for finding the roots of real-valued functions. For a real-valued function $f(x)$ that is reasonably well-behaved, a successively better approximation to the root, given an initial guess of x_0, is given by[24]:

$$x_{n+1} = x_n - \frac{f(x_n)}{f'(x_n)} \tag{27}$$

Notice that the change in the variable x between steps is given by $-\frac{f(x_n)}{f'(x_n)}$. Thus, as x_n approaches the exact root, the correction term $-\frac{f(x_n)}{f'(x_n)}$ tends to zero. Hence, an arbitrarily good solution can be obtained and the procedure is terminated when the correction term is smaller than a pre-specified value, commonly referred to as the tolerance.

In our case, the function whose root we want to find is the difference between the Black-Scholes call price and the observed call price, here denoted by $D = c_t - \hat{c}_t$. Given that all parameters needed to calculate the Black-Scholes price apart from σ are known, this can be expressed as (recall expressions (12) and (13) from the theoretical framework section):

$$\begin{gathered} D(\sigma) = c_t(\sigma) - \hat{c}_t = e^{-q(T-t)} S_t \boldsymbol{\Phi}(d_1) - e^{-r(T-t)} K \boldsymbol{\Phi}(d_2) - \hat{c}_t \\ d_1 = \frac{\log\left(\frac{S_t}{K}\right) + \left(r - q + \frac{\sigma^2}{2}\right)(T-t)}{\sigma\sqrt{T-t}} \\ d_2 = \frac{\log\left(\frac{S_t}{K}\right) + \left(r - q - \frac{\sigma^2}{2}\right)(T-t)}{\sigma\sqrt{T-t}} \end{gathered} \tag{28}$$

Thus, all we need in order to use the Newton-Rhapson method is $D'(\sigma) = \frac{\partial D(\sigma)}{\partial \sigma}$. Note that since $\hat{c}_t$ is a constant, $\frac{\partial D(\sigma)}{\partial \sigma} = \frac{\partial c_t(\sigma)}{\partial \sigma}$. This is the derivative of the Black-Scholes call option price with respect to the volatility, typically referred to as vega and denoted by ν:

$$\nu = S_t e^{-q(T-t)} \phi(d_1) \sqrt{T-t} \tag{29}$$

[24] Note that $f(x)$ in this setting is an arbitrary real-valued function and not in any way connected to the pricing function $f(t, S_t)$ that was introduced earlier.

In expression (29) above, d_1 is calculated in the same way as in expression (28) and $\phi(\cdot)$ denotes the probability density function of the standard normal distribution, elaborated on earlier. Thus, we obtain the implied volatility for an option by the iterative process:

$$\sigma_{n+1} = \sigma_n - \frac{D(\sigma_n)}{\nu(\sigma_n)} \tag{30}$$

We use the initial guess of $\sigma_0 = 20\%$ and a tolerance of 10^{-6}, i.e. the process is terminated when $\left| -\frac{D(\sigma_n)}{\nu(\sigma_n)} \right| \leq 10^{-6}$. The obtained implied volatilities are then plotted against the corresponding standardized exercise prices.[25] All of this is easily done in MATLAB, where vector notation allows us to obtain all implied volatilities for a specified maturity $T - t$ at a specific date t at the same time. The ATM implied volatility is obtained by linearly interpolating the implied volatilities for the exercise prices closest to the value of S_t. As the interpolation distance is usually short, a linear approximation should provide reasonable estimates, even though the implied volatility curve is typically convex, so linear interpolation could give rise to over-estimation.

5.2. RND

Even though risk neutral histograms are not used in the actual analysis, we present a few of them in the paper, as they provide an intuitive explanation of RNDs to the reader unfamiliar with the concept. When constructing the risk neutral histograms, we do not extrapolate outside of the available exercise price range and thus end up with truncated distributions. However, we do interpolate in order to obtain estimates for all the necessary option prices in the range between the lowest and the highest available exercise prices. We use linear interpolation in the implied volatility domain.[26] For every desired exercise price that we lack data for, we interpolate between the two implied volatilities corresponding to the two exercise prices closest to the desired exercise price and then convert the obtained implied volatilities to prices using the Black-Scholes formula. The histograms are then obtained by using expression (15).

For the two-lognormal method, we solve the minimization problem:

$$\begin{gathered}\min_{\mu_1,\sigma_1,\mu_2,\sigma_2,\theta} \sum_{i=1}^{n} \left(c_t(K_i) - \hat{c}_t(K_i)\right)^2 + \left(\theta e^{\mu_1 + \frac{1}{2}\sigma_1^2} + (1-\theta) e^{\mu_2 + \frac{1}{2}\sigma_2^2} - F_{t,T}\right)^2 \\ \text{subject to} \\ \mu_1, \mu_2 \in \mathbb{R} \\ \sigma_1, \sigma_2 \geq 0 \\ \theta \in [0,1]\end{gathered} \tag{31}$$

[25] The standardization is carried out by dividing all exercise prices K at time t by the level of the index at the time, S_t, and multiplying by 100. Hence, a standardized ATM exercise price is always equal to 100. The standardization is performed in order for the volatility smiles to be easily comparable over time as the level of the index changes.

[26] As has already been mentioned, the volatility smile is typically convex, thus making linear interpolation unsuitable. However, we usually interpolate over relatively short distances, where a linear approximation should not be too bad.

Obviously, this is a special case of expression (26) presented in the theoretical framework section. The only difference is that the sum of squared put option pricing errors is not present, as we do not have any put options in our final data set. The minimization is done in MATLAB by using the command fmincon, choosing the interior point optimization algorithm. The starting guesses that we use are $\mu_1 = \mu_2 = \log(F_{t,T})$, $\sigma_1 = \sigma_2 = \sigma_{ATM}\sqrt{T-t}$, $\theta = \frac{1}{2}$, where σ_{ATM} is the estimated ATM implied volatility.

Having estimated the RND, we proceed by calculating a number of descriptive statistics for it. Specifically, we look at expected value, standard deviation, annualized percentage standard deviation, skewness and kurtosis. All of these are calculated at time t under the equivalent martingale measure $\mathbb{Q}$. The expected value, also known as the mean and the first moment, is given by[27]:

$$\mathbb{E}_t^{\mathbb{Q}}[S_T] = \int_0^\infty S_T q_t(S_T) dS_T \tag{32}$$

As has already been explained, it should hold that $\mathbb{E}_t^{\mathbb{Q}}[S_T] = F_{t,T}$. Thus, comparing the obtained expected value of the RND to the corresponding futures price provides a rough indication of the goodness of fit.[28]

The second moment of a distribution is known as the variance. The square root of the variance is the standard deviation, which is a measure of dispersion around the mean. The standard deviation is calculated as:

$$\mathbb{D}_t^{\mathbb{Q}}[S_T] = \sqrt{\int_0^\infty \left(S_T - \mathbb{E}_t^{\mathbb{Q}}[S_T]\right)^2 q_t(S_T) dS_T} \tag{33}$$

It should be pointed out that using the unadjusted standard deviation as an estimate of the volatility is problematic, as it depends on both the level of the index and on the time to maturity. Thus, comparing standard deviations without first adjusting them can lead to incorrect conclusions. For this reason, we adjust the standard deviation and obtain a measure that we call annualized percentage standard deviation, calculated as:

$$\%\mathbb{D}_t^{\mathbb{Q}}[S_T] = \frac{\mathbb{D}_t^{\mathbb{Q}}[S_T]}{S_t}\sqrt{T-t} \tag{34}$$

[27] For all of the descriptive statistics presented here, the interval that the integral should be taken over is $(-\infty, \infty)$. As has already been mentioned, however, since $q_t(S_T) = 0 \; \forall \; S_T \in (-\infty, 0)$, integrating over $[0, \infty)$ will yield the same result.

[28] Obviously, there are more accurate goodness of fit measures, e.g. the root mean square error (RMSE). However, explicitly evaluating and analyzing the goodness of fit is not the focus of this is paper. Here, we find it sufficient to point out that the goodness of fit appears to be satisfactory for all of the obtained RNDs.

The next measure that we look at is skewness. Skewness is calculated as the standardized third moment and measures the asymmetry of a distribution:

$$\mathbb{S}_t^{\mathbb{Q}}[S_T] = \left(\frac{1}{\mathbb{D}_t^{\mathbb{Q}}[S_T]}\right)^3 \int_0^{\infty} \left(S_T - \mathbb{E}_t^{\mathbb{Q}}[S_T]\right)^3 q_t(S_T) dS_T \tag{35}$$

A skewness of zero corresponds to a symmetric distribution, e.g. the normal distribution.[29] Negative skewness means that there is more probability mass in the left tail than in the right tail. For this reason, a distribution with negative skewness is referred to as left-skewed. Similarly, a distribution with positive skewness has more probability mass in the right tail and is referred to as right-skewed. Notice that since skewness is a standardized moment, it is dimensionless and hence does not need to be adjusted for comparability.

Finally, we look at the kurtosis of the distribution, calculated as the standardized fourth moment.

$$\mathbb{K}_t^{\mathbb{Q}}[S_T] = \left(\frac{1}{\mathbb{D}_t^{\mathbb{Q}}[S_T]}\right)^4 \int_0^{\infty} \left(S_T - \mathbb{E}_t^{\mathbb{Q}}[S_T]\right)^4 q_t(S_T) dS_T \tag{36}$$

Kurtosis measures the "peakedness" of a distribution. A higher kurtosis means that the distribution has a higher peak around the mean and fatter tails. Distributions with a kurtosis equal to three are referred to as mesokurtic. The most common example of a mesokurtic distribution is the normal distribution. If the kurtosis is higher than three, the distribution is referred to as leptokurtic. Conversely, if it is lower than three, the distribution is referred to as platykurtic. Like skewness, kurtosis is a standardized moment and hence does not need to be adjusted for comparability.

All of the integrals presented above are calculated numerically in MATLAB using the trapz command, which approximates the integral by using the trapezoidal method. The reason for using numerical integration is that the analytical expressions for all moment but the first are rather cumbersome, since we are dealing with a weighted sum of two distributions. However, recall that the mean can be calculated analytically by the rather simple formula in expression (24). Comparing the analytically obtained mean to the one obtained through numerical integration, we find that the difference appears no earlier than in the sixth decimal place. Thus, for all practical purposes, numerical integration yields sufficiently accurate results.

6. Results

In this section, we present our results. We start by examining the implied volatilities, looking at the evolution of the ATM implied volatility over time as well as the evolution of the volatility smile as a whole. We then proceed by looking at the implied RND functions, analyzing their shape as well as

[29] It should be pointed out that a skewness of zero does not necessarily imply a symmetric distribution (though this is typically the case). However, the converse is always true – a symmetric distribution always has a skewness that is equal to zero.

various summary statistics. Throughout this section, we will only present figures as illustrative examples of our analysis. All of the figures for the volatility smiles and the RND functions are available in Appendix A and Appendix B respectively.

6.1. Implied volatility

We begin by analyzing the implied volatilities for the last month of every quarter in the period December 2006 to December 2008. We then move on to look at the implied volatilities for September 2008. One general observation that is largely independent of the time period studied or the maturity at hand is that the overall shape of the implied volatility function is convex and that we observe a volatility smirk rather than a volatility smile. This is consistent with the findings of Rubinstein (1994) discussed earlier, i.e. that equity indices exhibit reverse skew in their implied volatility functions. Thus, we conclude that the dynamics proposed by the Black-Scholes framework do not accurately describe the price process of the index at hand.

6.1.1. December 2006 to December 2008

6.1.1.1. ATM implied volatility

In order to get a rough impression as to whether implied volatilities predicted the crash of September 15 2008 or reacted to it, we look at the evolution of the ATM implied volatility for the maturities of three, six and nine months (see Figure 6 below). We can see that for all of the maturities, there is a slow rise in the ATM implied volatility leading up to September 2008. Naïvely, this could be interpreted as the options market having predicted the crash. However, this is a dangerous conclusion. For one thing, other events, e.g. the bankruptcy of several subprime lenders in March 2007 and the bailout of Bear Stearns in March 2008, which unsettled the market and potentially could have increased the volatility, occurred before the crash, so the observed increase could just be a reaction to them. Furthermore, the increase in ATM implied volatility after the crash (i.e. from September 2008 to December 2008) is of a far greater magnitude than the relatively modest increase of the period leading up to the crash, thus suggesting that the options market reacted to the crash rather than predicted it. Still, coming to any far-reaching conclusion based solely on quarterly ATM implied volatility is dangerous, as ATM implied volatility could potentially give an incomplete picture. For this reason, we proceed by looking at the evolution of the entire implied volatility curve for the period at hand.

Figure 6 – ATM implied volatility for the period December 2006 to December 2008

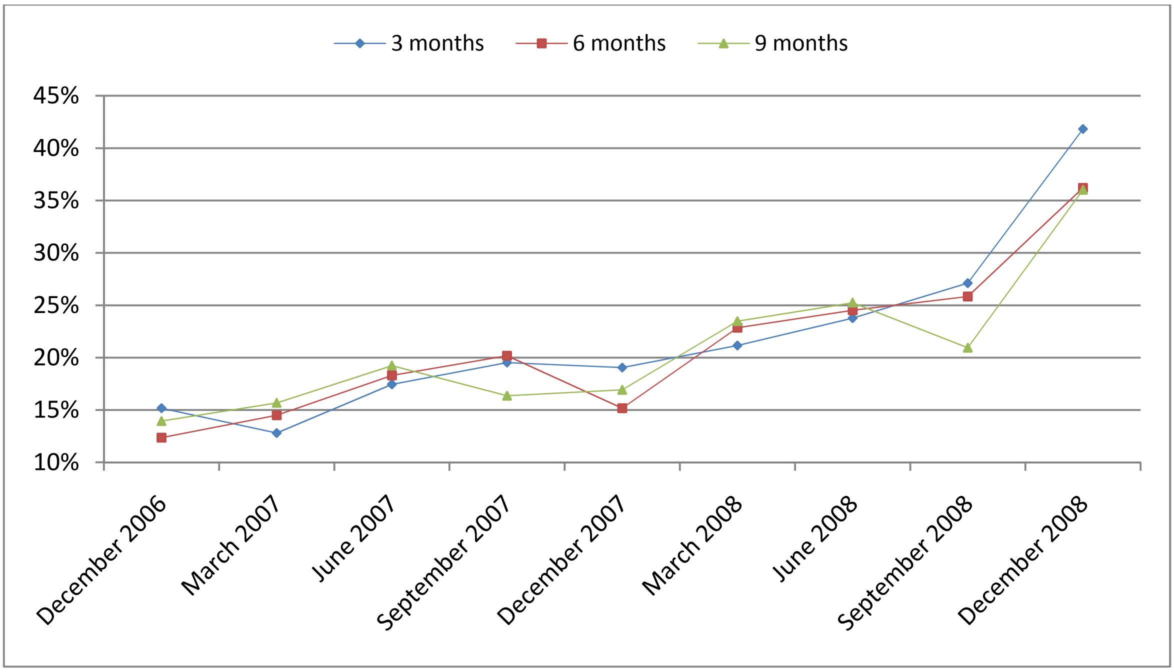

6.1.1.2. Volatility smile

When looking at the shape of the volatility smile and its evolution over time for the three month maturity, we find that it is rather steep for December 2006. However, the steepness declines and does not come back to a similar level until December 2008, i.e. after the crisis. Recall that the steepness for an implied volatility curve with a reverse skew shape is related to the left-skewness of the RND, where a steeper curve indicates more left-skewness. Thus, the decreased steepness indicates that the RND became less left-skewed during the financial crisis. A word of caution when looking at the steepness of a volatility smile is in order. One may think that the volatility smile for e.g. June 2007 is just as steep as the one for December 2006, as the two look similar at first glance. However, this is a result of there being more observations for options with low moneyness (here, moneyness is defined as the strike price divided by the time t value of the underlying asset, i.e. $\frac{K}{S_t}$) for June 2007 than for December 2006. As the volatility smile is the steepest for options with low moneyness, months where reliable data on options of this kind is available will look to have steeper volatility smiles at first glance. Thus, one needs to pay close attention to the moneyness when assessing the steepness of the implied volatility curve and how it changes over time. Still, interesting findings can be made when looking at the shape of the volatility smile. We note that the volatility smile has a very strange appearance in the month of March for both 2007 and 2008. The curve is almost flat and even slightly concave for low exercise prices. To see whether this was an anomaly for the specific dates, other dates in March 2007 and 2008 were tried without significantly affecting the results. Events that could be expected to affect the volatility smile occurred in both of these months, i.e. the bankruptcies of several subprime lenders in March 2007 and the bailout of Bear Stearns in March 2008, so the strange shape could have arisen due to these events. However, we do not see the same strangely flat shape for

September 2008, when the crash occurs, indicating that something other than negatively perceived events may have played a part in the strange shape of the March 2007 and March 2008 volatility smiles.

For maturities of six and nine months, the findings are rather similar to those for the three month maturity. However, there are some differences. In particular, there are more months where the implied volatility curve is rather flat and generally, the curve is less steep than for the three month maturity.

To summarize, we find that the implied volatility curve for the three month maturity is generally steeper than for six months and nine months. Also, there are fewer months where the curve is almost flat for no apparent reason. Thus, we feel that the three month maturity contains more relevant information than the longer maturities in this case. Additionally, as has already been mentioned, Ederington and Lee (1996) find that the effect of new information on the implied volatility curve is greater for shorter maturities, thus making them more appropriate to look at for our purposes. For this reason, we choose to use the three month maturity when looking at the volatility smile for every trading day of September 2008 and when retrieving the implied RND functions later on. Another reason for choosing three months is that there is generally more reliable data available for three months than for six and nine months in our data set.

6.1.2. September 2008

6.1.2.1. ATM implied volatility

Looking at the evolution of the three month ATM implied volatility throughout September 2008 (see Figure 7 below), we can see that it was relatively stable up until the crash of September 15. After the crash, the ATM implied volatility began to rise. However, the big increase in ATM implied volatility did not occur until the end of the month. This suggests that not only did the options market react to the crash rather than predict it, but the reaction was not immediate, at least not to its full extent. It is also noteworthy that there is actually a dip in ATM implied volatility during the week after the crisis, right before the sharp increase occurs, suggesting that investors were initially unsure of how to react to the crash. The probable reason for this is that the Emergency Economic Stabilization Act was presented on September 19, likely having a calming effect. However, the act was voted down on September 29, bringing the ATM implied volatility up to a higher level.[30]

[30] The Emergency Economic Stabilization Act was amended and passed into law under the name Troubled Asset Relief Program (TARP) on October 3 2008. However, this did little to soothe the market. For an in-depth discussion of the government policies implemented to handle the stock market crash, please refer to Taylor (2009).

Figure 7 – ATM implied volatility for every trading day of September 2008, three months

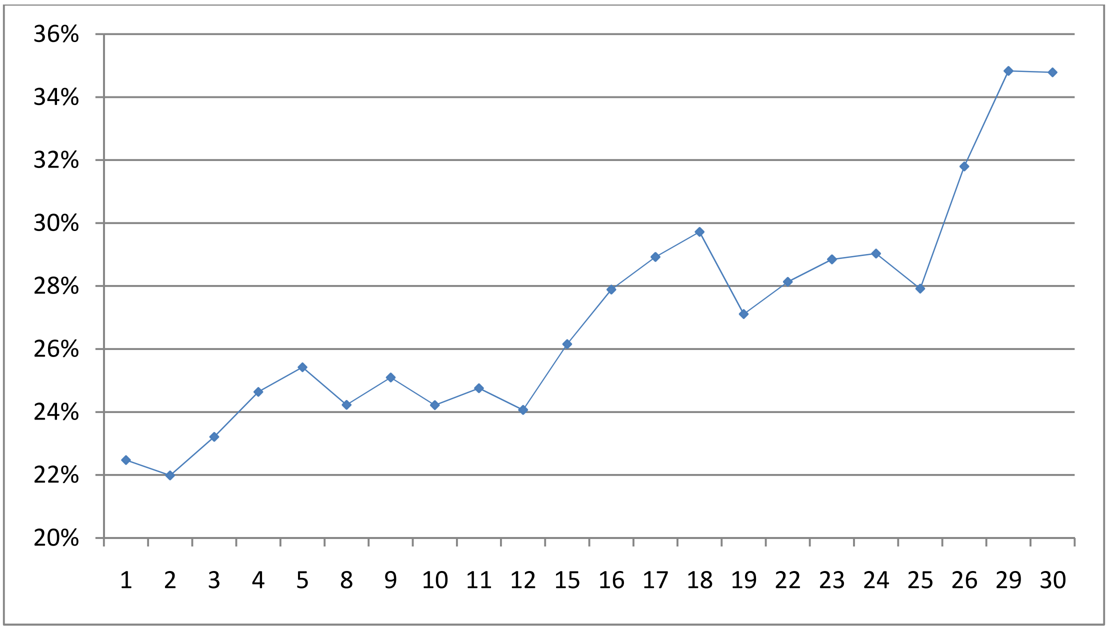

6.1.2.2. Volatility smile

When looking at how the shape of the volatility smile changes during September 2008, we find that while the slope does change somewhat throughout the month, it is very hard to find a general tendency. If anything, the volatility smile appears to be the steepest at the beginning of the month, though that is a finding that is far from certain, as the changes are just too small to be judged qualitatively, especially when having to consider that the moneyness of the available options changes from day to day.

