Impact of the Financial Crisis on Risk Aversion: Evidence from Option Prices

Lily Zhang[1]

Submitted to the Department of Economics of Amherst College in partial fulfillment of the requirements for the degree of Bachelor of Arts with Honors

Faculty Advisor: Geoffrey Woglom

April 18, 2012

[1] I would like to thank Professor Geoffrey Woglom for his invaluable help and endless patience – this thesis would not have happened without him. I would also like to thank Eric Weisberg, my friends, and my parents for their support throughout this entire process.

Abstract

This paper seeks to determine whether or not risk aversion has increased following the 2008 financial crisis by examining evidence from the S&P 500 Index and its related European options. Risk aversion is calculated from the subjective probabilities, which describe the perceived likelihood of outcomes, and the risk-neutral probabilities, which are those same probabilities adjusted for the marginal utility at the given level of wealth. The paper investigates existing methods and proposes new approaches for inferring both subjective and risk-neutral probabilities. After evaluating each method, the preferred estimates obtained are used to estimate marginal utility, which is then used to study risk aversion. The paper concludes that risk aversion has apparently decreased after the financial crisis, but this result may be sensitive to some assumptions inherent in the analysis.

Keywords: Risk Aversion, Financial Crises, Option Pricing, Black-Scholes, Kernel Density Estimation

1 Introduction

There is a common conception that risk aversion increases during and immediately following market crashes. This view, if true, would indicate some sort of fundamental shift in the financial markets as a result of a crash. Since traditional economic theory assumes that preferences are stable, a change in risk aversion, as a property of utility functions, may call into question this assumption. This paper seeks to corroborate or refute the existence of a change in risk aversion by examining historical S&P 500 Index and option prices.

Risk aversion can be found as the difference between the subjective probabilities, which describe the perceived likelihood of outcomes, and risk-neutral probabilities, which are those same probabilities adjusted for the marginal utility at the given level of wealth, of future states of the world (i.e., stock prices). Subjective probabilities can be approximated with historical returns or inferred from option prices via the Black-Scholes option pricing formula. Risk-neutral probabilities may be derived directly from observed option prices or through the use of the Black-Scholes formula as well. Each method has its own weaknesses. In particular, the Black-Scholes techniques seem promising, but any analysis based on the Black-Scholes formula is invalid due to an empirical phenomenon – the volatility smile – that the formula produces.

The purpose of this paper is twofold: 1) to explore and evaluate different techniques for inferring subjective and risk-neutral probabilities from historic security and option prices, and 2) to examine the effect of the 2008 financial crisis on risk aversion through the use of the most suitable method. Regarding the first objective, the paper finds that substituting a non-central t distribution in the place of the normal distribution used in the Black-Scholes formula to describe stock returns appears to eliminate the volatility smile problem, and fitting available option prices to this modified formula gives the best estimates of both subjective and risk-neutral probabilities. Using the aforementioned technique, the paper finds evidence of decreased risk aversion following the financial crisis. However, this conclusion must be taken with caution because some of the assumptions needed to estimate risk aversion with option prices may not hold.

The paper proceeds as follows. Section 2 presents the theoretical background behind inferring risk aversion from index and option prices. Section 3 proposes different practical methods for calculating the factors needed to determine risk aversion and discusses their relative merits. Section 4 overviews the data. Sections 5 through 7 report results. Section 8 concludes.

2 Theory

This section outlines the theoretical framework that underlies this paper. The discussion commences with a broad overview of options, subjective probabilities, and risk-neutral probabilities. Option prices are then shown to imply risk-neutral probabilities. Finally, risk aversion determines the difference between subjective and risk-neutral probabilities.

