

Quantifying Option Implications

by

Michael Bauer, Xiaowen Chang, and Michael Conway

Abstract

We introduce relevant financial concepts and describe how mathematical tools can be used to extract information about the market's expectations and risk preferences from daily, observable options market prices on the S&P 500 index. This information takes the form of a probability density function, known as the Risk-Neutral Density (RND). Assuming no prior knowledge, we introduce our major tools, including splines and the Generalized Extreme Value (GEV) Distributions, and show how they can be used in a financial context. Finally, we illustrate some of the applications of the RND.

We would like to thank Lafayette College, Professor Qin Lu, Professor Jeffrey Liebner, and the National Science Foundation for their contributions.

1. Introduction

Market prices, containing the knowledge of billions of people, guide nearly every human decision. They dictate what we treasure and what we readily discard. Prices facilitate trade, which leads to specialization, and thus technological advancement. The significance of an asset's price is intuitive, but the prices of options—popularly traded financial derivatives—on that asset contain much more information. Options on the S&P 500 index can tell us about the market's expectations, future prices, and people's risk preferences—much more than can be learned from the index price alone. This information can be quantified in the form of a probability density curve, called the Risk-Neutral Density (RND). This paper explores how to extract the RND from market options prices in the most informative way possible and demonstrates some of its many useful applications.

In Section 2, “Financial Background,” we introduce the S&P 500 stock index, call and put options, and the basics of financial derivative valuation. Section 3, “RND,” summarizes the literature relevant to the adopted approach. Section 4, “Splines,” introduces splines in theory and how we use them. Section 5, “GEV,” discusses the Generalized Extreme Value (GEV) distribution family and its role in our research. Section 6, “Data,” briefly details the data for our study. We conclude by describing our own contributions to the literature.

2. Financial Background

The S&P 500 is what is known as a *stock index*. Its value is a weighted average of the stock prices of 500 different US companies. This number fluctuates up and down throughout the trading day, just like a stock. The S&P 500 value is generally regarded as the proxy for the US stock market. When the market is up in value, the S&P 500 is up, and when the market drops, so does the S&P 500. We are only concerned with the index's daily close price for our research, so when we refer to the January 22, 1996 value of the S&P 500, we are referring to the value at the close of trading. This value generally ranges from \$400 to \$2000 in our research data which spans from January 1996 to August 2013.

A *financial derivative* is any asset that derives its value from another underlying asset.

An *option* is a type of derivative. Many different varieties of options are traded, but for our research we focus on plain European options. Each option contract specifies an underlying asset such as the S&P 500, a strike price (X), whether the option is a call or a put, and an expiration date (T). Note that we assume the starting time 0, so T is also the amount of time to expiration.

We begin with calls. The owner of a *call option* has the right, but not the obligation, to buy the underlying security for the specified strike price (X) on expiration day (T) from the seller of the call option. We denote the price of the underlying security on expiration day as S_T . If S_T is greater than X , the owner of the call can buy the underlying security for price X from the seller of the call option and sell it in the market for S_T , making a

profit of $S_T - X$. However, if S_T is less than X , the owner will choose not to exercise the option and it will expire as worthless. So a call option gives the owner a payoff of $\max(S_T - X, 0)$ on expiration day. It should be noted that options are a zero sum game. In the case of a call, if S_T is greater than X , whoever sold the option is obligated to sell the underlying asset for the price X .

A put is slightly different. The owner of the *put option* has the right to sell the underlying security for X on expiration day, T , to the seller of the put option. Certainly, this right will only be exercised if S_T is less than X . Therefore, a put gives the owner a payoff of $\max(X - S_T, 0)$ on expiration day.

A call gives the owner unlimited upside and no downside at the expiration date, so naturally it should cost a premium for the owner at initialization time 0. The premium of the call paid by the owner to the seller is called the price of the call. Similarly, a put gives the owner limited upside and no downside at the expiration date, so naturally it should cost a premium for the owner at initialization time 0. The premium of the put paid by the owner to the seller is called the price of the put. In general, the price of a financial asset is its expected future payoff, discounted for time and riskiness.

We discount the payoff for time because a dollar today is worth more to an investor than a dollar in 3 months. If an investor has a dollar today, he can put it in the bank for three months, which will return his original dollar plus interest in 3 months.

We discount for risk because most investors would prefer a guaranteed return, so they must be compensated for the possibility that an asset expires as worthless. Discounting for risk is more difficult because every investor has a different risk tolerance.

Many options are settled in cash at expiration. In the case of a call, this means that when S_T is greater than X , the owner of the option receives $S_T - X$ in cash, rather than having to buy the underlying asset for X and sell it for S_T . This is the case for the options on the S&P 500, which we consider in our research.