While looking at the implied volatility is a good starting point, it should be clear that the analysis conducted above is nowhere near sufficient, as it is highly qualitative in nature and relies heavily on subjective interpretation of shapes. Though the approach of analyzing the implied volatility curve can be formalized, we choose to leave it here and instead look at the implied RND function, as it carries more information. Specifically, it is much easier to look at higher moments and thus quantify what we try to infer from the shape of the volatility smile.

6.2. RND

As was previously mentioned, a more convex volatility smile indicates fatter tails of the implied RND and a steeper volatility smile indicates a more left-skewed RND (for a volatility smile with a reverse skew). In the previous section, we tried to qualitatively look at steepness to say something about the skewness, but qualitatively looking at convexity to say something about the kurtosis is hard and was not even attempted. Here, we take a much more formal approach by actually calculating skewness and kurtosis, as well as the other previously explained descriptive statistics, for all of the retrieved RND functions. Just like for the implied volatilities, we first look at the last month of every quarter during the period December 2006 to December 2008 and then at every trading day of September 2008.

However, a general finding that does not depend on the time period examined is that all of the RNDs are left-skewed (to varying degree). This clearly contradicts the Black-Scholes assumption of GBM dynamics for the underlying price process, as GBM dynamics are consistent with a (single) lognormal RND, which is always right-skewed.[31] Another general finding is that the estimated distributions always seem to provide a good fit to the data, as $\left|\mathbb{E}_t^{\mathbb{Q}}[S_T] - F_{t,T}\right|$ is always very low (see Table 1 and Table 2 in the following subsections). It never exceeds 1.3 and is typically smaller than 0.1. As has been said, this is a rather rough estimate of the goodness of fit, but we deem it sufficient in this case, as it at least provides an indication.

6.2.1. December 2006 to December 2008

Table 1 – Descriptive statistics for the period December 2006 to December 2008

Descriptive statistics (and futures prices)						
Date	F	$\mathbb{E}$	$\mathbb{D}$	$\%\mathbb{D}$	$\mathbb{S}$	$\mathbb{K}$
2006-12-18	4156	4156.71	299.67	3.58%	-1.05	4.61
2007-03-16	3914	3914.00	362.65	4.57%	-0.65	3.03
2007-06-22	4510	4510.61	391.34	4.39%	-0.82	3.94
2007-09-21	4408	4408.00	417.53	4.80%	-0.85	3.58
2007-12-21	4424	4424.26	404.89	4.64%	-0.79	3.55
2008-03-20	3428	3427.99	485.22	6.99%	-0.41	2.56
2008-06-20	3453	3453.06	394.08	5.77%	-0.65	3.09
2008-09-19	3273	3272.90	429.08	6.62%	-0.47	3.00
2008-12-22	2359	2358.99	547.47	11.27%	-0.77	2.94

We begin by looking at the annualized percentage standard deviation for the time period at hand (refer to Table 1 above for the values of all of the descriptive statistics for the period), letting it serve as a proxy for volatility in this context. We note an upward jump in the level of this statistic when going from December 2006 to March 2007. After this, it stays relatively stable until March 2008, when we observe an increase again. This is followed by a slight decrease in June 2008, though the level is still higher than for the period before March 2008. In September 2008, we again see an increase, followed by a very large increase in December 2008. We note that the resulting volatility is more than three times higher at the end of the time period than at the beginning of it. When trying to make sense of these findings, it is useful to recall the events that occurred during these months. March of 2007 is the first month we look at after the first indicators of the forthcoming crash, i.e. the bankruptcy of several subprime lenders. This appears to have had an increase on the volatility. It is then stable until March 2008, when another early indicator of the crash – the bailout and the resulting acquisition of Bear Stearns by J.P. Morgan – occurred, again increasing the volatility. The volatility then decreases, before increasing again in September 2009, i.e. after the stock market crash. An unprecedented increase in

[31] The reason why implied distributions tend to be left-skewed is an interesting topic in its own right. One plausible explanation is portfolio-insuring behavior. For more on this topic, see Grossman and Zhou (1996).

volatility then occurs in December 2008. Thus, it appears that the volatility of the RND function reacts to events that the market perceives negatively by increasing. It does not seem to predict these events, which is particularly clear when noting the decrease in volatility in June 2008, the last point in time prior to the stock market crash considered in this section. This decrease is best explained as the market perceiving June 2008 to be a relatively better month than March 2008. However, if the options market had predicted the stock market crash, we would not expect the volatility to decrease in June 2008, as it should have reflected the expected increase in volatility of September 2008. Furthermore, the large increase in volatility in December 2008, i.e. after the stock market crash occurred, indicates that the full extent of the implications of the crash was not realized immediately, which is also more consistent with reaction than with prediction. Thus, when looking at the volatility of the RND function, we find that it indicates that the market reacted to rather than predicted the crash. One should note that there is a slight upward trend in volatility leading up to the crash, but as we have argued, this appears to be a result of the options market reacting to unsettling events prior to the crash rather than predicting the crash itself.

We next look at the skewness of the implied RND function. The distribution is rather left-skewed in December 2006, with a skewness of below -1. However, it becomes less left-skewed in March 2007. The skewness then stabilizes on a level somewhere between the December 2006 and the March 2007 levels from June 2007 to December 2007, only to become significantly less left-skewed in March 2008. The left-skewness then increases again in June 2008, but decreases to the March 2008 level following the crash of September 2008. The left-skewness then increases in December 2008, though the distribution is still less left-skewed than at the beginning of the period. The evolution of the skewness of the distribution is consistent with the results of the implied volatility section, where we found that the volatility smile for a time to maturity of three months is initially rather steep and that the steepness then declines somewhat, increasing again for December 2008. Notice that left-skewness tends to decrease during months where events perceived negatively by the market occur (i.e. March 2007, March 2008 and September 2008). For the relatively calm period from June 2007 to December 2007, the skewness changes very little after becoming more left-skewed in June 2007 when compared to March 2007. This suggests that negatively perceived events are associated with a decrease in left-skewness. Furthermore, just like the annualized percentage standard deviation, the evolution of the skewness seems to indicate that the options market reacted to rather than predicted the crash. Note that the left-skewness in March 2008 was very similar to the left-skewness in September 2008. Thus, if the stock market crash was predicted by the options market, we would expect left-skewness not to change much in June 2008 and stay close to the September 2008 level. However, it increases, indicating that the market seemed to believe in a return to more normal market conditions. An interesting finding is that the distribution becomes more left-skewed in December 2008 when compared to September 2008. Thus, in terms of skewness, the market actually resembles the normal market conditions more after the crash than during the preceding financial turmoil (however, the overall shape of the RND is very different for December 2008 when compared to the pre-crisis RND; this will be elaborated on below).

The final descriptive statistic that we examine in this section is the kurtosis. A general observation is that the kurtosis decreases over time and is at a much lower level at the end of the period (December 2008) than at the beginning of it (December 2006). Looking at the evolution of the kurtosis, we see large dips for March 2007 and March 2008. After both of these months (i.e. for June 2007 and June 2008 respectively), we can see that the kurtosis increases again, though the increase is smaller in magnitude than the preceding decrease. Strangely, however, we do not see a dip for September 2008. Instead, the kurtosis is rather stable at values around three for the period June 2008 to December 2008. Still, the general finding is that negatively perceived events are associated with decreased kurtosis. However, one should note that kurtosis is sensitive to observations that are far away from the mean, i.e. the tails of the distribution. Most reliable data is for options with moneyness ranging from 90% to 110%, meaning that we generally have few observations in the tails of the distribution, and hence, the uncertainty concerning the estimated tails is rather large. For this reason, the obtained kurtosis should be interpreted with care.

Our findings are largely consistent with those of previous studies. Like e.g. Äijö (2008) and Gemmill and Saflekos (2000), we find that negatively perceived events ("bad" news, to use Äijö's terminology) are associated with increased volatility and decreased kurtosis. We also find that changes in the RND function tend to occur after an event rather than before it. Hence, we find that the RND function did not appear to predict the stock market crash of September 2008, consistent with the finding of Gemmill and Saflekos that RND functions do not predict market crashes. Another finding consistent with Gemmill and Saflekos is that the RND takes on a more bimodal shape after the crash. This can clearly be seen when comparing the RND for any of the months before the crash to the RND of December 2008 (see Figure 8 below). While the changes in the shape of the RND functions that we observe prior to December 2008 are captured by the changes in standard deviation, skewness and kurtosis, December 2008 exhibits a more bimodal shape, indicating that the market experienced difficulty in reaching a new consensus. This cannot be seen by simply looking at the descriptive statistics.

Figure 8 – RND functions for September 12 2008 and December 22 2008, three months

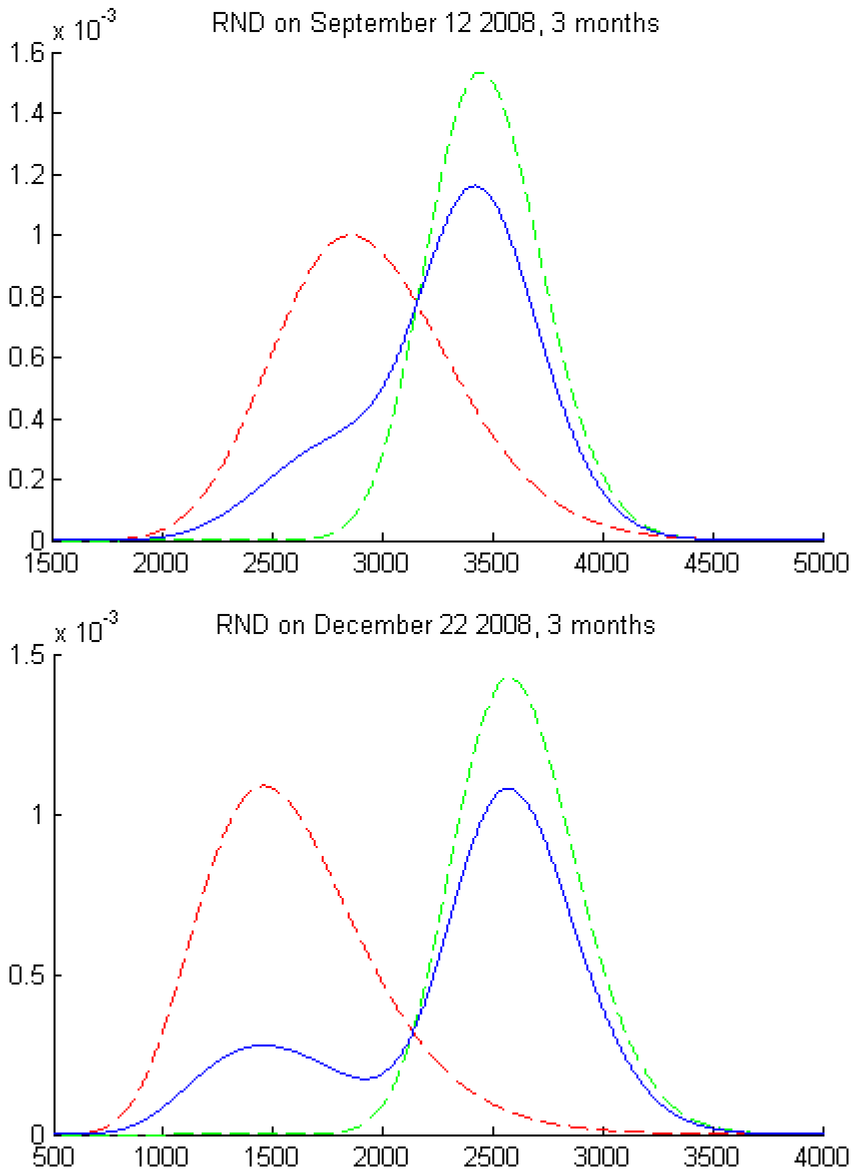

While those of our findings that have been discussed so far are consistent with previous research, our results for the evolution of the skewness are not. Both Äijö and Gemmill and Saflekos find that negatively perceived events (such as crashes) are associated with increased left-skewness. However, we observe a decrease in left-skewness as a reaction not only to the crash itself, but to the other "bad news" occurring before the crash as well. Interestingly, Birru and Figlewski (2010) come to the very same conclusion when looking at the implied RND function for the S&P 500 equity index during the crash of September 2008. Like us, they find that volatility increases, that the kurtosis decreases and that the distribution becomes less left-skewed. As was just stated, the first two findings are consistent with the studies on previous stock market crashes, whereas the last one appears to be a finding specific for this crisis. As the decrease in left-skewness appears to be a unique finding for this particular stock market crash, it is interesting to note that we make it on European data, whereas Birru and Figlewski made it on US data.

6.2.2. September 2008

Table 2 – Descriptive statistics for every trading day of September 2008

Descriptive statistics (and futures prices)						
Date	F	$\mathbb{E}$	$\mathbb{D}$	$\%\mathbb{D}$	$\mathbb{S}$	$\mathbb{K}$
2008-09-01	3389	3389.18	404.80	6.61%	-0.59	3.26
2008-09-02	3440	3440.04	400.68	6.41%	-0.56	3.17
2008-09-03	3398	3398.00	403.45	6.51%	-0.48	3.04
2008-09-04	3300	3300.02	420.69	6.96%	-0.47	2.95
2008-09-05	3203	3203.05	427.20	7.23%	-0.44	2.86
2008-09-08	3301	3300.99	415.72	6.72%	-0.50	2.98
2008-09-09	3284	3283.99	416.32	6.75%	-0.47	2.96
2008-09-10	3256	3256.85	413.12	6.70%	-0.52	3.14
2008-09-11	3236	3236.47	417.83	6.79%	-0.52	3.12
2008-09-12	3295	3296.29	405.63	6.44%	-0.51	3.16
2008-09-15	3160	3159.99	424.31	6.91%	-0.44	2.88
2008-09-16	3099	3099.00	431.78	7.13%	-0.34	2.63
2008-09-17	3020	3019.63	450.18	7.57%	-0.31	2.86
2008-09-18	3011	3011.00	446.99	7.52%	-0.43	2.79
2008-09-19	3273	3272.90	429.08	6.62%	-0.47	3.00
2008-09-22	3207	3206.99	418.53	6.50%	-0.43	2.89
2008-09-23	3162	3162.00	426.52	6.67%	-0.46	2.88
2008-09-24	3144	3143.98	420.30	6.58%	-0.43	2.88
2008-09-25	3227	3227.00	424.94	6.43%	-0.49	3.02
2008-09-26	3183	3183.00	451.38	6.90%	-0.32	2.60
2008-09-29	3036	3036.25	470.21	7.41%	-0.41	2.87
2008-09-30	3062	3062.01	477.85	7.41%	-0.40	2.84

When looking at the obtained RND functions for every trading day of September 2008, we note that there is a trend for left-skewness as well as kurtosis to decrease throughout the month (refer to Table 2 above for the values for all of the descriptive statistics for this period). However, the decrease is not smooth for either of these two statistics. For this reason, it is difficult to come to any conclusions based on the day-to-day changes in them. However, it should be noted that they both stay within a fairly limited interval throughout the month. More can be said about the annualized percentage standard deviation. We see that it increases following the crash of September 15 and stays at the higher level until September 19, when it decreases again. It then stays at this level until September 29, when it again comes back to the level of the first few days following the crisis. This evolution of the volatility is consistent with the market reacting to the developments concerning the Emergency Economic Stabilization Act. When it is announced on September 19, it appears to have a calming effect on the market, decreasing the volatility. However, when it is voted down on September 29, the volatility jumped back up. This result is consistent with what we found when looking at the development of the ATM implied volatility earlier and is more consistent with the options market reacting to the crash rather than predicting it. The finding that the impact on volatility is the clearest immediate effect of a crash is consistent with Gemmill and Saflekos (2000). It should also be noted that while the volatility for the days following the crash in September 2008 is higher than for any of the earlier dates that we look at, it is still much lower than in December 2008, indicating that the initial reaction to the crash did not take the full extent of it into account. Thus, the reaction was, to some extent, delayed, again clearly indicating that the options market reacted to rather than predicted the crash.[32]

7. Conclusions

In this paper, we have tried to determine whether the options market predicted or reacted to the stock market crash of September 2008. To do this, we have used option data for the EURO STOXX 50 equity index and looked at implied volatilities and RND functions. The obtained volatility smiles exhibit a reverse skew and hence clearly show that the functional form for the RND implied by the Black-Scholes dynamics does not adequately reflect the functional form of the actual RND implied by the market. We estimate the market RND by using a mixture of two lognormal distributions, as it is sufficiently flexible to capture features of the true RND function that are missed by the Black-Scholes single lognormal RND, such as fat tails and left-skewness.

Our general finding is that the options market did not predict the crash. Instead, we find that it reacted to it. This result is consistent with most other RND studies conducted on stock market crashes. Specifically, we find that the reaction is characterized by increased volatility, decreased left-skewness, decreased kurtosis and a tendency toward a bimodal shape of the RND function. Apart from the result regarding the skewness, these findings are consistent with studies conducted on earlier market crashes, such as the stock market crash of 1987 and the 1997 Asian stock market crash. However, studies on previous crashes find that the market's reaction in terms of skewness is that left-skewness increases. As was just stated, we make the opposite finding for the September 2008 stock market crash. Since Birru and Figlewski (2010) also observe that the September 2008 stock market crash lead to decreased left-skewness when looking at US data, we believe that this is a finding specific to this particular stock market crash, rather than an anomaly in our data set. We cannot think of any convincing explanation as to why the effect on skewness for this crash should be different from previous crashes, and Birru and Figlewski do not offer one either. For this reason, we believe that this is a potentially interesting area for further research.

[32] It should be pointed out that while the reaction was delayed, we find it highly unlikely that it was delayed for as long as until December 2008. However, as we do not look at October or November, we do not see the full extent of the reaction until December in our data.

When interpreting our results, it is important to note that the RND framework has certain limitations. The most obvious limitation is related to the availability of the data. Since deep OTM and deep ITM options are generally less liquid, reliable price estimates for them are often unavailable. For this reason, the estimation of the tails of the RND is somewhat unreliable. This is problematic, since the tails are likely to be best suited to gauge market expectations of extreme swings. Thus, to see whether the options market predicted a crash, we would ideally like to have a more accurate picture of the left tail of the RND. Another limitation of our study is that we do not consider the potential effects of estimation errors in our analysis. Specifically, we only look at how certain descriptive statistics change over time and do not test to see if any of the changes are statistically significant. However, it should be pointed out that the other studies in the field are conducted in a similar way and no study (at least that we know of) provides tests of statistical significance, which is criticized by Jackwerth (1999).

It is important to remember that the implied density functions that we look at are risk neutral. This means that they are likely to differ from the real world density functions, as investors are risk averse. Thus, option prices will incorporate not only probability beliefs about future outcomes, but also risk aversion and separating these two factors is far from trivial. This should not constitute a problem if risk aversion is relatively stable over time, as changes in the RND will mainly reflect changes in probability beliefs in this case. However, it is not unfeasible that the risk aversion could increase significantly during a crisis, which may or may not invalidate the RND framework for our type of analysis. Rubinstein (1994) finds that as long as the risk aversion stays within reasonable bounds, the shape of the real world density function is qualitatively quite similar to the RND. Still, we feel that more research of this kind is warranted. Hence, we would suggest the connection between the RND and the real world probability density during the late-2000s financial crisis as a topic for future research.