2.1 Options

An option, as a derivative financial product, is a contract that grants the right, but not the obligation, to buy or sell a certain number of shares of a given stock, the *underlying*, at a predetermined *strike price*. The transfer of stock only occurs if the owner of the option chooses to *exercise* it on a specific *expiration* date[2]. A call option is the right to buy stock, while a put option is the right to sell stock. Since it conveys the ability but not the commitment to buy stock, call options are typically only exercised if, at expiry, the price of the stock (also known as the *spot price*) is greater than the strike price. For example, at expiration, a call option with a strike of $25 will have payoffs represented by the following diagram.

Figure 2.1. Payoff diagram for a 25 call.

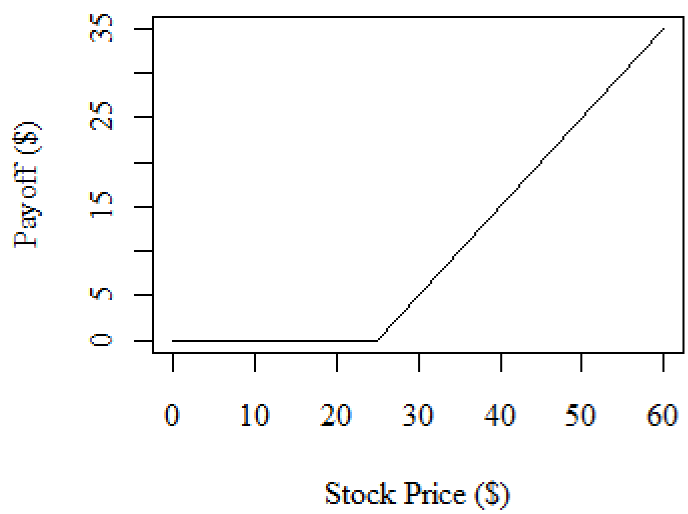

Note that in the above figure, the payoff function of the call option when stock price exceeds $25 appears to be a continuous straight line with slope 1. However, both stock and option prices are actually discrete – they must be in increments of whole cents. Thus, the right half of the payoff diagram is really a step function with steps $0.01 apart. This distinction is important to make, as the later analysis requires discrete states of the world, where each state is a possible stock price.

---

[2] European-style options must be exercised at expiry. American-style options may be exercised any time before expiration as well. This paper is concerned with European options.

This paper only examines call options and so will sometimes use the term "options" to refer to call options; a similar analysis can be conducted with put options.

2.2 Subjective Probabilities

Subjective probabilities refer to the probabilities that one believes to represent the likelihood of the realization of a certain outcome. These are not the same as the *objective probabilities* that describe the actual probabilities of occurrence, which is unobservable for non-experimental data. The following example illustrates the distinction between the two terms: suppose there is an unfair coin that will come up heads 60% of the time. The objective probability distribution is therefore 0.6H/0.4T. However, if the coin is falsely believed to be fair, then the subjective probabilities will be 0.5H/0.5H.

Asset prices, including option prices, are affected by subjective probabilities, not objective ones (which are unknown). However, it should be noted that, under the assumption of perfectly rational expectations, subjective probabilities should be unbiased estimates of objective probabilities.

2.3 Arrow Securities and Risk-Neutral Probabilities

An *Arrow security*[3] is a hypothetical state-contingent security that pays $1 if its specific state of the world occurs in the next time period. If, at t_0, one held an Arrow security for each possible distinct state of the world at t_1, one would be guaranteed a payout of $1 at t_1. This means that, by the no-arbitrage principle, the total sum of the prices of these Arrow securities at t_0 must be $1 + r$, where r is the risk-free rate from t_0 to

[3] Also known as an Arrow-Debreu security

t_1. Equivalently, this is to say that the present discounted value of the sum of all the Arrow securities must be $1.

Further, the discounted price of a single Arrow security is called the *risk-neutral probability*. That is,

$$\frac{P_i}{1+r} = RNP_i \tag{2.1}$$

where P_i is the price of the Arrow security and RNP_i is the risk-neutral probability for state i.