For example, suppose today is February 2nd, the S&P 500 is at \$1600, and you think the market is going to go up in the next 2 months; you might choose to buy a call option with a strike price $X = \$1610$, expiring on April 2nd. This is called *speculating*.

If the market does go up, and the S&P 500 value on April 2nd, S_T , proves to be \$1700, you will exercise the option and realize a profit of $\$(1700 - 1610) = \90 .

If however the S&P 500 drops to \$1540 on April 2nd, you simply won't exercise the call option.

Options are also used for *hedging*. Suppose again that today is February 2nd, the S&P 500 is at \$1600, and you own a diverse portfolio of stocks. If the value of your portfolio is highly correlated with the value of the S&P 500, you might choose to hedge, or protect, your portfolio against a market drop. Here, you could buy a put with strike price $X = \$1590$ expiring on April 2nd. If the S&P 500 drops to $S_T = \$1300$ on April 2nd, your portfolio goes down along with the index; however, you will exercise the put option for a payoff of $(1590 - 1300) = \$290$ and balance the loss from the portfolio. If the S&P 500

stays strong, finishing at \$1600 on April 2nd, the option expires as worthless, but it is likely your portfolio has remained strong, too. The put price is your hedge expense.

Since options are useful, determining fair prices for options is important. The *Risk-Neutral Principle* is one of the cornerstones of option pricing and financial derivatives pricing in general. It tells us that the price of an option should be its expected payoff in a risk-neutral world, discounted for time by the corresponding risk-free interest rate (Usually US Treasury bill/note/bond interest rates or London Bank Offer Rate, LIBOR for short, can be regarded as the risk-free interest rate). In a risk-neutral world, all investors are indifferent to risk. For example, given the choice between an asset that pays a sure \$1, or an asset with a 50% chance of paying \$2 and a 50% chance of paying nothing, a risk-neutral investor will be indifferent. That is, he only cares about the expected payoff, which is \$1 for both assets in the example.

By the risk-neutral principle, the formula for the price of a call option (C) is:

$$C = e^{-rT} \int_X^{\infty} (S_T - X) f(S_T) dS_T \quad (1)$$

Here, e^{-rT} is the discount factor for time where T is the time to expiration, and r is the risk-free interest rate, which is the interest paid on a US Treasury bill for the same period of time or the LIBOR rate. Note that in a risk-neutral world, all investments should earn the risk-free rate. In the formula, $f(S_T)$ is the risk-neutral density, RND for short, which assumes that the underlying asset grows at the risk-free rate. Intuitively, for each possible value the underlying asset S_T can take within the interval $[S_T, S_T + \Delta S]$ on expiration day, it assigns the probability of $f(S_T)\Delta S$ in a risk-neutral world. Essentially,

the risk-neutral density combines investors' subjective probability beliefs with their risk preferences.

Similarly, for a put option we have:

$$P = e^{-rT} \int_0^X (X - S_T) f(S_T) dS_T \quad (2)$$

Investment professionals generally use the famous Black-Scholes option pricing formula, assuming $f(S_T)$ is a lognormal density function which is equivalent to the stock prices following Geometric Brownian Motion. We will not discuss Geometric Brownian Motion here since most undergraduate financial mathematics textbooks cover the material. In particular, we recommend Hull's book for details. The Black-Scholes formula has an output of call/put price and has 5 inputs:

S_0 : the price of the underlying asset today

X : the strike price

r : the risk-free interest rate

T : amount of time to expiration

σ : volatility of the underlying asset

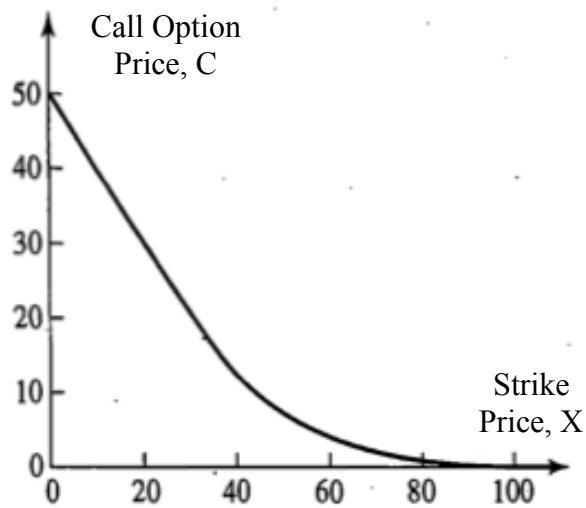
The volatility estimates how much the price of the underlying asset fluctuates; it is the only input of the Black-Scholes formula that isn't directly observable. Even though the historical volatility can be calculated by annualizing the standard deviation of daily Log returns for a period of time, the volatility we need here is the future projected volatility since options are concerned about future prices.