It should be noted that the fact that implied RNDs appear to lack predictive power does not mean that they are of no practical use. Since the RND function reacts to events such as crashes, it is a useful tool for assessing the market sentiment and how it changes over time. This could be of use for a wide variety of market participants, particularly central banks and other decision-making organs, as they need to have a good understanding of the market sentiment in order to develop effective policies and implement them in an efficient way.

8. References

Äijö, Janne, 2008, Impact of U.S. and UK macroeconomic news announcements on the return distribution implied by FTSE-100 index options, *International Review of Financial Analysis* 17, 242-258.

Aït-Sahalia, Yacine and Andrew W. Lo, 1998, Nonparametric Estimation of State-Price Densities Implicit in Financial Asset Prices, *Journal of Finance* 53, 499-547.

Aparicio, Silio David and Stewart Hodges, 1998, Implied Risk-Neutral Distribution: A Comparison of Estimation Methods, Working paper, Warwick University.

Arrow, Kenneth J., 1964, The Role of Securities in the Optimal Allocation of Risk-Bearing, *Review of Economic Studies* 31, 91-96.

Arrow, Kenneth J. and Gérard Debreu, 1954, Existence of an equilibrium for a competitive economy, *Econometrica* 22, 265-290.

Authers, John, 2010, *The Fearful Rise of Markets: Global Bubbles, Synchronized Meltdowns, and How To Prevent Them in the Future*, FT Press, Upper Saddle River, New Jersey.

Bahra, Bhupinder, 1997, Implied risk-neutral probability density functions from option prices: theory and application, Working paper No. 66, Bank of England.

Bakshi, Gurdip, Charles Cao and Zhiwu Chen, 1997, Empirical Performance of Alternative Option Pricing Models, *Journal of Finance* 52, 2003-2049.

Bates, David S., 1991, The Crash of ’87: Was it Expected? The Evidence from Options Markets, *Journal of Finance* 46, 1009-1044.

Bhabra, Ghurmeet S., Liliana Gonzalez, Myeong Sup Kim and John G. Powell, 2001. Volatility Prediction During Prolonged Crises: Evidence from Korean Index Options, *Pacific-Basin Finance Journal* 9, 147-164.

Birru, Justin and Stephen Figlewski, 2010, Anatomy of a Meltdown: The Risk Neutral Density for the S&P 500 in the Fall of 2008, Working paper, New York University.

Björk, Tomas, 2004, *Arbitrage Theory in Continuous Time*, 2nd Edition, Oxford University Press, Oxford.

Black, Fischer, 1976, Studies of Stock Price Volatility Changes, *Proceedings of the 1976 Meetings of American Statistical Association, Business and Economic Statistics Section*, 177-181.

Black, Fischer and Myron Scholes, 1973, The Pricing of Options and Corporate Liabilities, *Journal of Political Economy* 81, 637-659.

Bookstaber, Richard M. and James B. McDonald, 1987, A General Distribuion for Describing Security Price Returns, *Journal of Business* 60, 401-424.

Breeden, Douglas T. and Robert H. Litzenberger, 1978, Prices of state-contingent claims implicit in option prices, *Journal of Business* 51, 621-651.

Cox, John C. and Stephen A. Ross, 1976, The Valuation of Options for Alternative Stochastic Processes, *Journal of Financial Economics* 3, 145-166.

Ederington, Louis H. and Jae Ha Lee, 1996, The creation and resolution of market uncertainty: The impact of information releases on implied volatility, *Journal of Financial and Quantitative Analysis* 31, 513-539.

Figlewski, Stephen, 2009, Estimating the Implied Risk Neutral Density for the U.S.Market Portfolio, in *Volatility and Time Series Econometrics: Essays in Honor of Robert F. Engle*. Tim Bollerslev, Jeffrey R. Russell and Mark Watson, ed. Oxford University Press, Oxford, 323-353.

Figlewski, Stephen and Xiaozu Wang, 2000, Is the "Leverage Effect" a Leverage Effect?, Working paper, NYU Stern School of Business.

Fung, Joseph K. W., 2007, The Information Content of Option Implied Volatility Surrounding the 1997 Hong Kong Stock Market Crash, *The Journal of Futures Markets* 27, 555-574.

Garman, Mark B. and Steven W. Kolhagen, 1983, Foreign Currency Option Values, *Journal of International Money and Finance* 2, 231-237.

Gemmill, Gordon and Apostolos Saflekos, 2000, How Useful are Implied Distributions? Evidence from Stock-Index Options, *Journal of Derivatives* 7, 83-98.

Grossman, Sanford J. and Zhongquan Zhou, 1996, Equilibrium Analysis of Portfolio Insurance, *Journal of Finance* 51, 1379-1403.

Härdle, Wolfgang. 1992, *Applied Nonparametric Regression*, Econometric Society Monograph No. 19, Cambridge University Press, Cambridge.

Harrison, Michael J. and David M. Kreps, 1979, Martingales and Arbitrage in Multiperiod Securities Markets, Journal of Economic Theory 20, 381-408.

Hull, John C., 2006, *Options, Futures and Other Derivatives*, 6th Edition, Pearson Prentice Hall, Upper Saddle River, New Jersey.

Jackwerth, Jens Carsten, 1999, Option Implied Risk-Neutral Distributions and Implied Binomial Trees: A Literature Review, *Journal of Derivatives* 7, 66-82.

Kijima, Masaaki, 2003, *Stochastic Processes with Applications to Finance* , Chapman & Hall/CRC, London.

Lynch, Damien and Nikolaos Panigirtzoglou, 2008, Summary Statistics of Option-Implied Probability Density Functions and their Properties, Working paper No. 345, Bank of England.

Mandler, Martin, 2002, Comparing risk-neutral probability density functions implied by option prices – market uncertainty and ECB-council meetings, Working paper, Universität Giessen.

Melick, William R., and Charles P. Thomas, 1997, Recovering an Asset's Implied PDF from Option Prices: An Application to Crude Oil During the Gulf Crisis, *Journal of Financial and Quantitative Analysis* 32, 91-115.

Ritchey, Robert J., 1990, Call Option Valuation for Discrete Normal Mixtures, *Journal of Financial Research* 13, 285-295

Rubinstein, Mark, 1994, Implied Volatility Trees, *Journal of Finance* 49, 771-818.

Shimko, David C., 1993, Bounds on Probability, *RISK* 6, 33-37.

Taylor, John B., 2009, The Financial Crisis and the Policy Responses: An Empirical Analysis of What Went Wrong, *Critical Review* 21, 341-364.

Appendix A – Volatility smiles

Below, we present the implied volatility smiles for the last month of every quarter during the period December 2006 to December 2008. The volatilities for maturities of three, six and nine months are presented in Figure A1, Figure A2 and Figure A3 respectively. For December 2008, the implied volatility for nine months is calculated one day later than for three and six months. This is a consequence of insufficient data for the nine month maturity on the selected date. We also present the implied volatility for every trading day of September 2008 for contracts maturing on December 20 2008, i.e. a maturity of roughly three months (though the exact maturity obviously gets shorter over time). These are presented in Figure A4. For all of the plots, the red crosses mark the obtained implied volatilities. These are connected linearly by solid blue line segments in order to get an approximate picture of the volatility smile.

Figure A1 – Volatility smiles for the period December 2006 to December 2008, three months

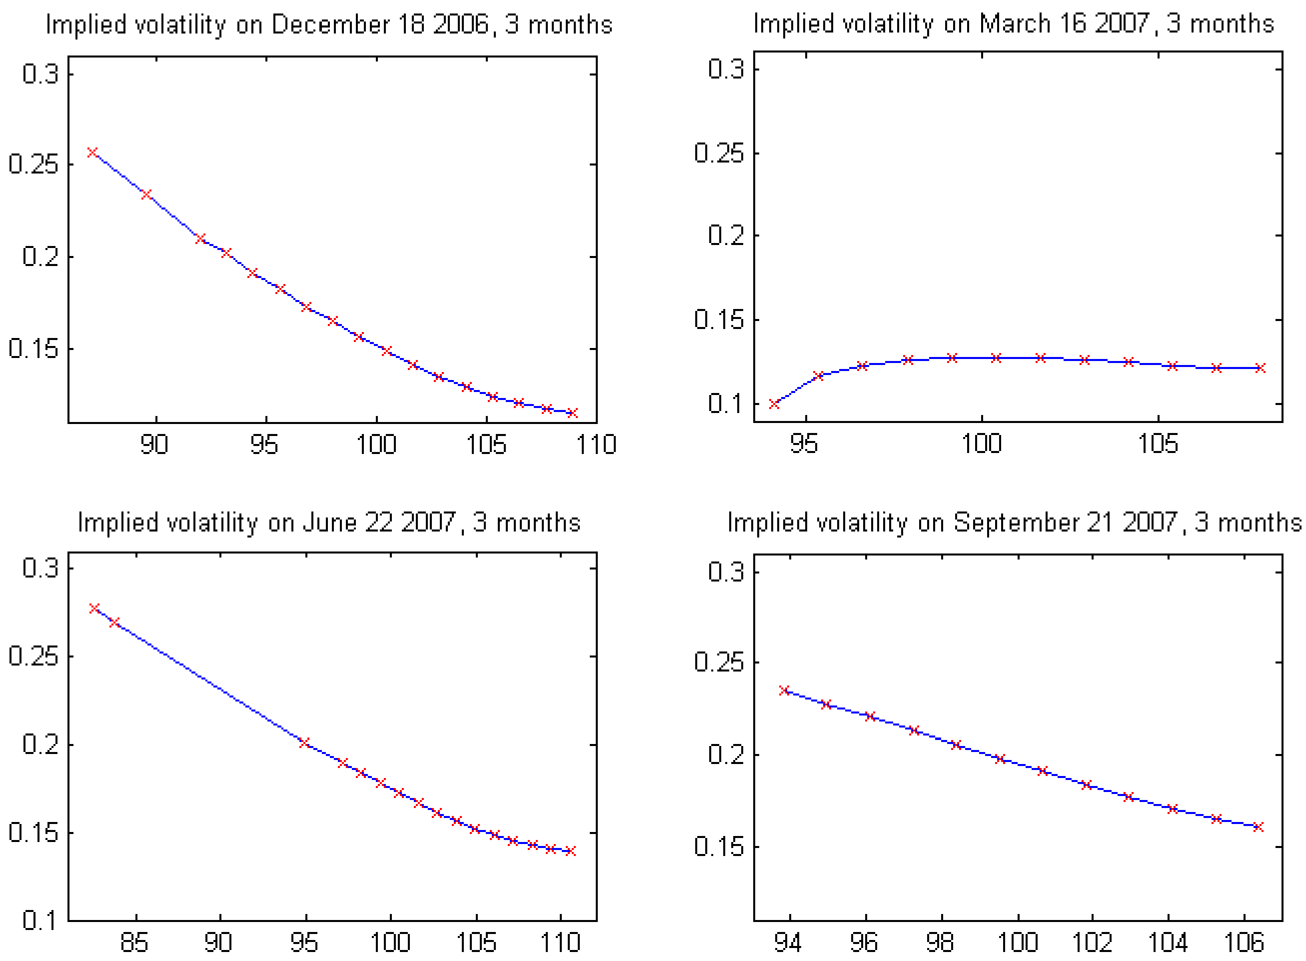

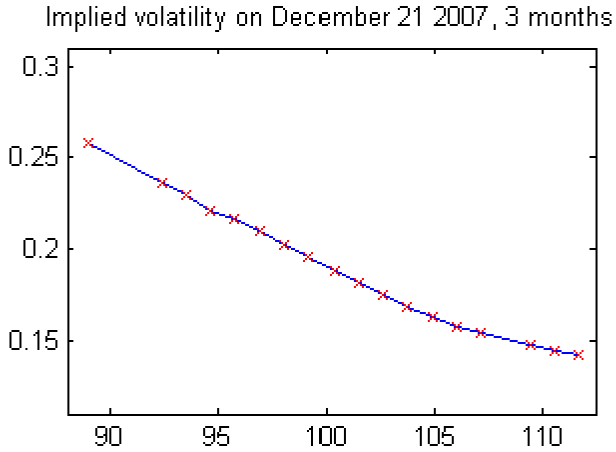
Implied volatility on December 21 2007, 3 months
0.3
0.25
0.2
0.15
90
95
100
105
110

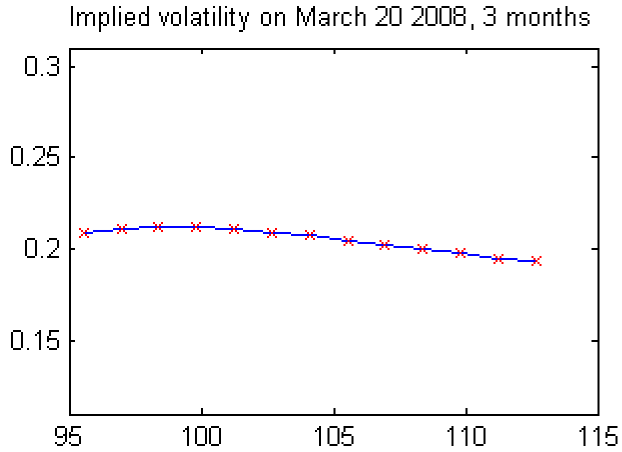
Implied volatility on March 20 2008, 3 months
0.3
0.25
0.2
0.15
95
100
105
110
115

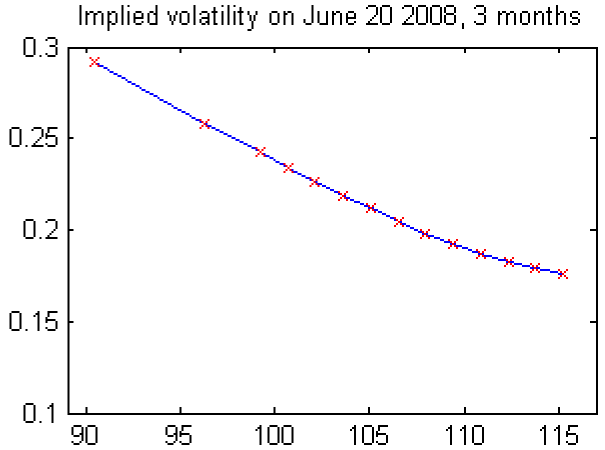
Implied volatility on June 20 2008, 3 months
0.3
0.25
0.2
0.15
0.1
90
95
100
105
110
115

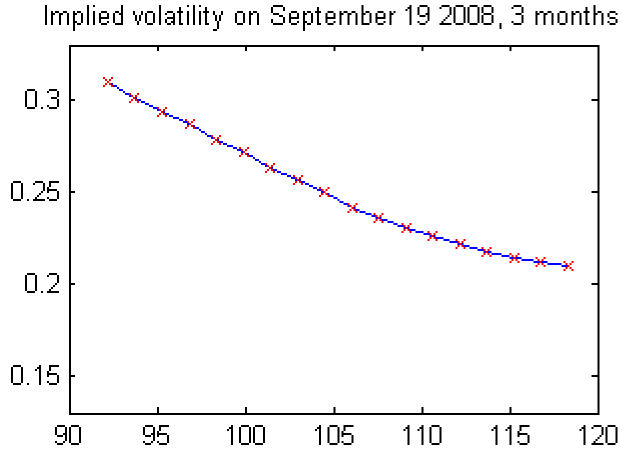
Implied volatility on September 19 2008, 3 months
0.3
0.25
0.2
0.15
90
95
100
105
110
115
120

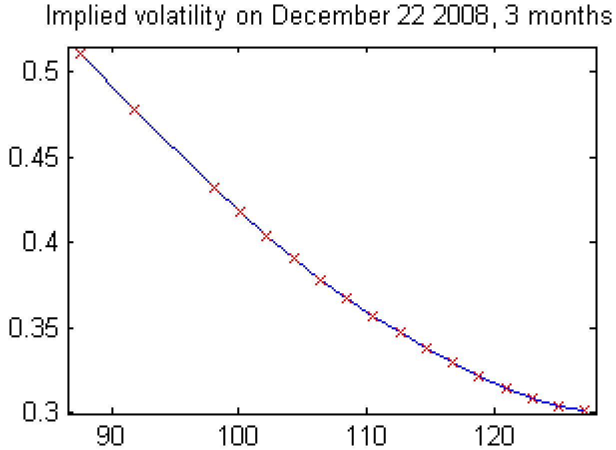
Implied volatility on December 22 2008, 3 months
0.5
0.45
0.4
0.35
0.3
90
100
110
120

Figure A2 – Volatility smiles for the period December 2006 to December 2008, six months