Risk-neutral probabilities are not actual probabilities – the term "probability" is only used to highlight two special properties of risk-neutral probabilities: each risk-neutral probability must be between 0 and 1, and together they must sum to one. However, risk-neutral probabilities are related to the true (that is, subjective) probabilities of the occurrences of their respective states of the world, adjusted by risk preferences:

$$RNP_i = p_i \alpha_i \tag{2.2}$$

where p_i is the subjective probability and α_i is a risk adjustment factor for state i.

The risk adjustment factor α_i is how much the risk-neutral probability has been affected by risk aversion in relation to the subjective probability. For state of the world i, $\alpha_i > 1$ indicates that one is willing to pay a risk premium to purchase the Arrow security[4], and $\alpha_i < 1$ means individuals will pay less than $1 for the Arrow security. In fact, α_i can

[4] This is because in state i, a lower level of wealth obtains and so $1 in state i is worth more than $1 today.

almost be viewed as today's equivalent, in utility terms, of $1 in the future state, given that the state occurs.

While Arrow securities do not exist in real life, they – and, relatedly, risk-neutral probabilities – are still particularly interesting because they allow the pricing of other types of securities, including stocks, options, and more elaborate financial instruments. Any security can be understood as a state-contingent contract that has various pre-defined payoffs based on the state of the world achieved and, thus, can be replicated with a basket of Arrow securities.

For instance, consider a call option with a strike price of $10. This contract will pay nothing if the stock price at the time of expiration is $10 or less. The call option will pay $0.01 if the final stock price is $10.01; it will pay $0.02 if the stock ends up at $10.02, and so on, with the option value increasing one-for-one with the resulting stock price. This call option has the same payout structure as a bundle of Arrow securities:

$$\mathrm{C}_{\$10.00} = 0.01\mathrm{A}_{\$10.01} + 0.02\mathrm{A}_{\$10.02} + 0.03\mathrm{A}_{\$10.03} + \cdots$$

where $\mathrm{C}_{\$10.00}$ is the payout for the call option, and A_i is the Arrow security that pays $1 if the stock finishes at price *i*.

Thus, if the prices of the Arrow security for each possible state of the world are known, then it is possible to calculate the price of any security as a linear combination of the prices of the appropriate Arrow securities. Similarly, it is possible to reverse this calculation and find the prices of Arrow securities using more complicated securities[5].

[5] This, of course, requires that the available securities span all states of the world. In theory, as Gisiger (2010) points out, this is assumed to hold.

The next section demonstrates how call option prices can be used to infer the prices of Arrow securities and, through that, the risk-neutral probabilities.

2.4 Option Prices Imply Risk-Neutral Probabilities

The following example based on Breeden and Litzenberger (1978) shows how the second difference of call prices over different strike prices imply risk-neutral probabilities.

Suppose at t_1 there are five possible states of the world, s_1 through s_5. For simplicity, let the states of possible stock prices be numerically ordered, so that s_2 is one level "higher" than s_1, and so on. At t_0, a discrete "call option" c^1 on the market will pay $1at t_1 for each state higher than s_1. Figure 2.2 illustrates the payoffs for c^1. Another call option contract c^2 that pays off $1 if state s_2 or higher occurs has similar payoffs. This is shown in Figure 2.3.

A third contract $c^{1\prime}$ can be constructed by buying c^2 and selling c^1. Essentially, $c^{1\prime} = c^1 - c^2$, and $c^{1\prime}$ can be seen as a first difference of the original call options. Its price is equal to the price of c^2 minus the price of c^1, and its payoff (see Figure 2.4) can be found by subtracting the payoff of c^2 from that of c^1. By the same method, $c^{2\prime}$ can also be defined as the difference between c^2 and c^3. The payoffs for $c^{2\prime}$ are shown in Figure 2.5.