By solving the previous integral assuming lognormal density function with above parameters, the Black-Scholes formula for call and put option prices can be derived as follows:

$$\begin{aligned}
 C &= S_0 N(d_1) - X e^{-rT} N(d_2) \\
 P &= X e^{-rT} N(-d_2) - S_0 N(-d_1) \\
 d_1 &= \frac{\ln\left(\frac{S_0}{X}\right) + \left(r + \frac{\sigma^2}{2}\right)T}{\sigma\sqrt{T}} & d_2 &= \frac{\ln\left(\frac{S_0}{X}\right) + \left(r - \frac{\sigma^2}{2}\right)T}{\sigma\sqrt{T}} \quad (3)
 \end{aligned}$$

Note that in the formula $N(x)$ denotes the standard normal cumulative distribution function. Now that we have a way to price calls and puts, we are interested in analyzing call/put price as function of strike price X while assuming other parameters are fixed. This is because for a fixed asset with S_0 and certain projected volatility σ , a fixed expiration time T and risk-free interest rate r , there are many contracts with different strikes existing. Therefore, we can generate a call curve and a put curve with call/put price as a function of strike price. We begin with the call curve $C(X)$. For a given underlying asset with price S_0 today, expiration date T , risk-free rate r , and volatility σ , we can generate the call curve as a function of strike price X . See Figure 1.

Figure 1.




Clearly, the call curve is decreasing as the strike price X increases. This is intuitive, as a call with strike price \$40 gives you the right to buy the underlying asset for \$10 less than a call with strike price \$50. We can do the same to produce a put curve—which is increasing by similar logic.

An interesting consequence of the Black-Scholes model is the concept of *implied volatility*. Besides volatility, all of the inputs into Black-Scholes are readily observable. Since the market price of the option is observable as well, we can solve for the volatility such that the Black-Scholes model outputs the market price. For example, suppose today the S&P 500 index is at \$1638.71 and the risk-free rate is 5%. We observe a put on the S&P 500 with X =\$1615 expiring in 21 days that is currently trading for \$16. The volatility that yields \$16 when inputted into Black-Scholes is called the implied volatility, which is 18%.

Implied volatility reveals an inconsistency between Black-Scholes prices and options market prices. Every day, for a fixed underlying asset such as the S&P 500 index, for each traded expiration date T , there are at least 20 different strike prices X trading for options on the S&P 500 index. We would think that each of these different options would imply the same volatility since they all have the same underlying asset, time to expiration T and interest rate environment, but in actuality the different strikes give different implied volatilities. When we plot implied volatilities σ against the strike price X , we get a downward sloping curve, referred to as the volatility skew. See Figure 2. The existence of the skew implies that something must be wrong with the Black-Scholes model, particularly its assumption of a lognormal RND of the underlying asset price S_T (Hull).

Figure 2.



3. Risk-Neutral Density

Given Black-Scholes' inconsistency, there has been a lot of interest in generating a Risk-Neutral Density (RND) that is consistent with observed options prices. There are a few different methods for doing this, but we prefer to take a nonparametric approach—proposed by Breeden and Litzenberger (1978). They showed that if we have a smooth call/put curve, then we can use simple calculus to derive the RND. We only illustrate the call here; the results for put are similar.

We know that the price of a call is its expected payoff under the RND, discounted for time:

$$C = e^{-rT} \int_X^{\infty} (S_T - X) f(S_T) dS_T$$

Taking the derivative with respect to strike price gives:

$$\begin{aligned} \frac{\partial C}{\partial X} &= \frac{\partial}{\partial X} [e^{-rT} \int_X^{\infty} (S_T - X) f(S_T) dS_T] \\ \frac{\partial C}{\partial X} &= e^{-rT} [-(X - X) f(X) + \int_X^{\infty} -f(S_T) dS_T] \\ \frac{\partial C}{\partial X} &= -e^{-rT} \int_X^{\infty} f(S_T) dS_T = -e^{-rT} [1 - F(X)] \end{aligned}$$

Here $F(X)$ is the Risk-Neutral CDF. Rearranging terms gives:

$$F(X) = e^{rT} \frac{\partial C}{\partial X} + 1 \quad (4)$$

Taking the derivative again gives the RND:

$$f(X) = e^{rT} \frac{\partial^2 C}{\partial X^2} \quad (5)$$

Clearly a call curve implies a specific RND.

In our research, this means that we need S&P 500 options prices for every possible strike price in the interval $[0, \infty)$. In reality, we don't have a well defined call curve in the interval $[0, \infty)$.