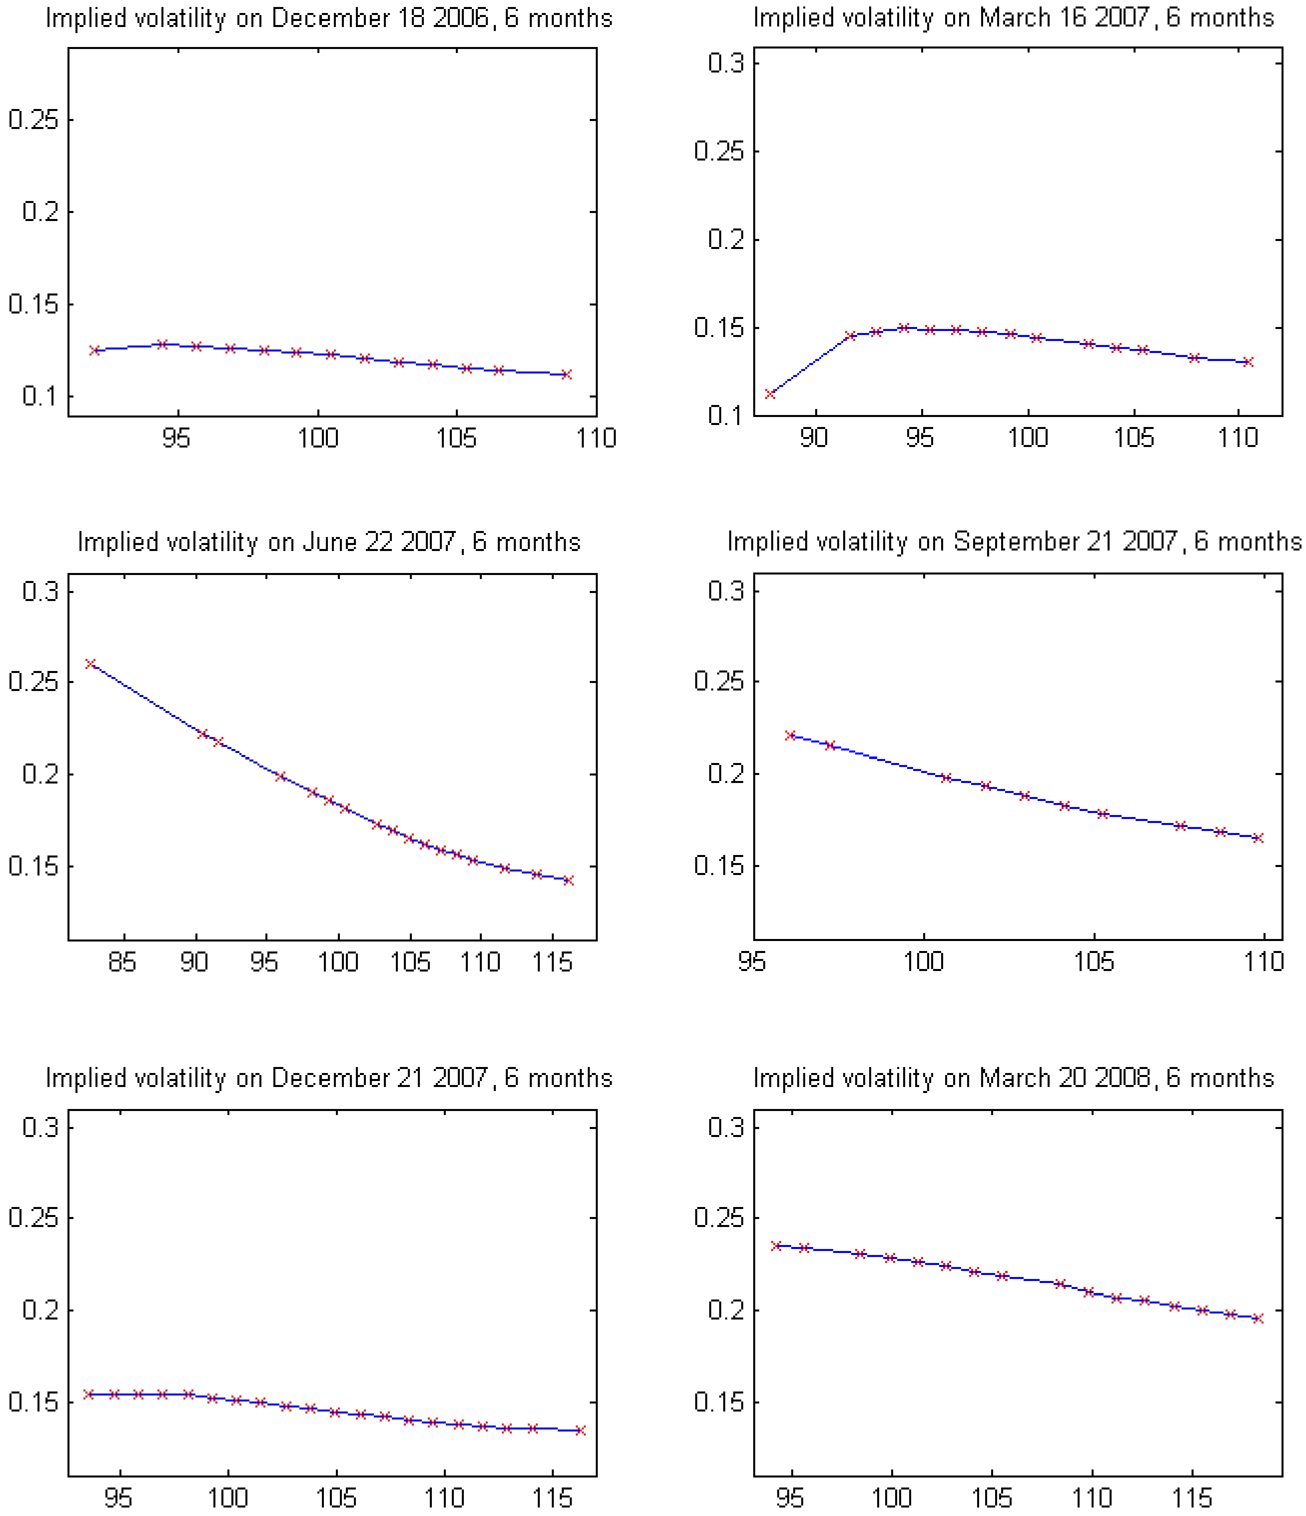

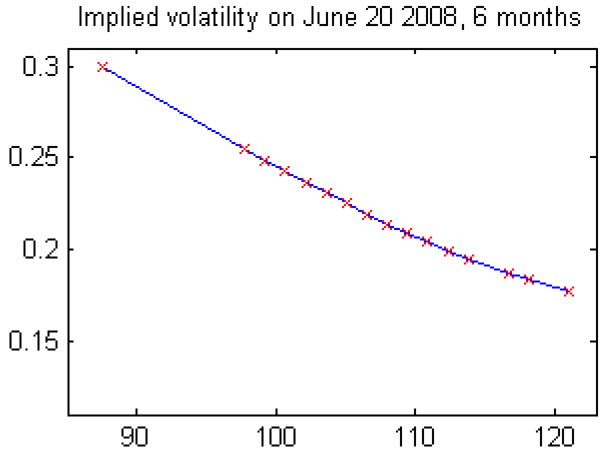

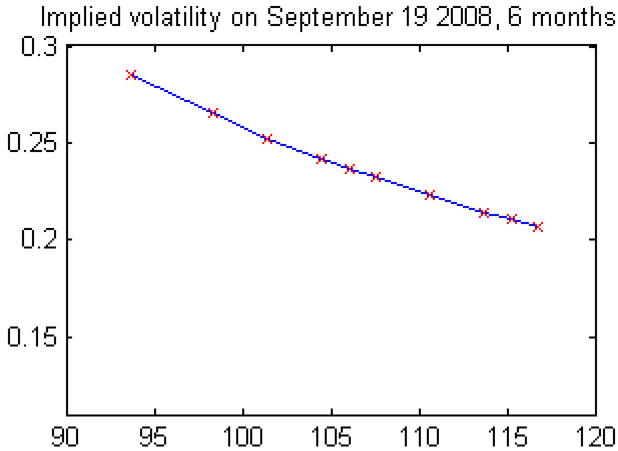

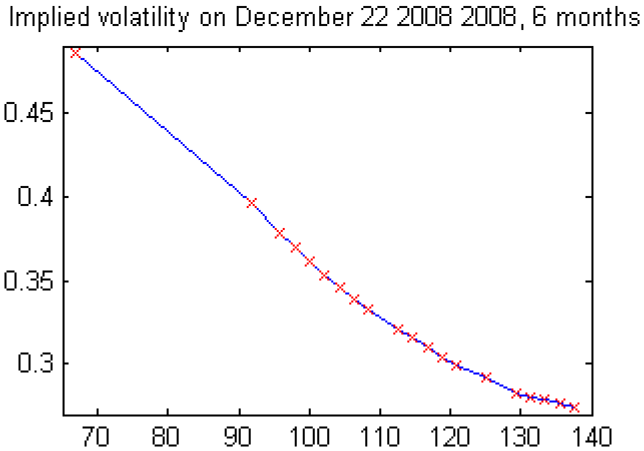

Figure A3 – Volatility smiles for the period December 2006 to December 2008, nine months

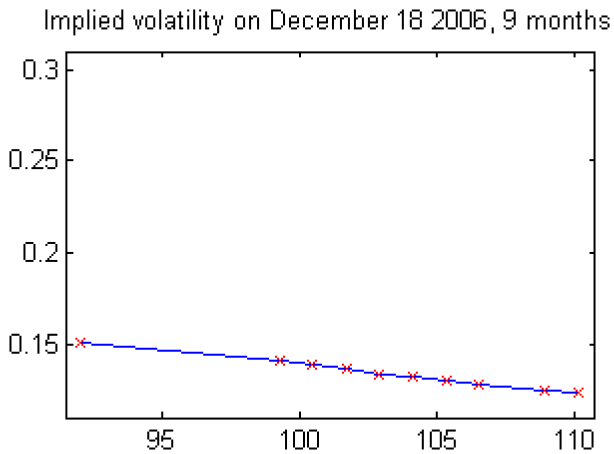

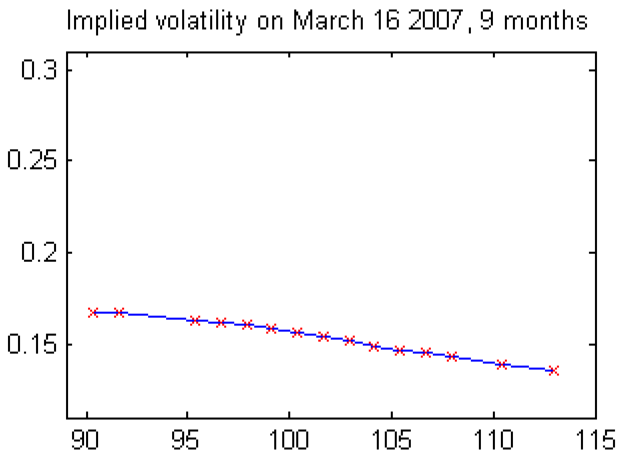

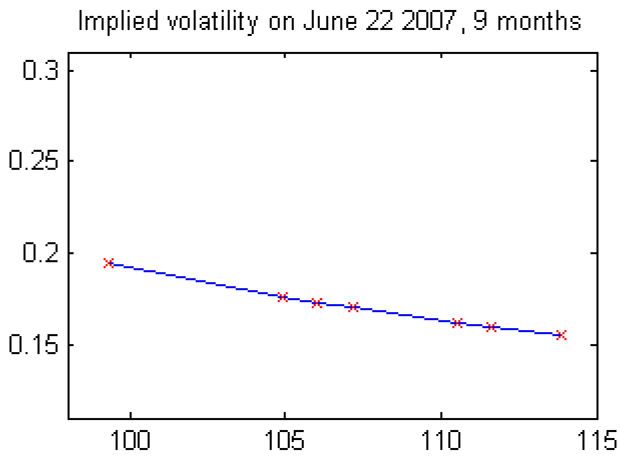
Implied volatility on June 22 2007, 9 months
0.3
0.25
0.2
0.15
100
105
110
115

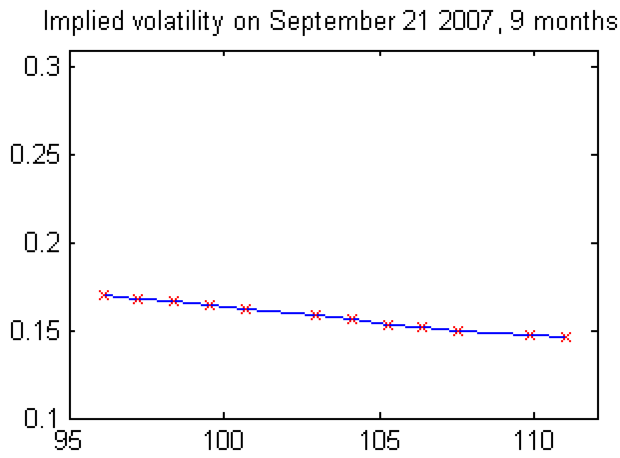
Implied volatility on September 21 2007, 9 months
0.3
0.25
0.2
0.15
0.1
95
100
105
110

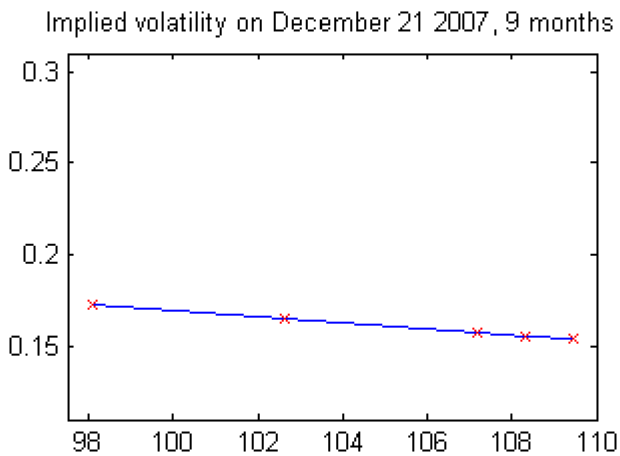
Implied volatility on December 21 2007, 9 months
0.3
0.25
0.2
0.15
98
100
102
104
106
108
110

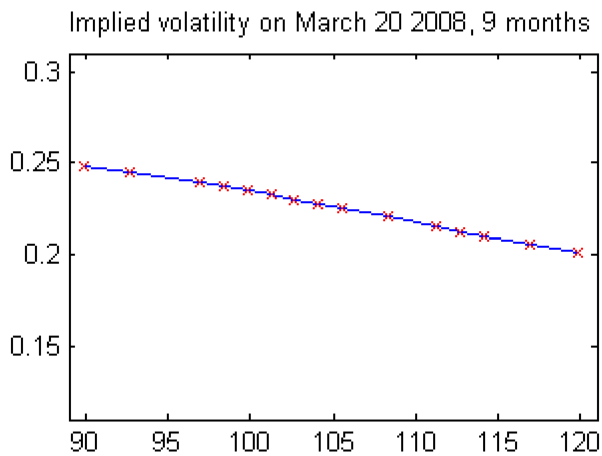
Implied volatility on March 20 2008, 9 months
0.3
0.25
0.2
0.15
90
95
100
105
110
115
120

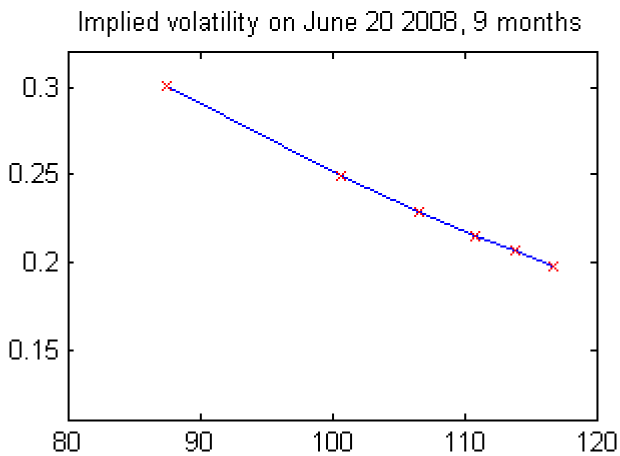
Implied volatility on June 20 2008, 9 months
0.3
0.25
0.2
0.15
80
90
100
110
120

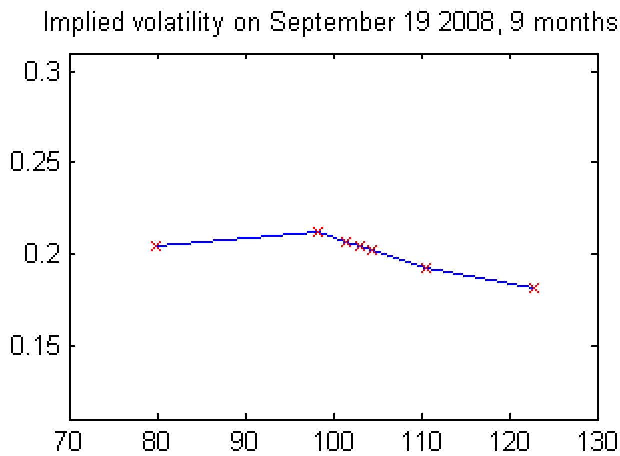
Implied volatility on September 19 2008, 9 months
0.3
0.25
0.2
0.15
70
80
90
100
110
120
130

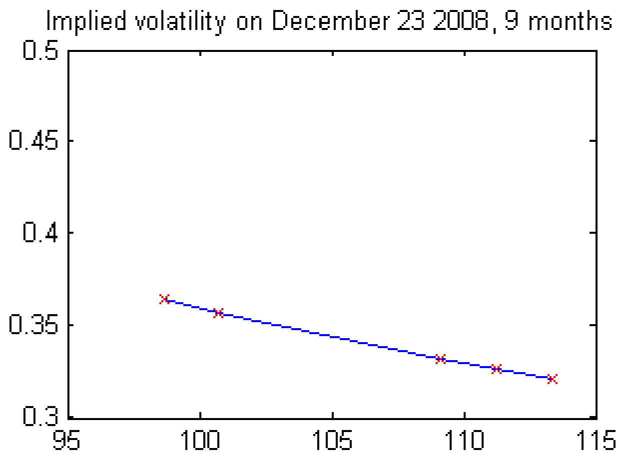

Figure A4 – Volatility smiles for every trading day of September 2008, three months

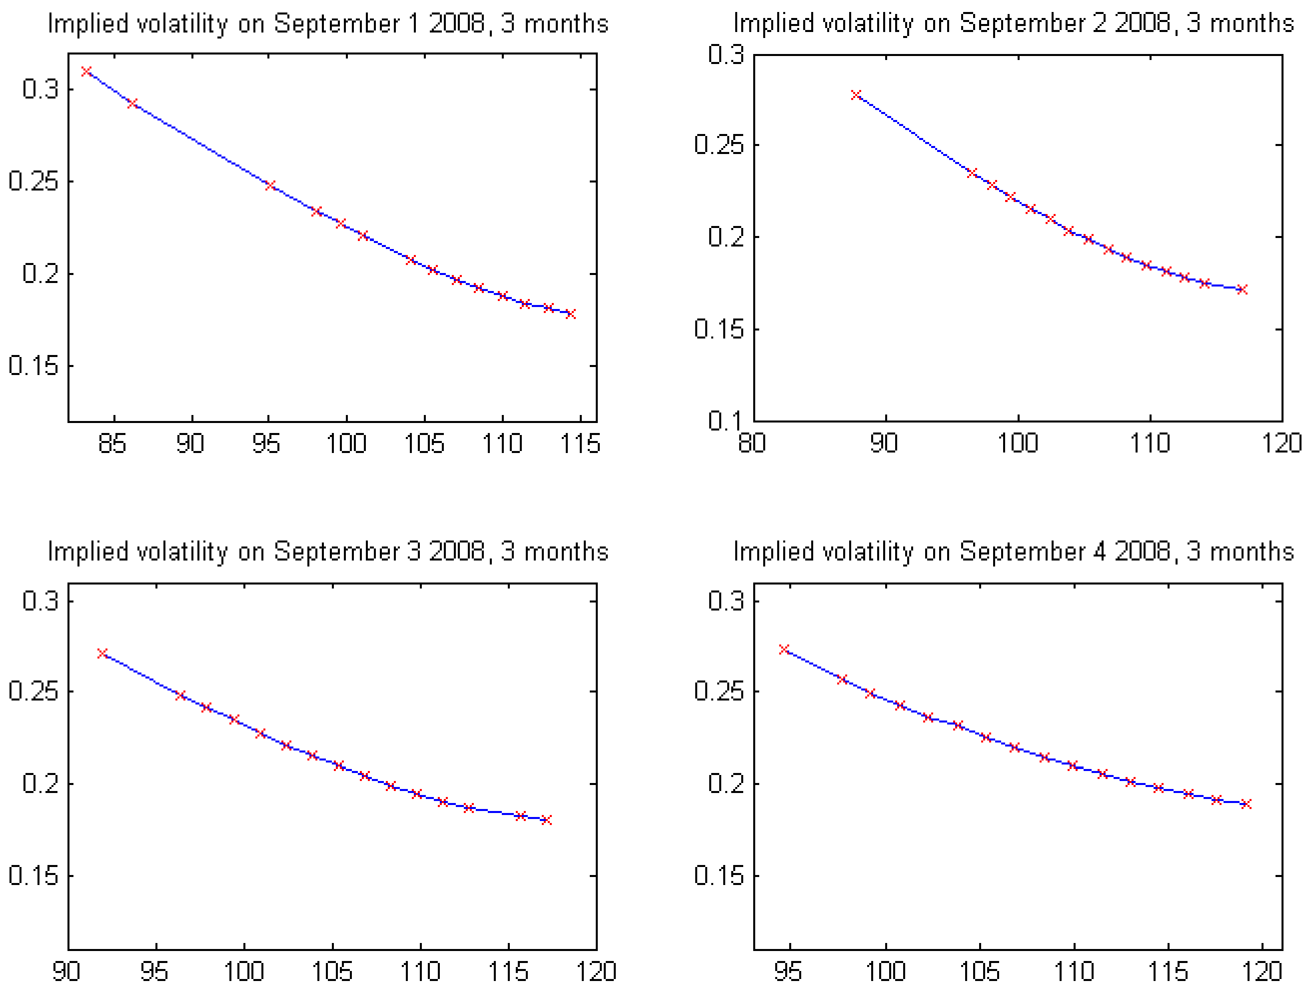

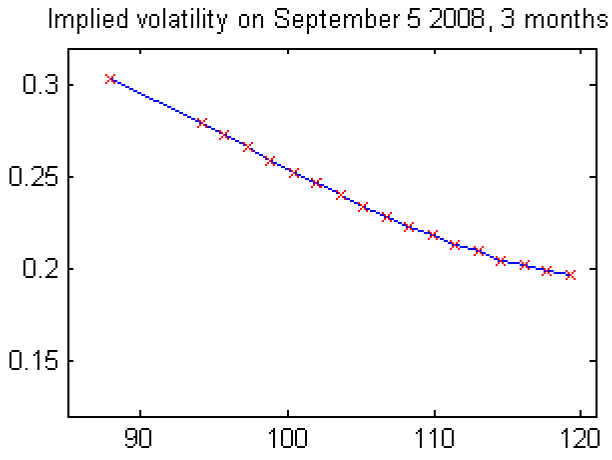
Implied volatility on September 5 2008, 3 months
0.3
0.25
0.2
0.15
90
100
110
120