Finally, another contract $c^{1\prime\prime}$ can be produced by buying $c^{1\prime}$ and selling $c^{2\prime}$. $c^{1\prime\prime}$ may be represented as $c^{1\prime\prime} = c^{1\prime} - c^{2\prime} = (c^1 - c^2) - (c^2 - c^3)$ and is therefore a second difference of the c^ns, where n is a state of the world. This new contract has payoffs according to Figure 2.7.

Figure 2.2. Payoff diagram for c^1.

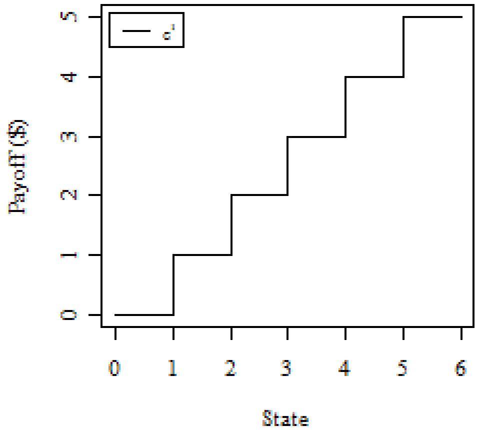

Figure 2.3. Payoff diagram for c^2.

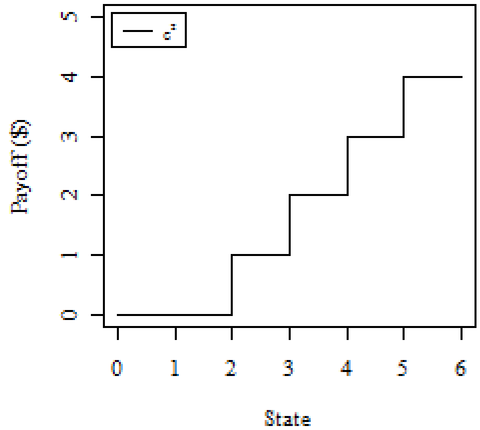

Figure 2.4. Payoff diagram for $c^{1\prime}$.

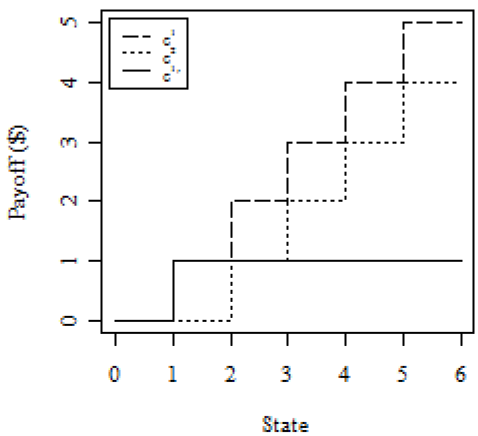

Figure 2.5. Payoff diagram for $c^{2\prime}$.

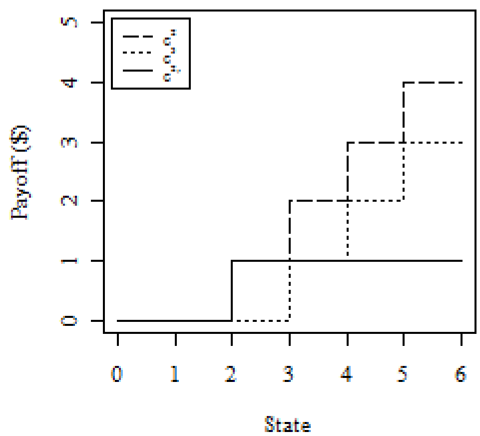

Figure 2.6. Payoff diagram for $c^{1\prime\prime}$.