Given a dense set of strike prices in increasing order $\{X_1, X_2, X_3, \dots, X_n\}$ and their associated call prices observed in the market, $\{C_1, C_2, C_3, \dots, C_n\}$ we can approximate the second derivatives and find the implied density:

$$f(X_i) = e^{rT} \frac{C_{i+1} - 2C_i + C_{i-1}}{(\Delta X)^2}, \quad i = 2, 3, \dots, n - 1 \quad (6)$$

Observe that this method gives the option implied RND values at all but the highest and lowest traded strike prices—so given n prices for different strikes, we can get RND values at $n-2$ points. Certainly, we expect the RND to be a continuous probability density function. Following this approach and interpolating between points doesn't sound awful, but in actuality it works very poorly. Most prices contain a lot of market noise, which can create characteristics that are undesirable in a probability density; in some days where there is very little trading volume, this approach might even give negative probability density values or unwanted choppiness in the RND. Clearly, it is a problem that options are only traded at a discrete set of strike prices—we address this in the next section.

Fitting the tails of the distribution is another issue that must be addressed. Even if we do interpolate between the points we have, we need to find a way to assign probabilities to the values between 0 and the 2nd strike and between the $n-1$ st strike and infinity. This issue is addressed in Section 5, where we elect to fit a generalized extreme value distribution to the tails.

4. Splines

We adopt Figlewski's approach (2010), which advocates a few changes to Shimko's method. Our contribution is empirical analysis. Figlewski has analyzed data from January 4, 1996 to February 20, 2008. We extend the data to August 2013 which includes the market crash and recover. Figlewski gave very few analyses other than the proposed method. In order to follow Figlewski's approach, we need to understand spline functions. A spline is a piecewise-defined polynomial curve. It is used to fit a curve to a number of points, $\{(x_1, y_1), (x_2, y_2), \dots, (x_n, y_n)\}$. There are many different types of splines, determined by their various characteristics. The first characteristic of the spline is its order—this is the degree of the polynomials used. Next is the number of knots—points where the different polynomials connect. For example, a cubic spline with 8 knots consists of 9 different cubics, each connecting to the next at one of the 8 knot points.

For a spline to be useful it must be close to the points it is fitting. Thus, it is useful to minimize the squared difference between the spline's value at every x_i , $1 \leq i \leq n$, and the value y_i . If we care more about the spline's fit at some points than others, we can use a weighting function. This assigns a weight to each point. The spline then minimizes the weighted squared differences.

On a given day, we have between 20 and 40 strike prices with their corresponding option prices, and we want to fit a curve to them, so a spline sounds like the natural solution. Unfortunately, the simple solution is not always the best solution. Options prices contain

a lot of market noise, which can produce undesirable affects in the RND, such as negative probabilities.

Shimko (1993) suggested calculating the options' implied volatilities (IV's) for these 20-40 options, including both calls and puts, and considering the implied volatilities as a function of strike price, fitting a spline to the IV's, and then translating the IV spline into a call/put curve. While an option's price may fluctuate rapidly throughout the day, its IV remains fairly constant. In addition, market call/put prices can range from \$0 to as high as \$1000, whereas IV's tend to remain between 10% and 40%, so we interpolate across a much smaller range. Therefore, transforming into IV space significantly cuts down on the impact of market noise. The Black-Scholes model is initially used in reverse to calculate the market IV's and later to transform the IV spline back into a call/put curve. It should be noted that we are simply using the Black-Scholes formula as a transformation, which does not require us to make any of its assumptions. So, the RND we get is still model-free.

As in Figleweski's approach, we fit a fourth order spline to the IV function, with one knot located at $X=S_0$, with a special weighting function. Each option has a quoted bid price at which people are willing to buy, and an ask price at which people are willing to sell. When we refer to an option's price, we refer to the midpoint of its bid and ask prices. For the weighting function, it assigns more weight to fitted points falling outside the bid-ask spread. The calculation process is an iterated process. This approach cuts down on the

effect of market noise while still extracting as much information as possible from the market prices. See Figure 3 for the IV's and Figure 4 for the curve fitting.

Figure 3.





Figure 4.



After converting the IV spline to a call/put curve, we can follow the Breeden and Litzenberger approach (equation 5) to get the RND values between the highest and lowest traded strike prices. See Figure 5 for a call curve and Figure 6 for the RND.

Figure 5.