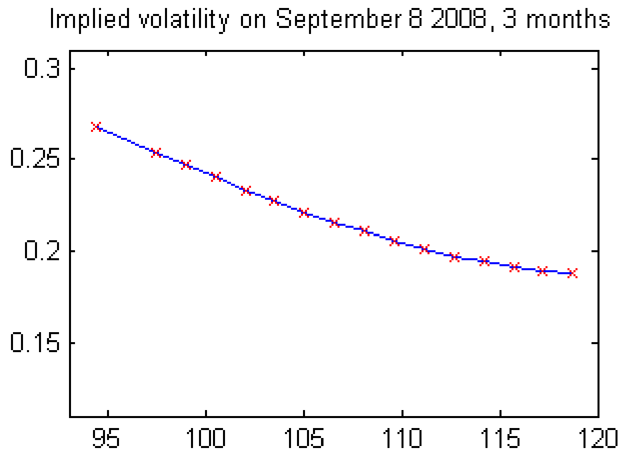
Implied volatility on September 8 2008, 3 months
0.3
0.25
0.2
0.15
95
100
105
110
115
120

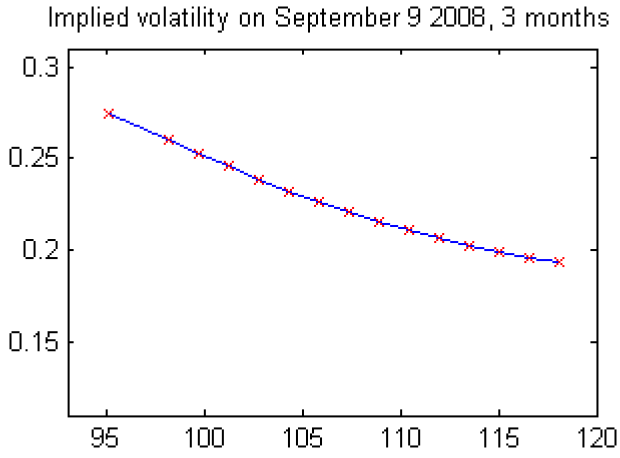
Implied volatility on September 9 2008, 3 months
0.3
0.25
0.2
0.15
95
100
105
110
115
120

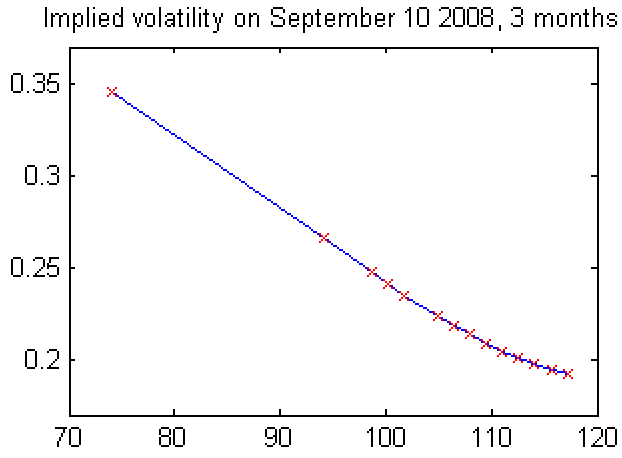
Implied volatility on September 10 2008, 3 months
0.35
0.3
0.25
0.2
70
80
90
100
110
120

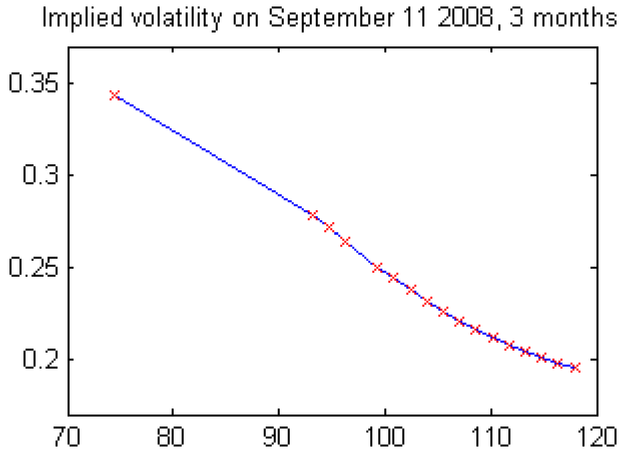
Implied volatility on September 11 2008, 3 months
0.35
0.3
0.25
0.2
70
80
90
100
110
120

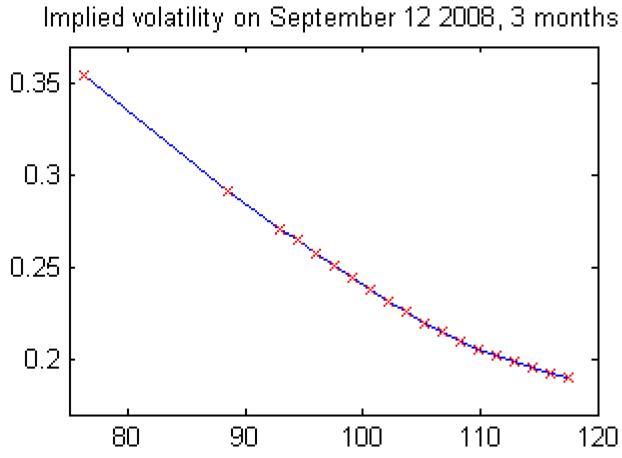
Implied volatility on September 12 2008, 3 months
0.35
0.3
0.25
0.2
80
90
100
110
120

Implied volatility on September 15 2008, 3 months

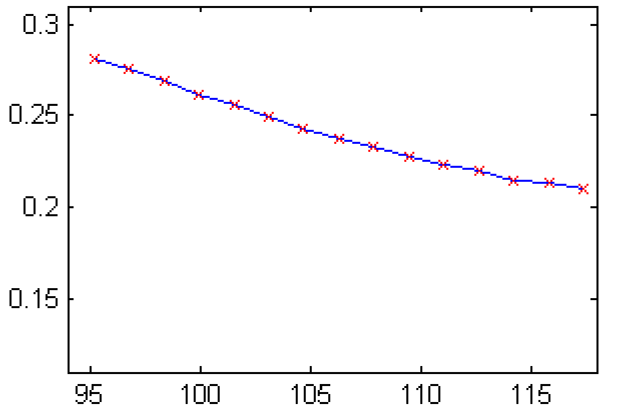

Implied volatility on September 16 2008, 3 months

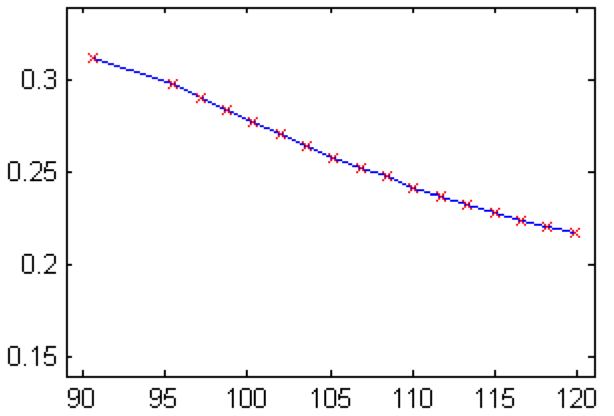

Implied volatility on September 17 2008, 3 months

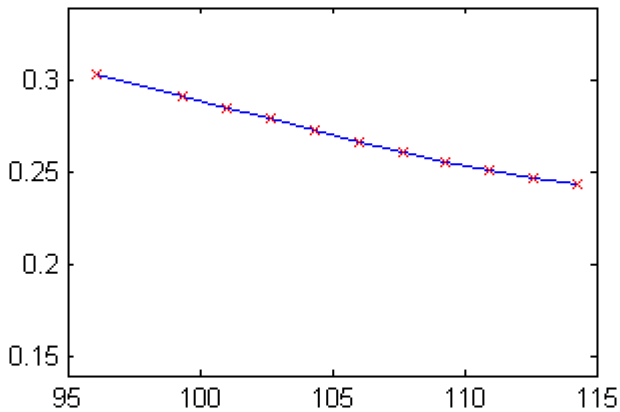

Implied volatility on September 18 2008, 3 months

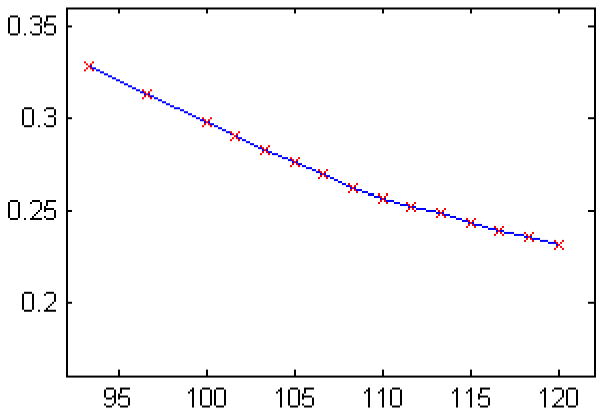

Implied volatility on September 19 2008, 3 months

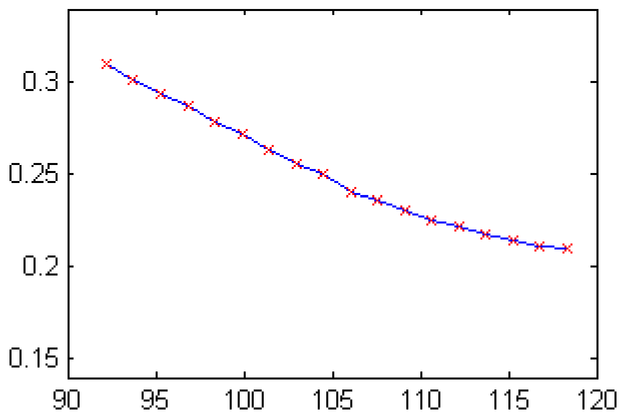

Implied volatility on September 22 2008, 3 months

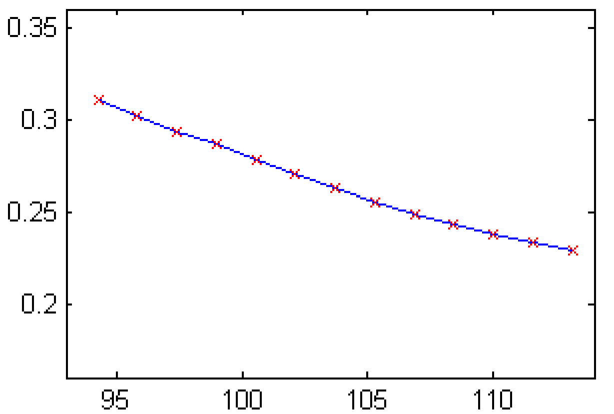

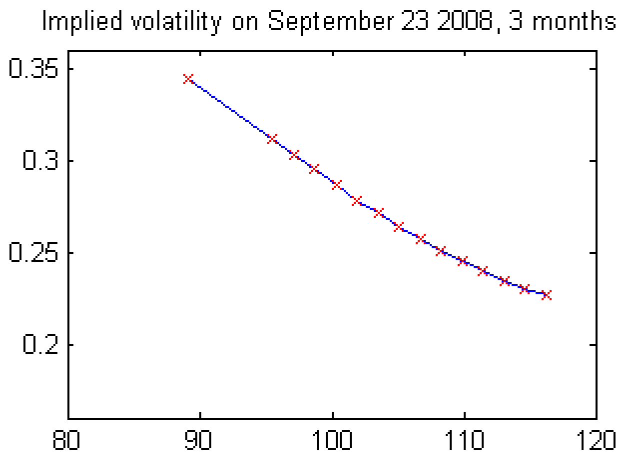
Implied volatility on September 23 2008, 3 months
0.35
0.3
0.25
0.2
80
90
100
110
120

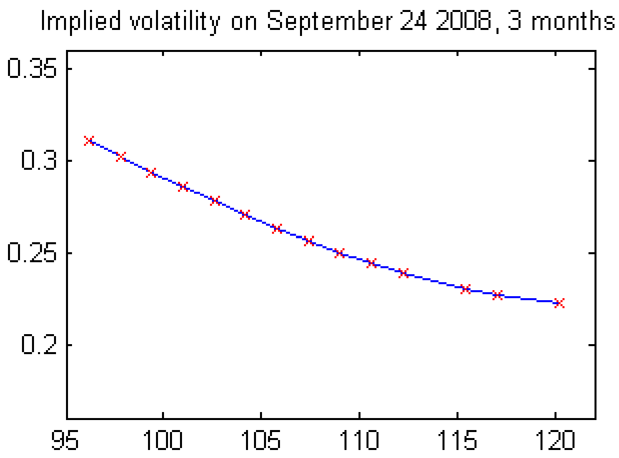
Implied volatility on September 24 2008, 3 months
0.35
0.3
0.25
0.2
95
100
105
110
115
120

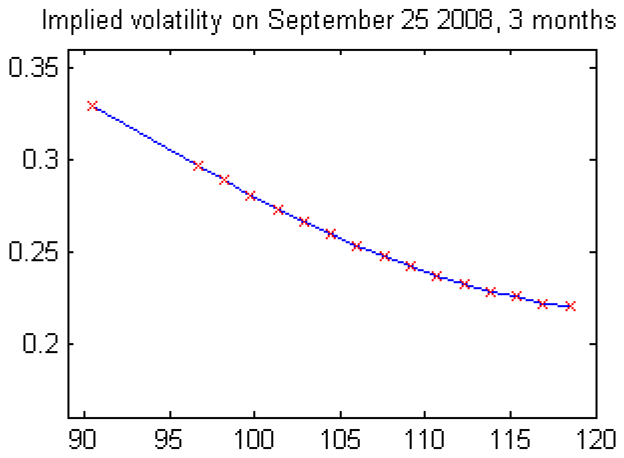
Implied volatility on September 25 2008, 3 months
0.35
0.3
0.25
0.2
90
95
100
105
110
115
120

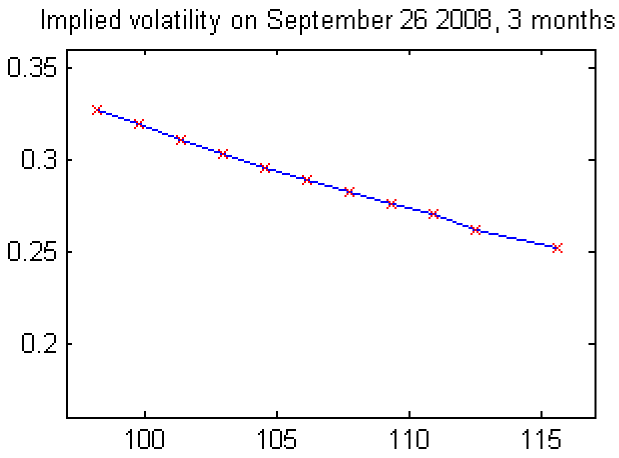
Implied volatility on September 26 2008, 3 months
0.35
0.3
0.25
0.2
100
105
110
115

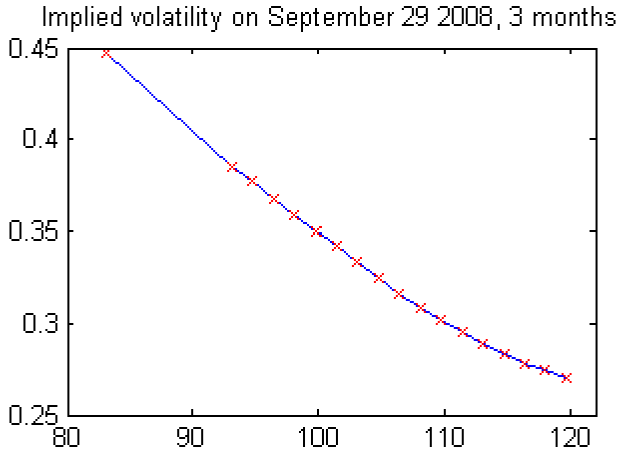
Implied volatility on September 29 2008, 3 months
0.45
0.4
0.35
0.3
0.25
80
90
100
110
120

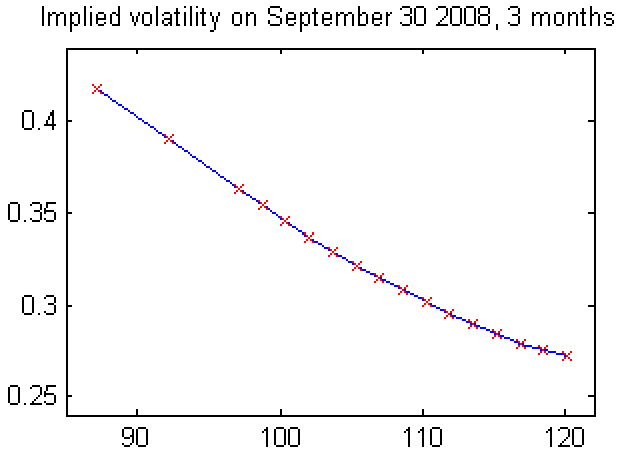
Implied volatility on September 30 2008, 3 months
0.4
0.35
0.3
0.25
90
100
110
120

Appendix B – RND functions

Here, we present all of the retrieved implied RND functions, along with the estimated parameters and the previously described descriptive statistics for each RND function. In each graph, the dashed green line corresponds to the first set of parameters, the dashed red line corresponds to the second set of parameters and the solid blue line is the weighted sum, i.e. the RND function itself. All of the presented RNDs have a maturity of (roughly) three months. In Figure B1, RND functions for the last month of every quarter during the period December 2006 to December 2008 are presented. RND functions for every trading day of September 2008 are presented in figure B2.