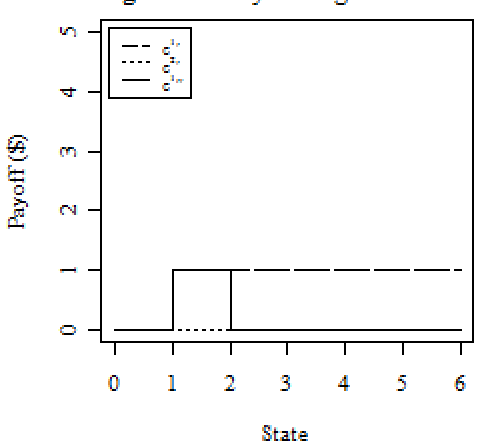

Figure 2.7. RNPs.

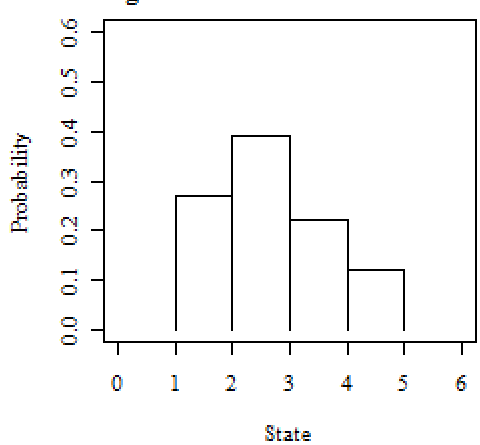

Clearly, $c^{1\prime\prime}$ pays \$1 if and only if state s_1 occurs at t_1. Thus, its price – which is the second difference of the original call option prices – is equal to the probability of s_1 adjusted for risk. That is to say, $c^{1\prime\prime}$ represents the price of the Arrow security of s_1. Other $c^{n\prime\prime}$s may be calculated in the same way to find the individual prices of the complete set of Arrow securities. By equation (2.1), these prices can be discounted to the present via the risk-free rate to produce the risk-neutral probabilities of all the possible states at t_1, which necessarily sum to one. Example risk-neutral probabilities are shown in Figure 2.7.

2.5 Deriving Risk Aversion from Risk-Neutral and Subjective Probabilities

As shown in equation (2.2), risk-neutral probabilities can be expressed as the product of subjective probabilities and an adjustment for risk aversion. Rearranging that equation gives risk-adjustment as a function of both probabilities:

$$\alpha_i = \frac{RNP_i}{p_i}, \tag{2.3}$$

Recall from the brief discussion earlier that risk-adjustment is also related to the utility function. More formally, as Jackwerth (2000) shows,

$$U'(i) = \left(\frac{\lambda}{1+r}\right)\left(\frac{RNP_i}{p_i}\right),$$

where λ is a constant representing the shadow price of the budget constraint and *U'(i)* is the marginal utility of wealth in state *i*. Manipulating this equation and substituting in α_i from equation (2.3) produces

$$\alpha_i = \left(\frac{1+r}{\lambda}\right)U'(i).$$

This final equation demonstrates that α_i – which, by equation (2.3), is also the ratio of the risk-neutral probability to the subjective probability – is actually directly proportional to marginal utility of the equivalent level of wealth.

3 Methodology

3.1 Finding Subjective Probabilities

Equation (2.3) indicates that estimates of subjective probabilities are needed in order to infer risk aversion from risk-neutral probabilities. This paper uses two methods to estimate the subjective probability distribution for the time periods studied. They are introduced in the following sections along with both their advantages and shortcomings.

3.1.1 Kernel Density Estimation with Historical Returns

One simple way to estimate the subjective probability distribution is to find a kernel density estimate of historical returns over a period of time (Jackwerth (2000)). Kernel density estimation is a technique for calculating a probability density function based on frequency data. One of its advantages lies in the fact that it does not presuppose a functional form for the final distribution. In fact, it requires only two parameters: a kernel function and a bandwidth. This paper will use the normal distribution as the kernel function and will select a bandwidth via a calculation proposed by Scott (1992)[6]:

$$bandwidth = 1.06 \frac{\hat{\sigma}}{\sqrt[5]{n}}$$

[6] This method is a commonly-used variant of the Silverman rule-of-thumb (Silverman (1986)), which used 0.9 as the multiplier instead.

where $\hat{\sigma}$ is the sample standard deviation and n is the sample size.