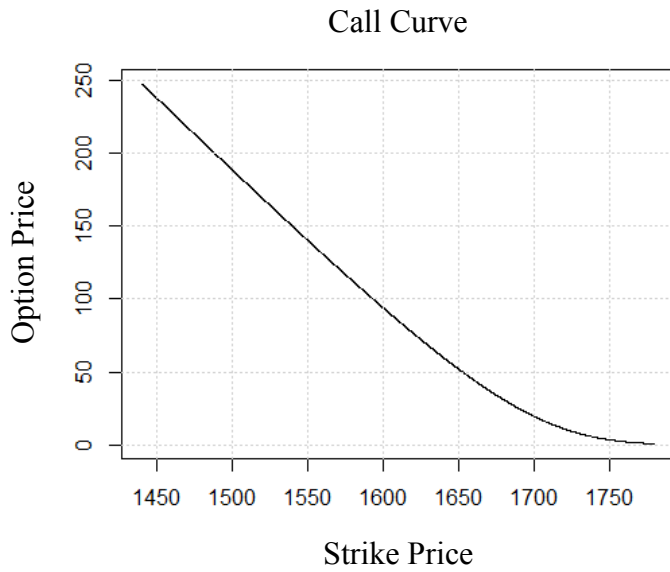
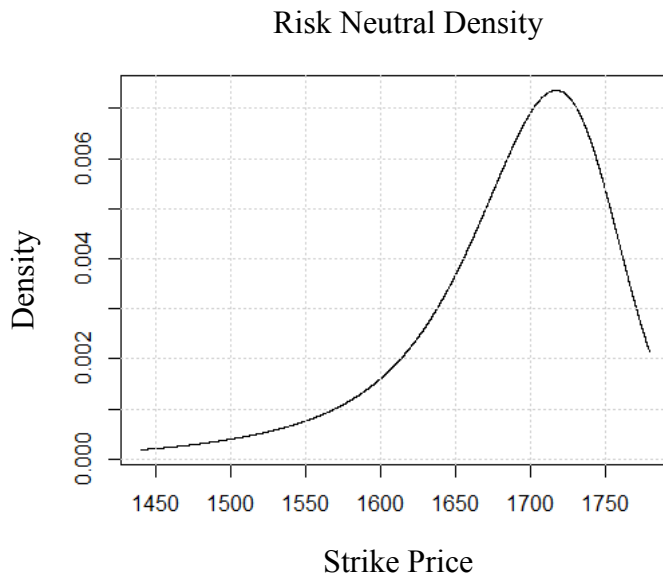


Figure 6.



Of course, it is possible that the S&P 500 could take on a value below the lowest or above the highest traded strike—otherwise people would have no reason to buy a put with the lowest X or a call with the highest X . We must find a way to assign probabilities to these possible index values, which we refer to as the left and right tails.

5. GEV

Figlewski advocates fitting Generalized Extreme Value (GEV) distributions to the left and right tails. The GEV has 3 parameters; μ , the location parameter, σ , the scale parameter, and ξ , the shape parameter. It is a flexible family of distributions, with 3 possible forms dictated by the value of the shape parameter. Its distribution function is:

$$F(z) = \exp [-(1 + \xi z)^{-1/\xi}]$$

$$z = \frac{S_T - \mu}{\sigma}$$

If $\xi = 0$, we get a Gumbel distribution—with a tail similar to that of a normal distribution. If $\xi > 0$, we get a Fréchet distribution with its characteristic fat tail extending out to ∞ . If $\xi < 0$, we get the Weibull distribution, which is the only family member with a tail that does not extend out to infinity.

For the right tail, we choose two points just below the highest traded strike price. We fit a GEV subject to the conditions that its density values at these two points equal the empirical RND's values at those points, and that the GEV's CDF value at the lesser of the two strike prices matches the empirical CDF value at that point. We follow a similar procedure for the left tail. See Figure 7 for an example of fitting; the black curve shows the middle part of our RND, the red curve is a GEV function which extends the tail of the black curve to the left and the green curve is a GEV function which extends the tail of the black curve to the right. See Figure 8 for the combined fitting.

Figure 7

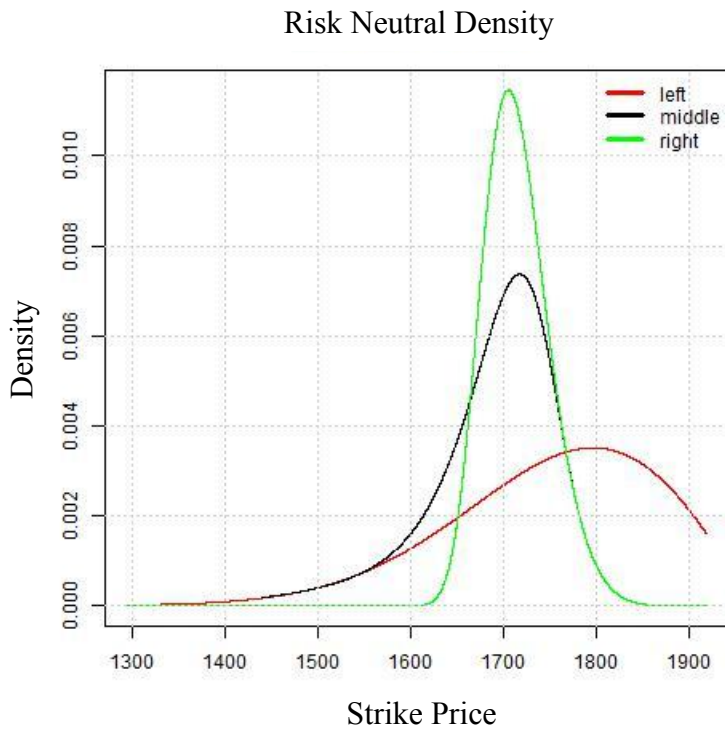



Figure 8.