Figure B1 – RND functions for the period December 2006 to December 2008, three months

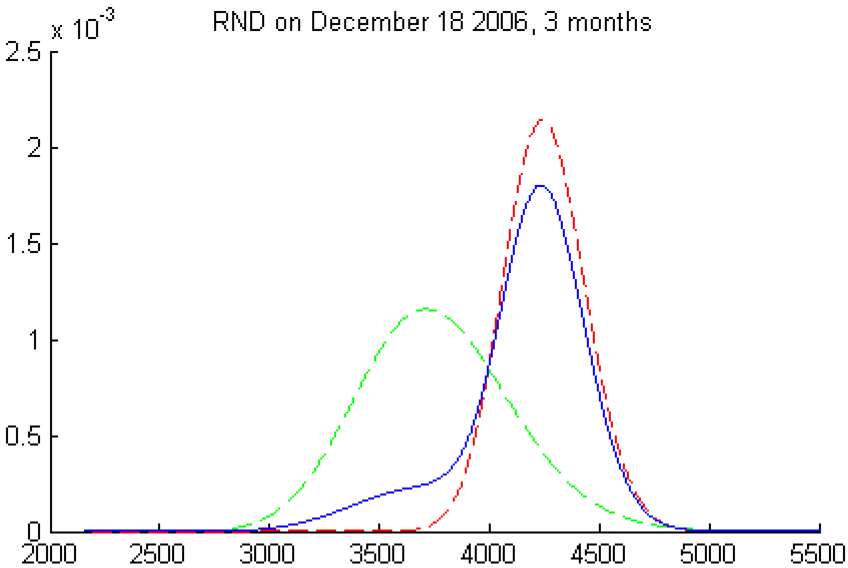

December 18 2006	
Parameters	
$\boldsymbol{\mu_1}$	8.2290
$\boldsymbol{\sigma_1}$	0.0924
$\boldsymbol{\mu_2}$	8.3544
$\boldsymbol{\sigma_2}$	0.0439
$\boldsymbol{\theta}$	0.1972
Descriptive statistics	
$\mathbb{E}$	4156.7090
$\mathbb{D}$	299.6706
$\%\mathbb{D}$	3.5829%
$\mathbb{S}$	-1.0507
$\mathbb{K}$	4.6064

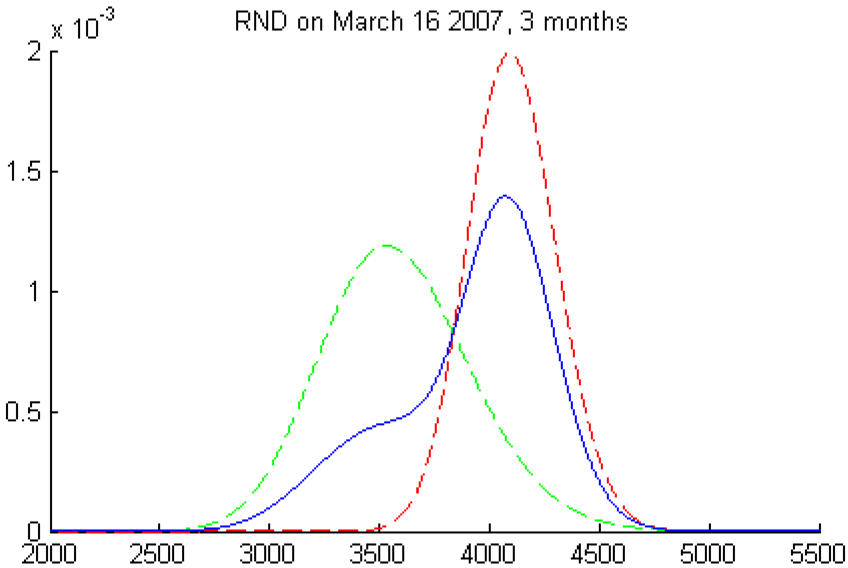

March 16 2007	
Parameters	
$\boldsymbol{\mu_1}$	8.1800
$\boldsymbol{\sigma_1}$	0.0946
$\boldsymbol{\mu_2}$	8.3192
$\boldsymbol{\sigma_2}$	0.0489
$\boldsymbol{\theta}$	0.3692
Descriptive statistics	
$\mathbb{E}$	3914.0015
$\mathbb{D}$	362.6529
$\%\mathbb{D}$	4.5695%
$\mathbb{S}$	-0.6525
$\mathbb{K}$	3.0254

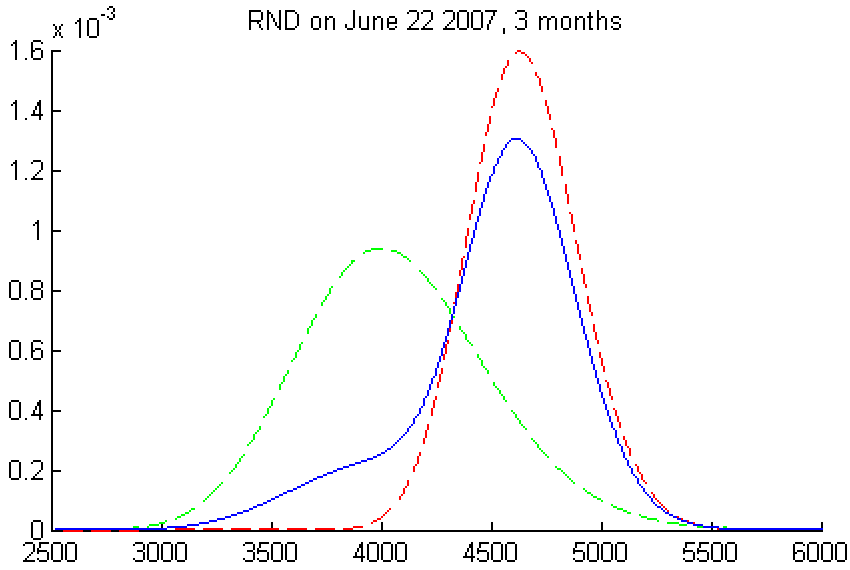

June 22 2007	
Parameters	
$\boldsymbol{\mu_1}$	8.3049
$\boldsymbol{\sigma_1}$	0.1056
$\boldsymbol{\mu_2}$	8.4428
$\boldsymbol{\sigma_2}$	0.0540
$\boldsymbol{\theta}$	0.2365
Descriptive statistics	
$\mathbb{E}$	4510.6144
$\mathbb{D}$	391.3381
$\%\mathbb{D}$	4.3861%
$\mathbb{S}$	-0.8226
$\mathbb{K}$	3.9374

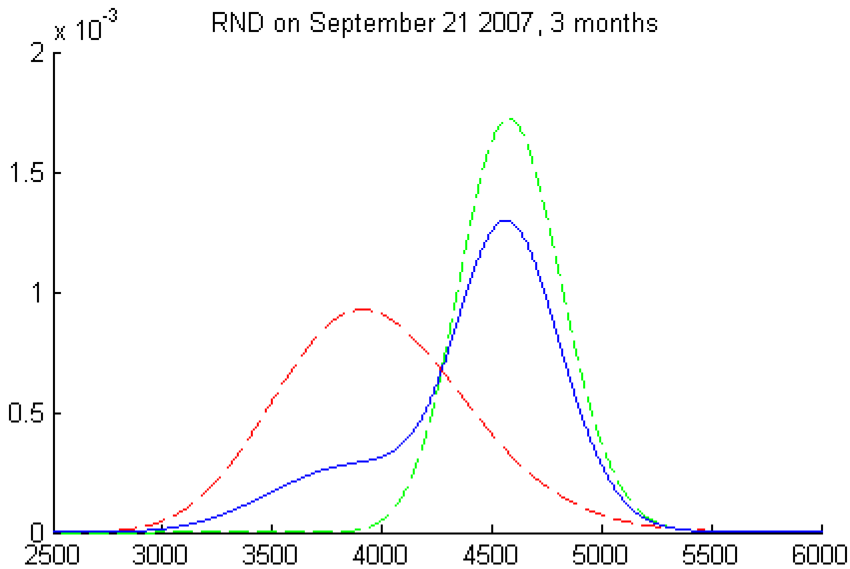

September 21 2007	
Parameters	
$\boldsymbol{\mu_1}$	8.4314
$\boldsymbol{\sigma_1}$	0.0505
$\boldsymbol{\mu_2}$	8.2847
$\boldsymbol{\sigma_2}$	0.1090
$\boldsymbol{\theta}$	0.6931
Descriptive statistics	
$\mathbb{E}$	4407.9957
$\mathbb{D}$	417.5276
$\%\mathbb{D}$	4.7964%
$\mathbb{S}$	-0.8456
$\mathbb{K}$	3.5798

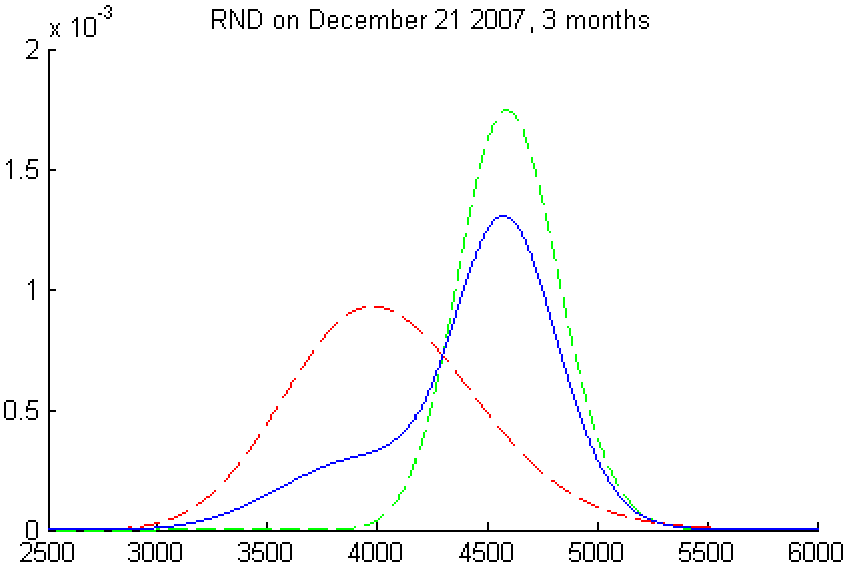

December 21 2007	
Parameters	
μ_1	8.4336
σ_1	0.0498
μ_2	8.3018
σ_2	0.1069
θ	0.6724
Descriptive statistics	
$\mathbb{E}$	4424.2599
$\mathbb{D}$	404.8866
$\%\mathbb{D}$	4.6361%
$\mathbb{S}$	-0.7892
$\mathbb{K}$	3.5523

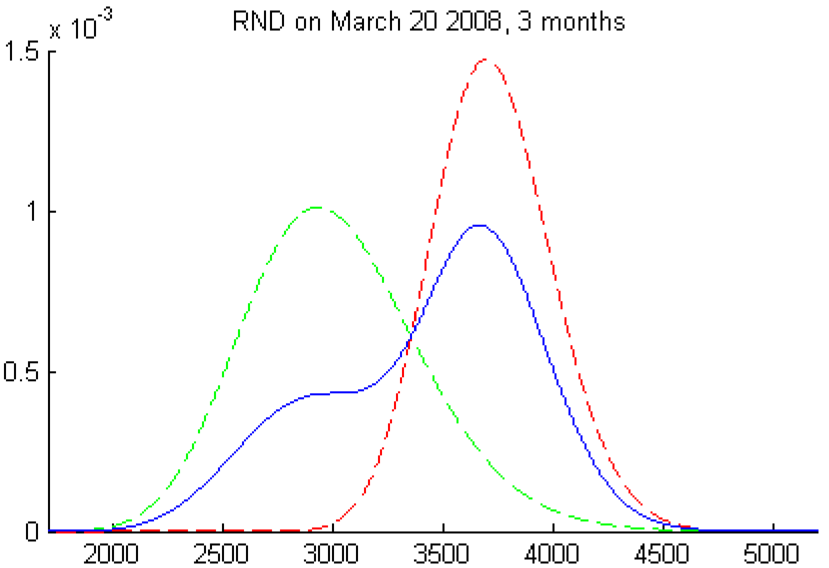

March 20 2008	
Parameters	
μ_1	8.0017
σ_1	0.1333
μ_2	8.2206
σ_2	0.0731
θ	0.4184
Descriptive statistics	
$\mathbb{E}$	3427.9948
$\mathbb{D}$	485.2163
$\%\mathbb{D}$	6.9860%
$\mathbb{S}$	-0.4096
$\mathbb{K}$	2.5579

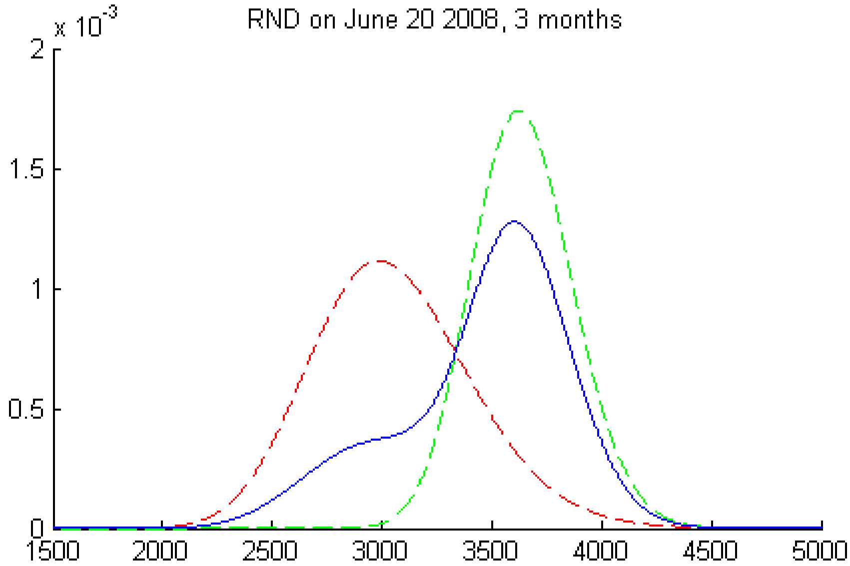

June 20 2008	
Parameters	
μ_1	8.1990
σ_1	8.0174
μ_2	0.0632
σ_2	0.1190
θ	0.6753
Descriptive statistics	
$\mathbb{E}$	3453.0633
$\mathbb{D}$	394.0824
$\%\mathbb{D}$	5.7739%
$\mathbb{S}$	-0.6457
$\mathbb{K}$	3.0944

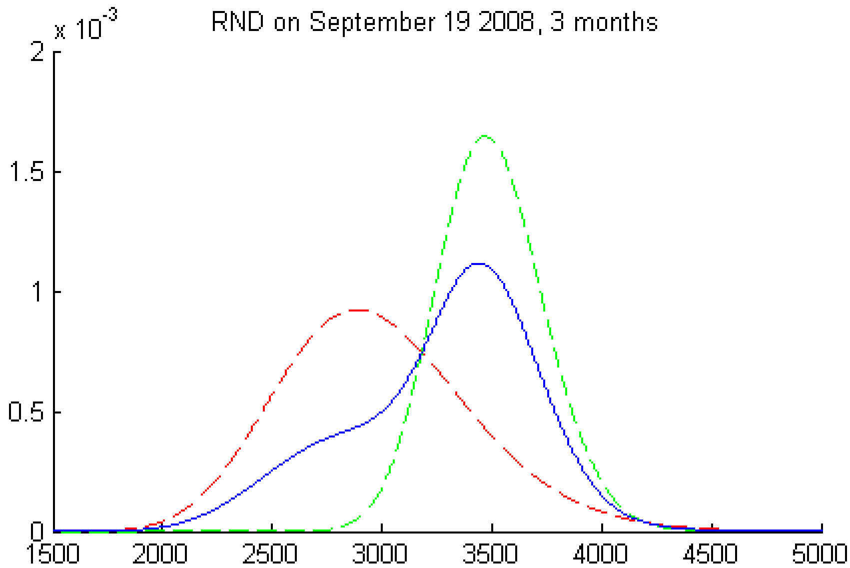

September 19 2008	
Parameters	
μ_1	8.1574
σ_1	0.0697
μ_2	7.9912
σ_2	0.1478
θ	0.5599
Descriptive statistics	
$\mathbb{E}$	3272.9015
$\mathbb{D}$	429.0790
$\%\mathbb{D}$	6.6212%
$\mathbb{S}$	-0.4670
$\mathbb{K}$	3.0008

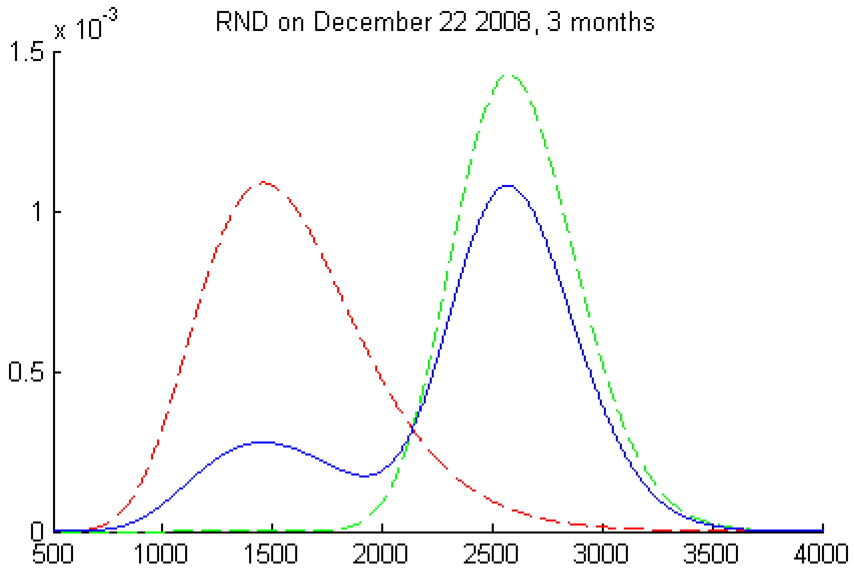

December 22 2008	
Parameters	
μ_1	7.8655
σ_1	0.1080
μ_2	7.3457
σ_2	0.2438
θ	0.7443
Descriptive statistics	
$\mathbb{E}$	2358.9922
$\mathbb{D}$	547.4694
$\%\mathbb{D}$	11.2749%
$\mathbb{S}$	-0.7725
$\mathbb{K}$	2.9372

Figure B2 – RND functions for every trading day of September 2008, three months

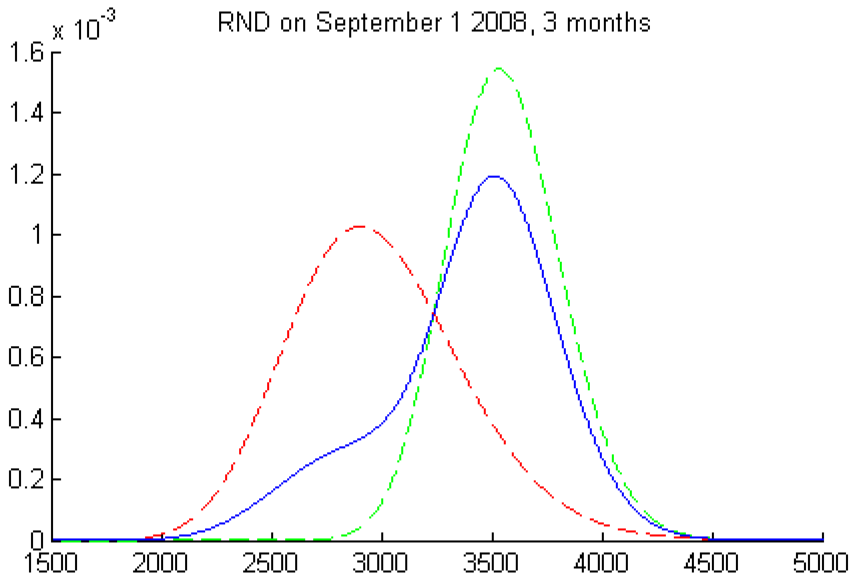

September 1 2008	
Parameters	
μ_1	8.1748
σ_1	0.0730
μ_2	7.9913
σ_2	0.1327
θ	0.7053
Descriptive statistics	
$\mathbb{E}$	3389.1821
$\mathbb{D}$	404.8018
$\%\mathbb{D}$	6.6052%
$\mathbb{S}$	-0.5856
$\mathbb{K}$	3.2554

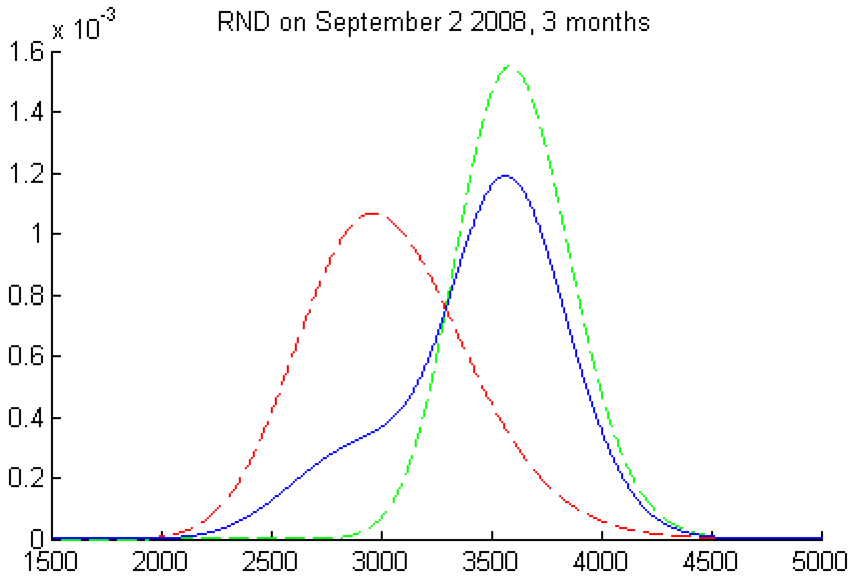

September 2 2008	
Parameters	
μ_1	8.1898
σ_1	0.0716
μ_2	8.0100
σ_2	0.1250
θ	0.7009
Descriptive statistics	
$\mathbb{E}$	3440.0362
$\mathbb{D}$	400.6783
$\%\mathbb{D}$	6.4089%
$\mathbb{S}$	-0.5564
$\mathbb{K}$	3.1734

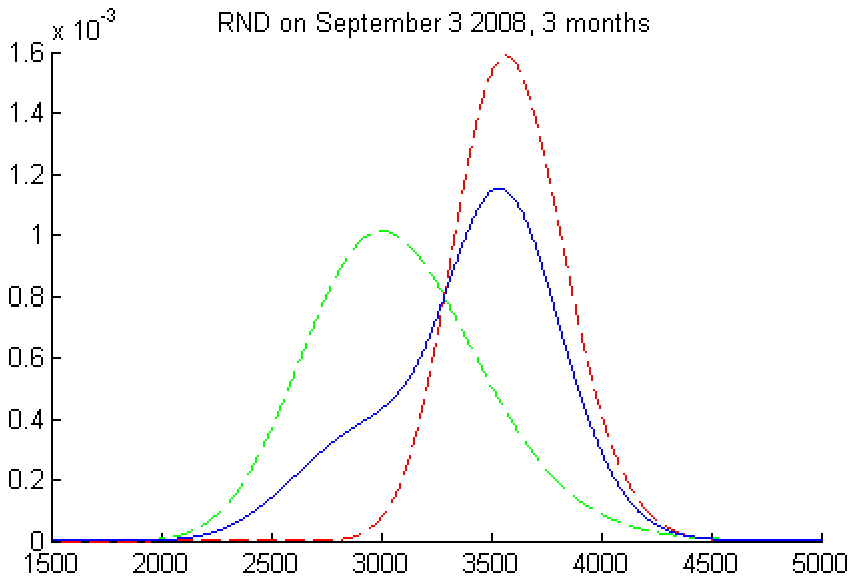

September 3 2008	
Parameters	
μ_1	8.0243
σ_1	0.1301
μ_2	8.1840
σ_2	0.0703
θ	0.3791
Descriptive statistics	
$\mathbb{E}$	3397.9982
$\mathbb{D}$	403.4544
$\%\mathbb{D}$	6.5142%
$\mathbb{S}$	-0.4750
$\mathbb{K}$	3.0397

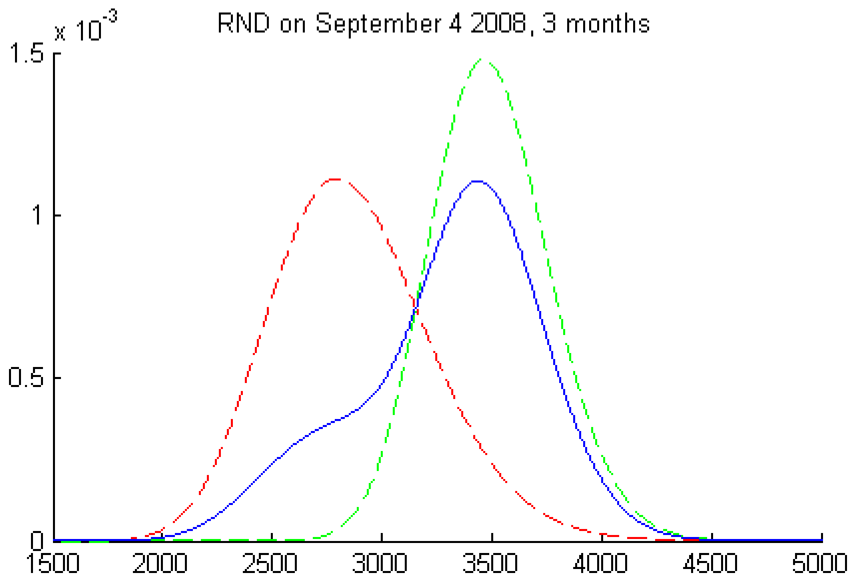

September 4 2008	
Parameters	
μ_1	8.1562
σ_1	0.0777
μ_2	7.9536
σ_2	0.1271
θ	0.6881
Descriptive statistics	
$\mathbb{E}$	3300.0156
$\mathbb{D}$	420.6857
$\%\mathbb{D}$	6.9553%
$\mathbb{S}$	-0.4690
$\mathbb{K}$	2.9497

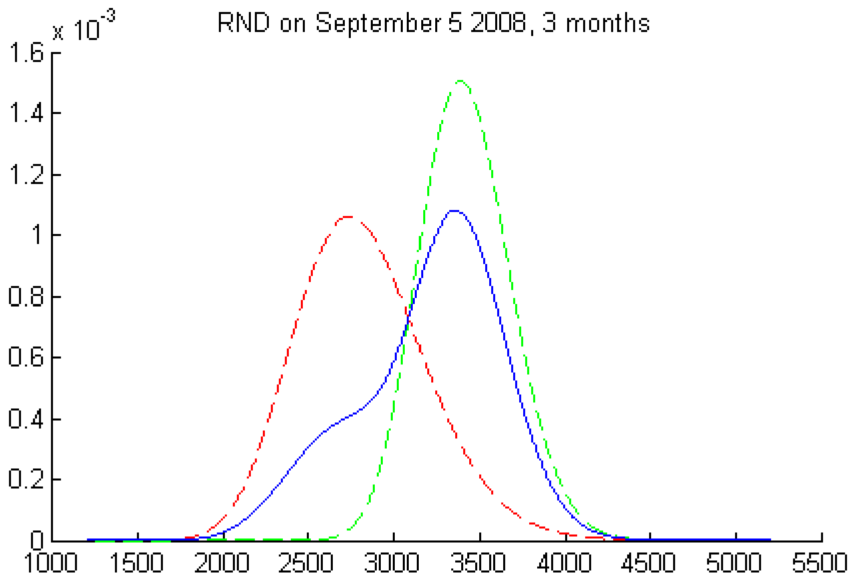

September 5 2008	
Parameters	
μ_1	8.1344
σ_1	0.0779
μ_2	7.9347
σ_2	0.1359
θ	0.6391
Descriptive statistics	
$\mathbb{E}$	3203.0495
$\mathbb{D}$	427.1959
$\%\mathbb{D}$	7.2263%
$\mathbb{S}$	-0.4446
$\mathbb{K}$	2.8600

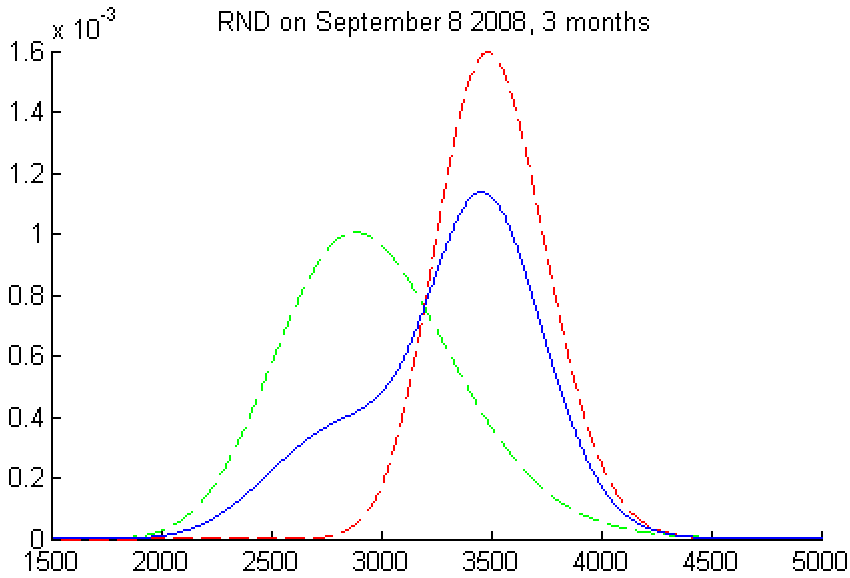

September 8 2008	
Parameters	
$\boldsymbol{\mu_1}$	7.9852
$\boldsymbol{\sigma_1}$	0.1364
$\boldsymbol{\mu_2}$	8.1610
$\boldsymbol{\sigma_2}$	0.0716
$\boldsymbol{\theta}$	0.3838
Descriptive statistics	
$\mathbb{E}$	3300.9885
$\mathbb{D}$	415.7171
$\%\mathbb{D}$	6.7244%
$\mathbb{S}$	-0.4957
$\mathbb{K}$	2.9780

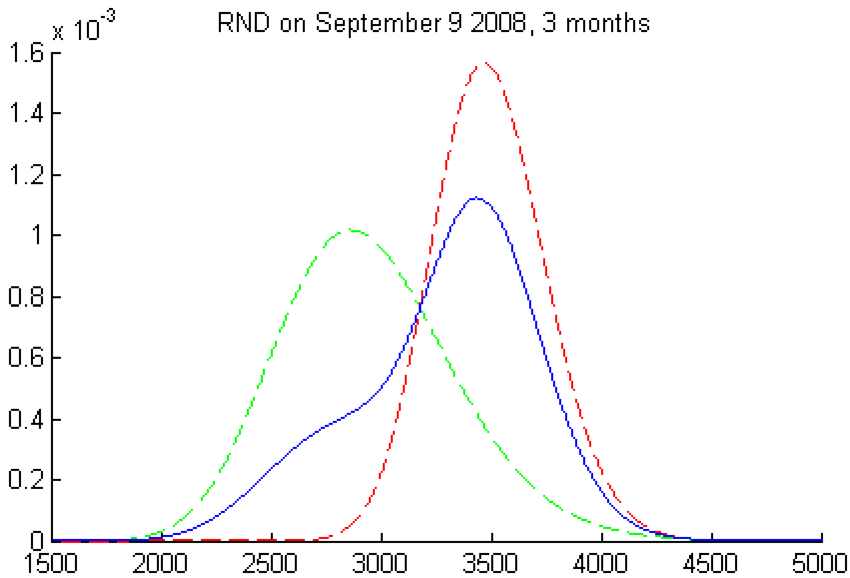

September 9 2008	
Parameters	
$\boldsymbol{\mu_1}$	7.9778
$\boldsymbol{\sigma_1}$	0.1359
$\boldsymbol{\mu_2}$	8.1558
$\boldsymbol{\sigma_2}$	0.0734
$\boldsymbol{\theta}$	0.3794
Descriptive statistics	
$\mathbb{E}$	3283.9904
$\mathbb{D}$	416.3164
$\%\mathbb{D}$	6.7486%
$\mathbb{S}$	-0.4732
$\mathbb{K}$	2.9558

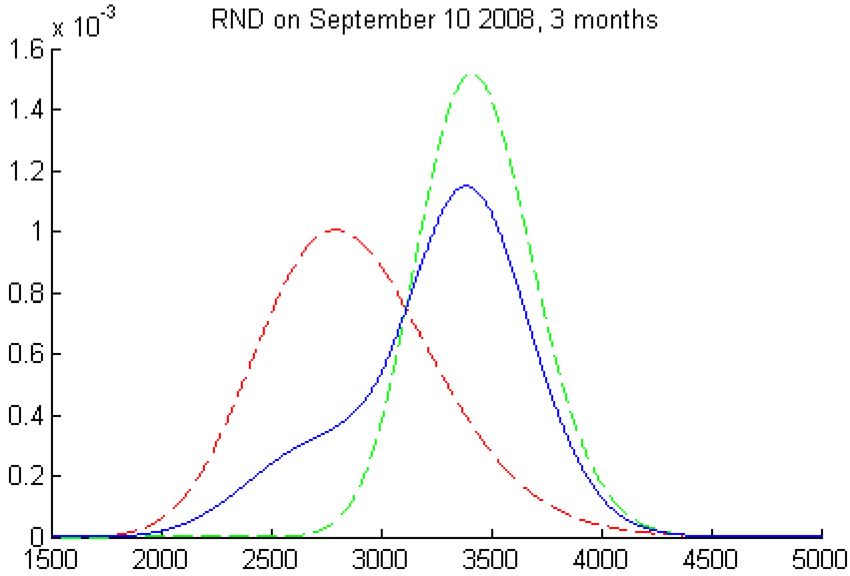

September 10 2008	
Parameters	
μ_1	8.1400
σ_1	0.0769
μ_2	7.9539
σ_2	0.1408
θ	0.6767
Descriptive statistics	
$\mathbb{E}$	3256.8457
$\mathbb{D}$	413.1230
$\%\mathbb{D}$	6.7032%
$\mathbb{S}$	-0.5248
$\mathbb{K}$	3.1446

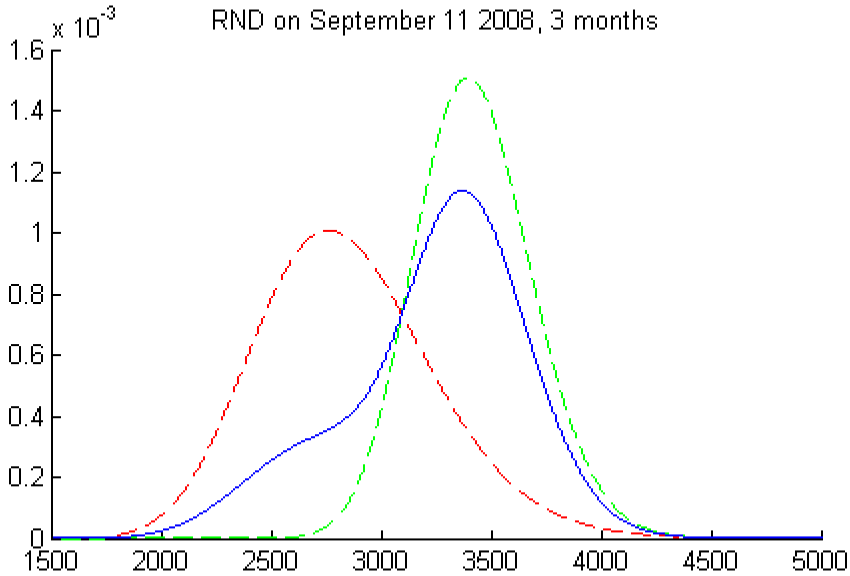

September 11 2008	
Parameters	
μ_1	8.1354
σ_1	0.0778
μ_2	7.9431
σ_2	0.1419
θ	0.6770
Descriptive statistics	
$\mathbb{E}$	3236.4702
$\mathbb{D}$	417.8265
$\%\mathbb{D}$	6.7875%
$\mathbb{S}$	-0.5242
$\mathbb{K}$	3.1155

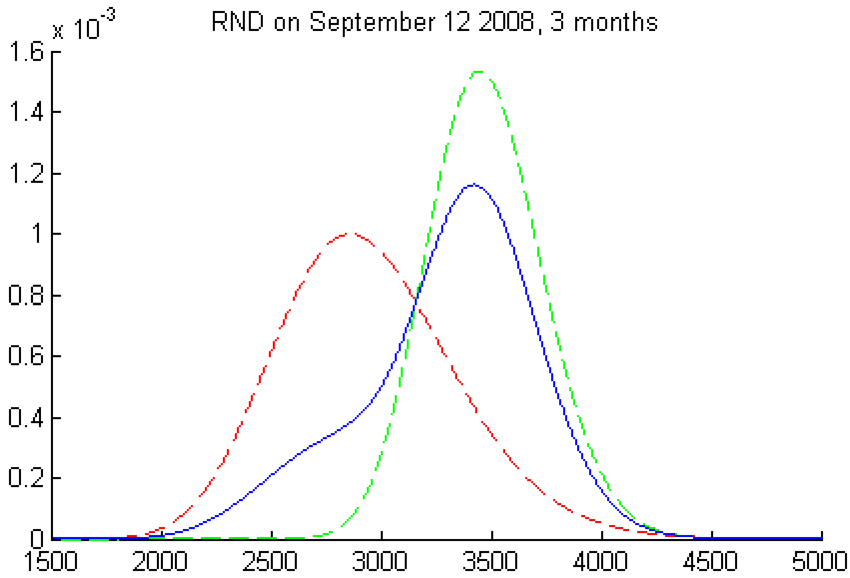

September 12 2008	
Parameters	
μ_1	8.1503
σ_1	0.0753
μ_2	7.9775
σ_2	0.1383
θ	0.6652
Descriptive statistics	
$\mathbb{E}$	3296.2857
$\mathbb{D}$	405.6329
$\%\mathbb{D}$	6.4446%
$\mathbb{S}$	-0.5069
$\mathbb{K}$	3.1612

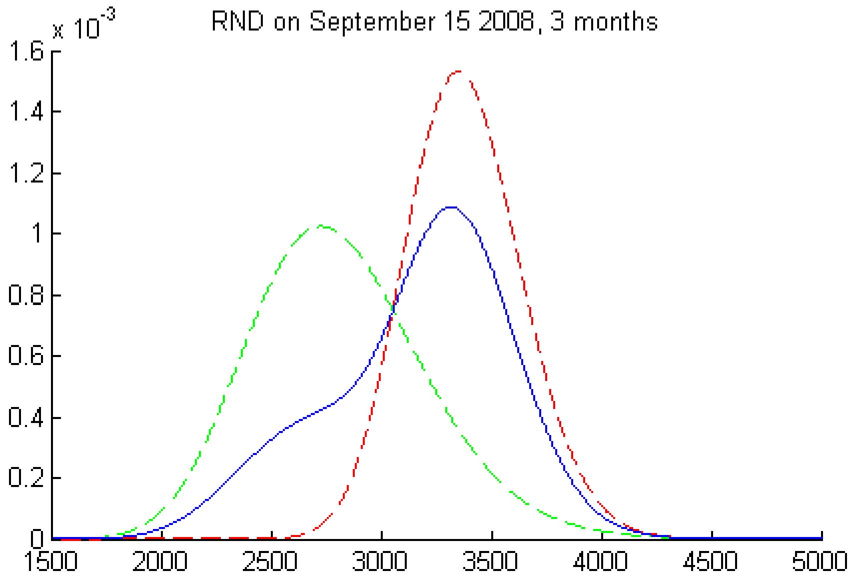

September 15 2008	
Parameters	
μ_1	7.9311
σ_1	0.1415
μ_2	8.1220
σ_2	0.0777
θ	0.3841
Descriptive statistics	
$\mathbb{E}$	3159.9857
$\mathbb{D}$	424.3070
$\%\mathbb{D}$	6.9055%
$\mathbb{S}$	-0.4377
$\mathbb{K}$	2.8825

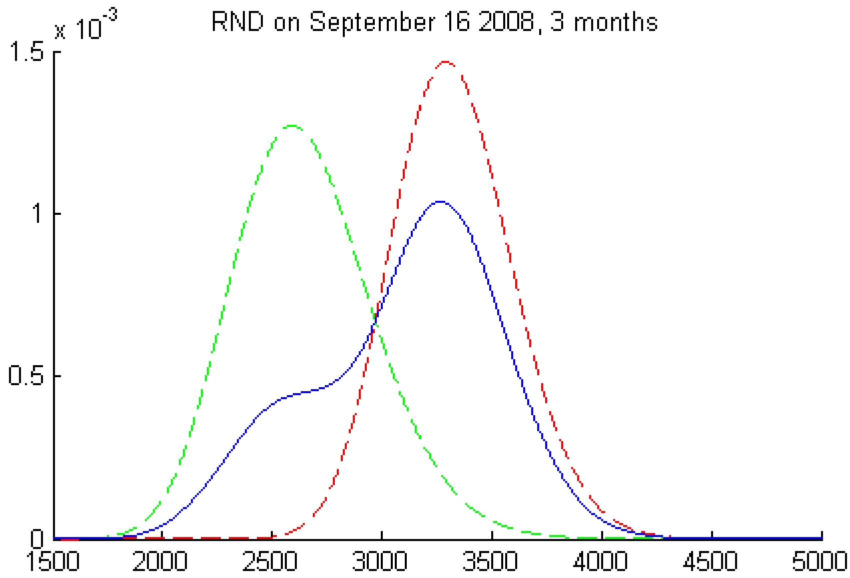

September 16 2008	
Parameters	
μ_1	7.8721
σ_1	0.1052
μ_2	7.5340
σ_2	0.1966
θ	0.6912
Descriptive statistics	
$\mathbb{E}$	3098.9961
$\mathbb{D}$	431.7782
$\%\mathbb{D}$	7.1325%
$\mathbb{S}$	-0.3397
$\mathbb{K}$	2.6308