6. Data

We began with the S&P 500 options data from January 1996 to August 2013. As in Figlewski's method, we focused on major quarterly expirations—those options expiring in March, June, September, and December. As per Figlewski's suggestion, for each trading day, we considered only the nearest expiration date, unless the nearest expiration date was less than 14 days away, in which case we would look at the following expiration's options for those 14 days; this is because the prices are very volatile close to expiration, making it very difficult to distinguish market noise from useful information. After throwing away a small number of days because of poor data due to low trading volume and other issues, we extracted 3800 RND's of S_T . Now that we can define annualized return as $R = \frac{\ln(S_T) - \ln(S_0)}{T}$, we can thus get RND for annualized return, which we will refer to as the Risk-Neutral Return Rate Density (RNRRD); this makes it easier to compare plots from different time periods.

7. Contributions to the Literature

The RND and RNRRD contain information about investors' risk preferences and their subjective probabilities about the S&P 500's value on expiration day. Literature supports that the historical stock market return distribution is not a normal distribution. The historical stock market return distribution is skewed to the left with a fat tail. However, RND and RNRRD show the perspective of people's projection about future returns. The densities' changes from day to day tell us about the market's response to different events. One of the difficulties we noticed with analyzing these plots is that a standard computer can only view a few plots at the same time, which makes it harder to get a big picture view of their dynamic behavior. We began by calculating the mean, standard deviation, skewness, and kurtosis of each RNRRD. To understand the plots' dynamic behavior better, we constructed time series plots of the standard deviation, skewness, and kurtosis.

Figure 9 is the time series of the standard deviations of RNRRD during the 3800 trading days. It represents the uncertainty people have of the future returns. We clearly see the peaks, which represent the market crash. We also see how long the market sentiment evolves. Even though the recent market crash happened in 2007-2008, we see that people continue to worry about the crash during 2009.

Figure 9.




Figure 10 represents the time series of the skewness of the RNRD during 3800 trading days. It is not a surprise that the RNRD are negatively skewed. We can also see that expectations are very negative during 2010-2013.

Figure 10.

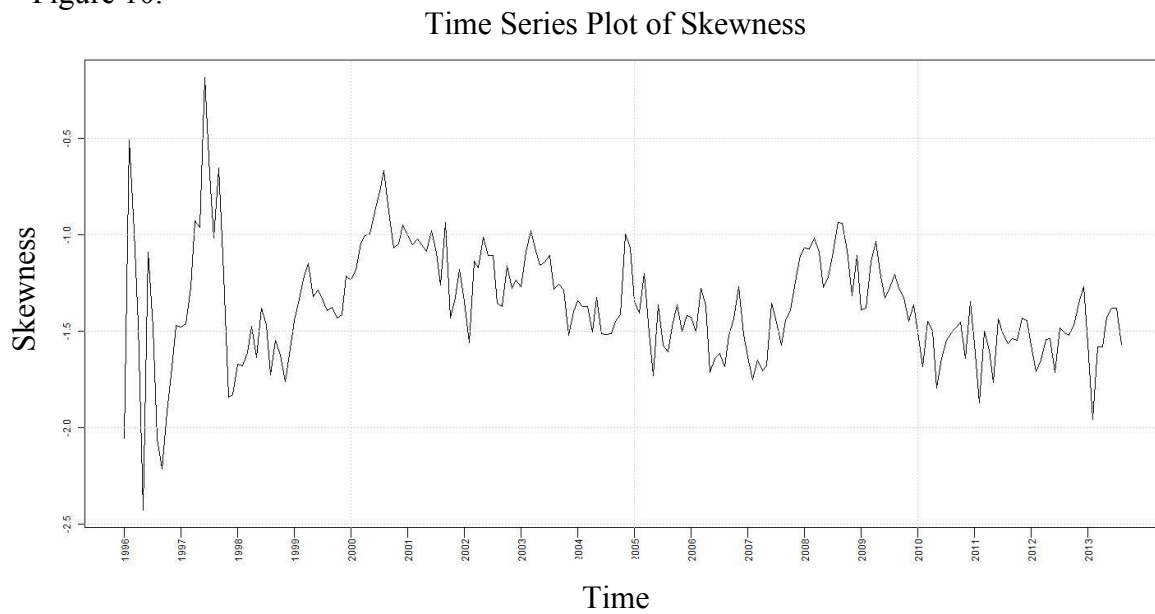



Figure 11 represents the time series of the kurtosis of RNRRD during 3800 trading days. Note that normal distribution has a kurtosis of 3. It is not surprising that RNRRD has a heavier tail than that of the normal distribution. In fact, people had a negative projection until the beginning of 2010.