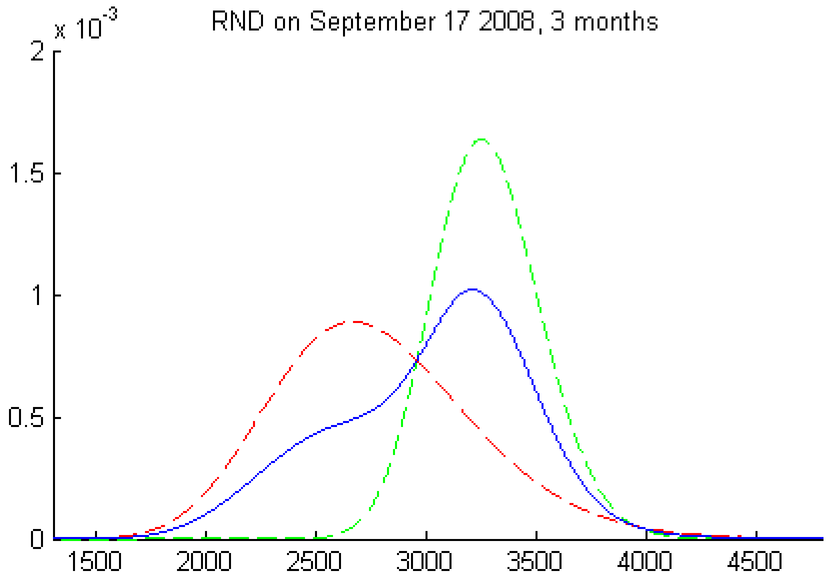

September 17 2008	
Parameters	
μ_1	8.0928
σ_1	0.0747
μ_2	7.9178
σ_2	0.1656
θ	0.4760
Descriptive statistics	
$\mathbb{E}$	3019.6309
$\mathbb{D}$	450.1752
$\%\mathbb{D}$	7.5677%
$\mathbb{S}$	-0.3145
$\mathbb{K}$	2.8603

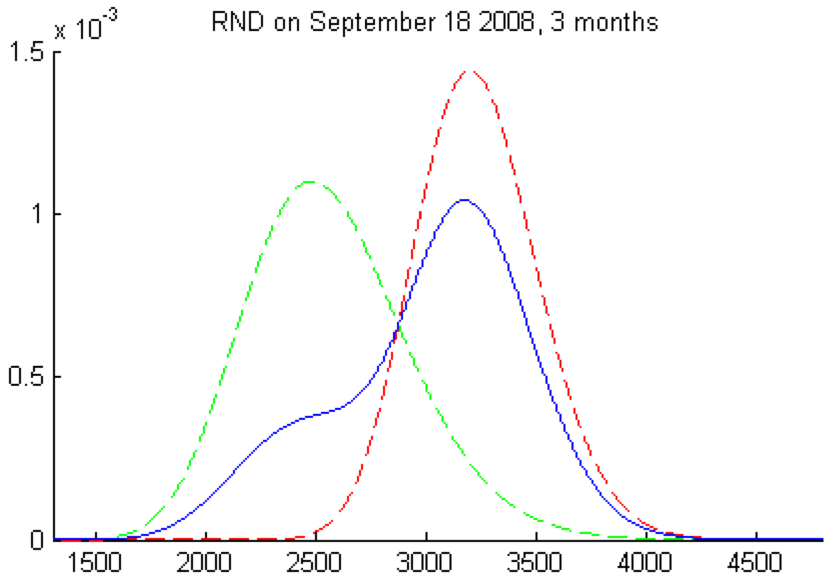

September 18 2008	
Parameters	
μ_1	7.8380
σ_1	0.1447
μ_2	8.0786
σ_2	0.0862
θ	0.3344
Descriptive statistics	
$\mathbb{E}$	3010.9987
$\mathbb{D}$	446.9928
$\%\mathbb{D}$	7.5190%
$\mathbb{S}$	-0.4330
$\mathbb{K}$	2.7907

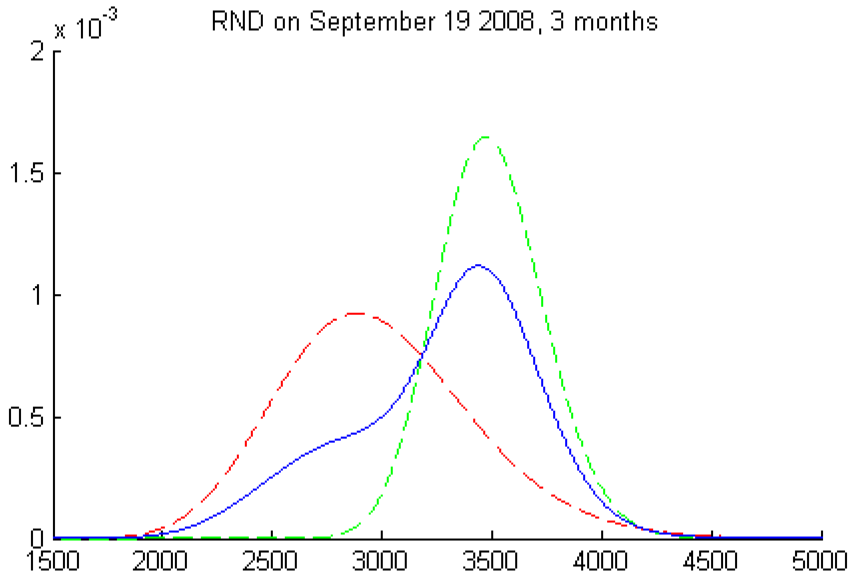

September 19 2008	
Parameters	
μ_1	8.1574
σ_1	0.0697
μ_2	7.9912
σ_2	0.1478
θ	0.5599
Descriptive statistics	
$\mathbb{E}$	3272.9015
$\mathbb{D}$	429.0790
$\%\mathbb{D}$	6.6212%
$\mathbb{S}$	-0.4670
$\mathbb{K}$	3.0008

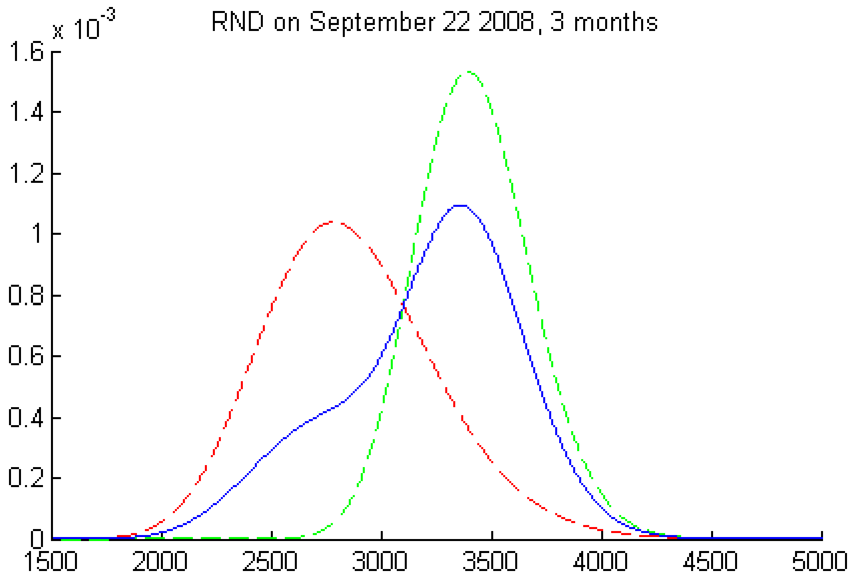

September 22 2008	
Parameters	
μ_1	8.1346
σ_1	0.0765
μ_2	7.9499
σ_2	0.1368
θ	0.6178
Descriptive statistics	
$\mathbb{E}$	3206.9893
$\mathbb{D}$	418.5328
$\%\mathbb{D}$	6.4954%
$\mathbb{S}$	-0.4306
$\mathbb{K}$	2.8884

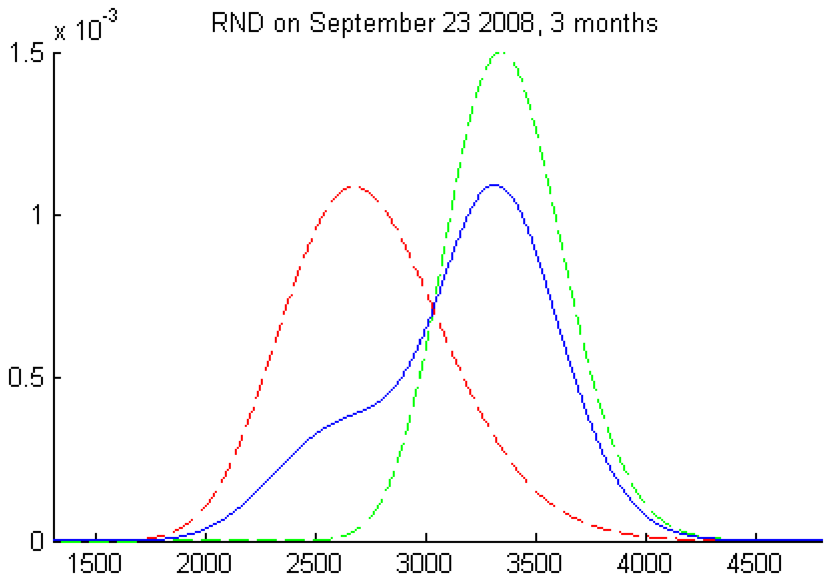

September 23 2008	
Parameters	
μ_1	8.1204
σ_1	0.0794
μ_2	7.9110
σ_2	0.1360
θ	0.6600
Descriptive statistics	
$\mathbb{E}$	3162.0049
$\mathbb{D}$	426.5174
$\%\mathbb{D}$	6.6701%
$\mathbb{S}$	-0.4602
$\mathbb{K}$	2.8795

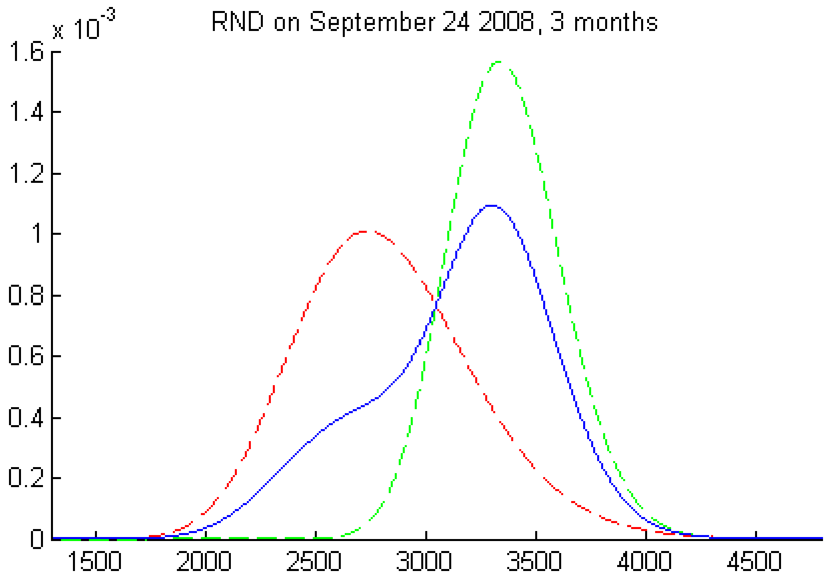

September 24 2008	
Parameters	
μ_1	8.1182
σ_1	0.0762
μ_2	7.9356
σ_2	0.1428
θ	0.5922
Descriptive statistics	
$\mathbb{E}$	3143.9795
$\mathbb{D}$	420.3043
$\%\mathbb{D}$	6.5773%
$\mathbb{S}$	-0.4255
$\mathbb{K}$	2.8793

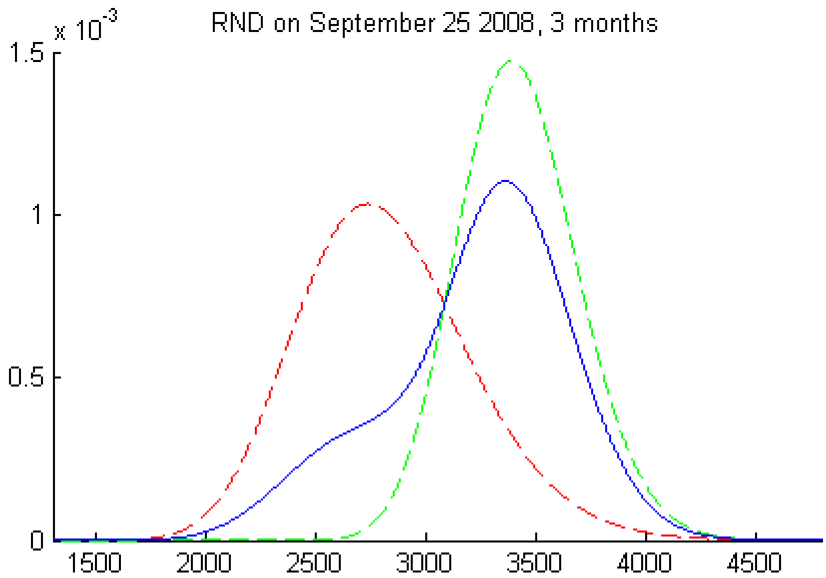

September 25 2008	
Parameters	
μ_1	8.1350
σ_1	0.0797
μ_2	7.9354
σ_2	0.1396
θ	0.6744
Descriptive statistics	
$\mathbb{E}$	3226.9978
$\mathbb{D}$	424.9442
$\%\mathbb{D}$	6.4316%
$\mathbb{S}$	-0.4855
$\mathbb{K}$	3.0185

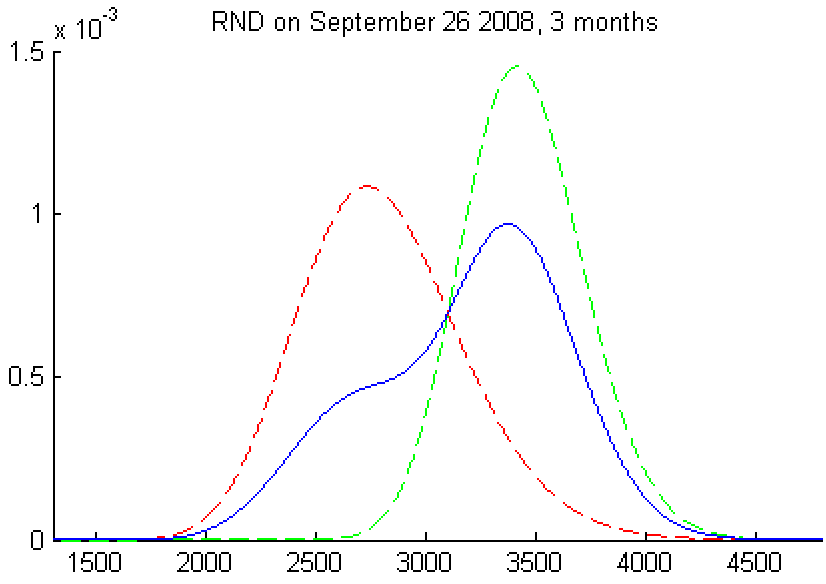

September 26 2008	
Parameters	
μ_1	8.1429
σ_1	0.0801
μ_2	7.9316
σ_2	0.1335
θ	0.5837
Descriptive statistics	
$\mathbb{E}$	3182.9994
$\mathbb{D}$	451.3765
$\%\mathbb{D}$	6.9007%
$\mathbb{S}$	-0.3175
$\mathbb{K}$	2.5986

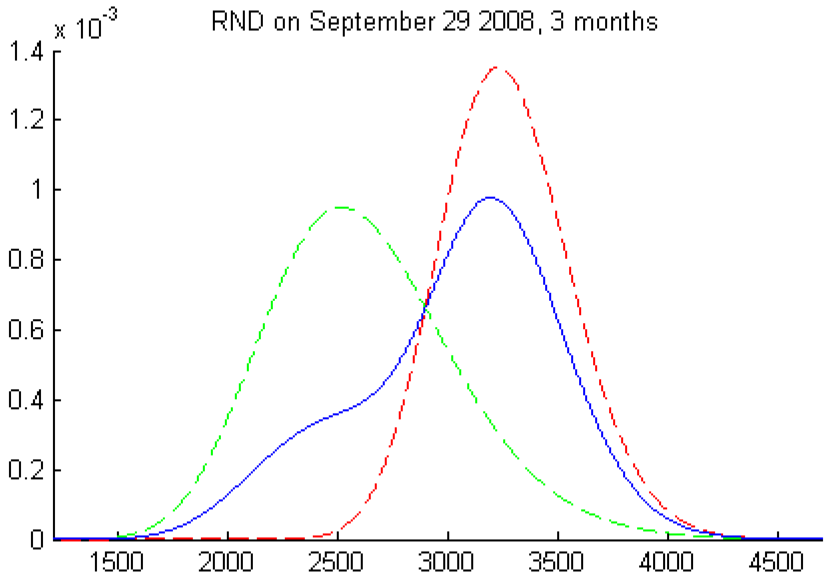

September 29 2008	
Parameters	
μ_1	7.8590
σ_1	0.1643
μ_2	8.0883
σ_2	0.0910
θ	0.3616
Descriptive statistics	
$\mathbb{E}$	3036.2530
$\mathbb{D}$	470.2059
$\%\mathbb{D}$	7.4087%
$\mathbb{S}$	-0.4071
$\mathbb{K}$	2.8672

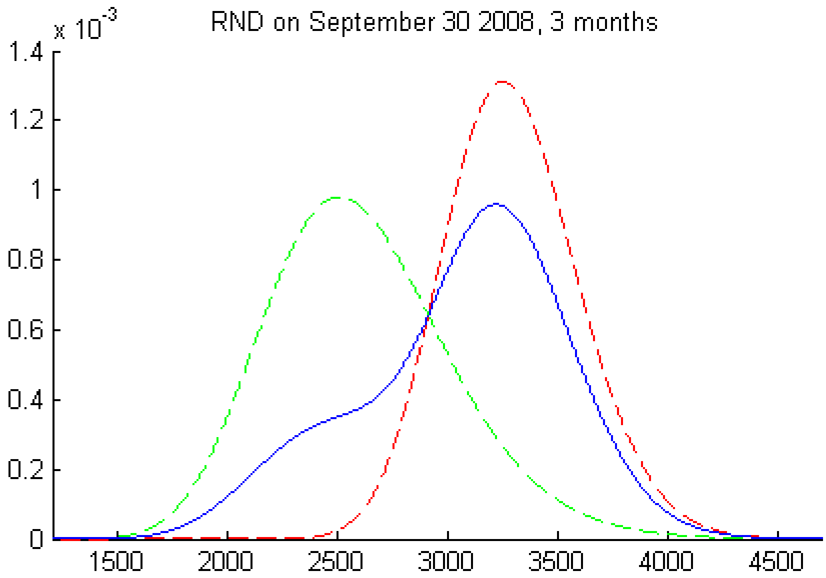

September 30 2008	
Parameters	
μ_1	7.8542
σ_1	0.1600
μ_2	8.0962
σ_2	0.0931
θ	0.3410
Descriptive statistics	
$\mathbb{E}$	3062.0143
$\mathbb{D}$	477.8492
$\%\mathbb{D}$	7.4092%
$\mathbb{S}$	-0.3978
$\mathbb{K}$	2.8409

Appendix C – Data cleaning

In Table C1 below, the number of options removed in every filter of the cleaning procedure is shown. Note that we perform the data cleaning procedure step by step, meaning that once an option has been caught in a filter, it is eliminated and not examined in subsequent filters. Thus, the options removed in a filter are those that do not meet the conditions of that particular filter, but that did not breach the conditions of any of the previous filters. Obviously, this does not affect the final data set, but it does affect the interpretation of the numbers, as these would be different if the filters were applied in a different order.

Table C1 - Cleaning procedure

Step	Filter	Options removed	Options left
0	Initial number of options	-	490508
1	Remove options with no traded volume and/or open interest	367159	123349
2	Remove options with less than six days to maturity	3531	119818
3	Remove options with negative bid and/or ask price	0	119818
4	Remove options where bid $>$ ask	0	119818
5	Remove options where bid and/or ask $> S_t$	9	119809
6	Remove options where bid and/or ask $< \max(S_t - K, 0)$	39033	80776
7	Remove options where ask/bid > 1.2	5602	75174
8	Remove options with bid and/or ask < 0.1	3514	71660
9	Remove put options	26546	45114
Total		**445394**	**45114**