Figure 11.

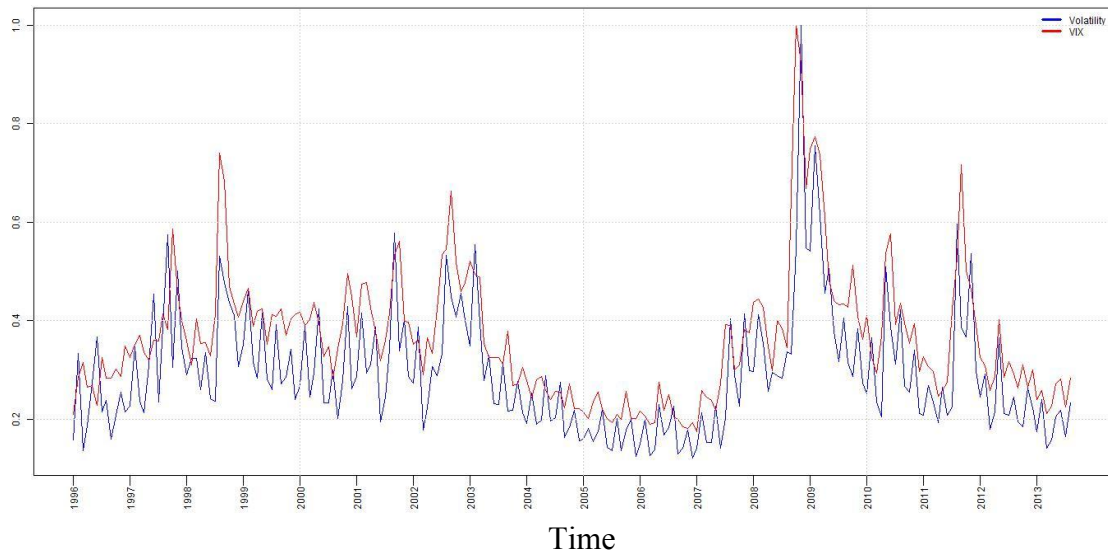
Time Series Plot of Kurtosis



The standard deviation of RNRRD is one measure of implied volatility. The Chicago Board Option Exchange VIX (2003) index is also meant to estimate future volatility. We ran regressions, plotted, and compared the two. We found that our volatility matches the VIX very closely. See Figure 12. Note that in Figure 12 both our volatility and the VIX were adjusted to an interval from 0 to 1 for comparison purposes.

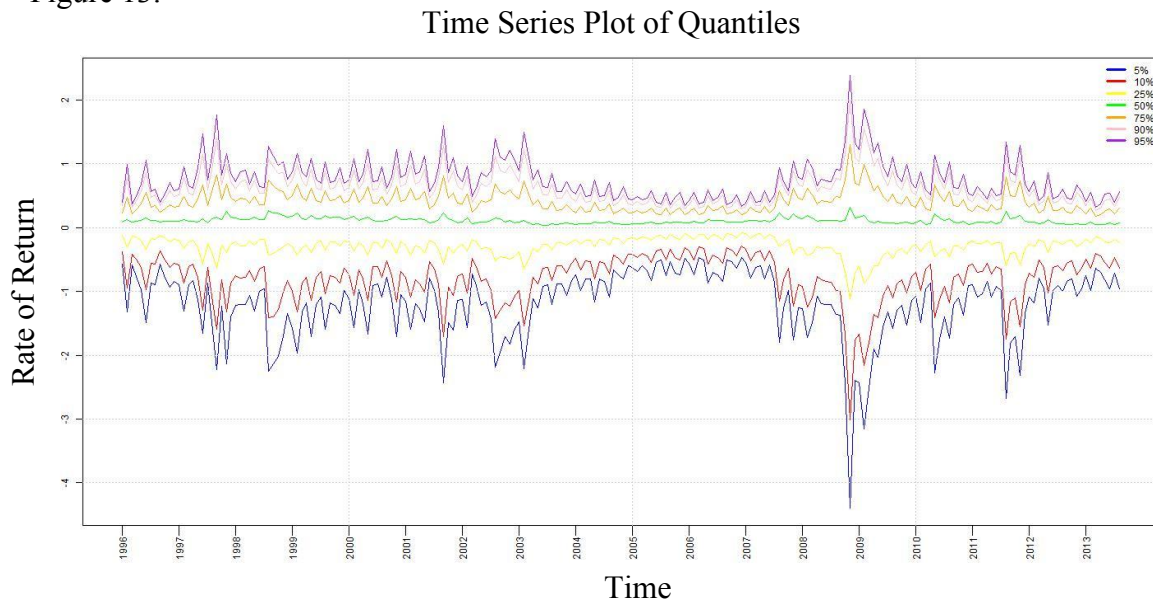
Figure 12.

Time Series Plot of Volatility and VIX



Inspired by Bahra's (1997) paper, we analyze the interquartile range, taking the visual representation a step further. We plotted the RNRRD's 5th, 10th, 25th, 50th, 75th, 90th, and 95th quantiles as a daily time series (see Figure 13). We feel that the quantiles give a concise summary of the RNRRD's. The spread of the quantile curves represent people's unsureness about the future. The market's true recovery begins in 2010. These quantiles, together with standard deviation, skewness, and kurtosis time series, contain all of the necessary information to begin an event analysis. This makes the enormous amount of data more tractable. We then used these plots to analyze events that we thought might have an impact on the US market.


Figure 13.



In general, market drops tend to increase the fat tail significantly, possibly indicating investors fear an even larger drop or investors are becoming less tolerant of risk.

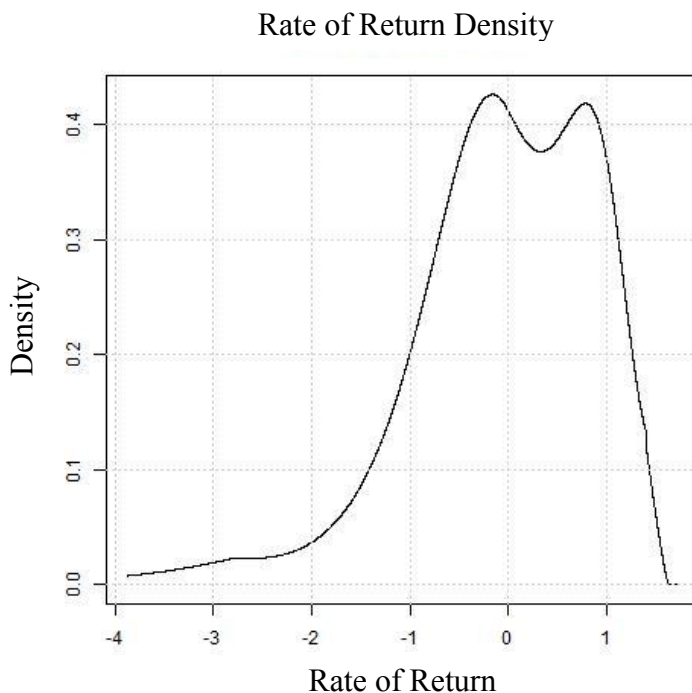
However, in strong bull markets, the left tail becomes less fat, with the RNRD almost resembling a normal distribution (see Figure 14).

Figure 14.



In the most extreme market downturns, a “double-hump” occurred occasionally (see Figure 15). There are many possible explanations for this phenomenon that relate to different combinations of investor optimism or pessimism and risk tolerance. In the future we would be interested in quantifying these to gain a better understanding of their meaning.

Figure 15.



In summary, RND and RNRRD analysis is very useful in the market. From data, we know that the assumptions of the Black-Scholes model are wrong, so we need to use the Black-Scholes formula with caution. For example, the Black-Scholes formula underestimated the left tail probability, so it undervalues the put options, especially the deep out of money put options. (The deep out of money put refers to the put which has a strike price that is way under the underlying asset price.) Also, from RND and RNRRD analysis, we can see when people begin to be optimistic about the market. If we push the

idea further, the RND and RNRRD can be used to price non-standard option contracts. The RND and RNRRD contain valuable information about people's perspective about the future. There is a big unsolved problem in finance, which inquires how one can decompose the RND into real life probability and risk preference. It has been solved partially for the discrete case under some assumptions (Ross 2011). Understanding the behavior of the RND is the first step in solving this problem.

Works Cited

- Bahra, B. "Implied Risk-Neutral Probability density Functions from Options Prices: Theory and Application." Working Paper, Bank of England. (1997).
- Breeden, Douglas and Robert Litzenberger (1978). "Prices of State-Contingent Claims Implicit in Option Prices." *Journal of Business* 51, 621-652.
- Chicago Board Options Exchange. VIX CBOE Volatility Index. (2003). Web. 23 July 2014. <<http://www.cboe.com/micro/vix/vixwhite.pdf>>.
- Figlewski, S. "Estimating the implied risk neutral density for the U.S. market portfolio." *Volatility and Time Series Econometrics: Essays in Honor of Robert F. Engle*. Oxford University Press, Oxford, UK. (2010).
- Hull, John. "Options, Futures, and Other Derivatives." Boston: Prentice Hall (2012).
- Optionmetrics. "Ivy DB File and Data Reference Manual, Version 2.0." 2003. Web. 30 July 2014.
- Ross, Stephen. "The recovery Theorem." *Journal of Finance*, to appear.
- Shimko, David. "The Bounds of Probability." *RISK* 6 (1993): 33-37.