The benefit to using this method is that it is easy to put into practice. Historical equity prices are easily acquirable. From there, it is only a straightforward calculation to find non-overlapping[7] n-day returns to prepare for the kernel density estimation.

One problem with this method is that relies on past performance being a good indicator of future predictions, which is not necessarily true. More precisely, two assumptions are made: 1) that the objective probability distribution did not change during or since the time the historical samples are taken, and 2) that subjective probabilities coincide with objective probabilities (i.e. rational expectations). Historical returns are based on past objective probabilities, and it is entirely possible that those objective probabilities have changed over time. Even if they have not changed, market participants may believe that they have. Thus, subjective probabilities of future stock returns may not actually be the same as the objective probabilities calculated from actualized returns. This problem may be ameliorated by analyzing returns from a relatively short timeframe immediately preceding the date to which the forecast is being made. The effectiveness of this strategy relies on the assumption that changes in market conditions, perceived or real, occur gradually, and so more recent historical samples will more closely reflect the correct probabilities.

Another potential problem lies in the possibility that the sample of historical returns used is not representative of the actual objective distribution. This could occur if

[7] Non-overlapping data is used to ensure independence. Unfortunately, depending on the size of n, this requirement could severely reduce the amount of usable data.

the timeframe examined happens to include not enough or too many low-probability events and is, unfortunately, exacerbated by smaller sample sizes.

3.1.2 Black-Scholes Implied Distribution

To lessen the reliance on historical data, this paper also proposes an alternative method that derives the subjective probability distribution from option prices through the Black-Scholes option pricing formula. The Black-Scholes formula assumes that the subjective underlying compound returns are normally distributed with standard deviation[8] σ. The idea behind this method is that by plugging the market price of any liquid[9] call option and the other readily observable parameters into this formula, it will be possible to infer the value of σ.

Unfortunately, this method will not produce the mean of stock returns for the subjective probability distribution[10], which must still be found from historical data. Thus, it will still produce an inaccurate estimate if past performance is not perfectly in line with future expectations. However, the degree to which this estimate may be affected is significantly smaller in comparison with the previous method; the general shape of the distribution will remain unaffected. Further, the error introduced may be decreased by using a longer time horizon for the historical sample[11].

---

[8] In option theory, this standard deviation of returns is referred to as the *volatility* of the security. Volatility is typically reported in annualized terms, unless otherwise stated.

[9] Liquidity is important as liquid contracts are less likely to be mispriced.

[10] Since the Black-Scholes formula is based on the no-arbitrage principle, it does not take into account drift in the underlying price.

[11] The intuition here is that over time, the shape (e.g. variance) of the subjective probability distribution may change, but since stock prices are assumed to follow a random walk, the overall

A second potential source of error lies in the possibility that general form of the subjective distribution used in the Black-Scholes formula is incorrect. For instance, Cassidy, Hamp, and Ouyed (2010) show that historical DJIA and S&P 500 returns are both better described with a Student's t distribution than with the normal distribution assumed by Black-Scholes. If this is the case, then the above method will not give an accurate subjective distribution. This problem is discussed in more depth in the following section.

3.2 Finding Risk-Neutral Probabilities

Calculating risk-neutral probabilities is a much more difficult process than is portrayed by Figures 2.2 through 2.7. The following sections detail the complications involved with this technique and discuss possible solutions.

3.2.1 An Empirical Difficulty

Although the procedure outlined in Section 2.3 for using option prices to find the risk-neutral distribution is theoretically sound, it is empirically impossible. For a stock or index, each penny-increment in price, from $0.00 to infinity, represents a potential future state of the world. To find the prices of every individual Arrow security for a given stock or index, the options on that security must be available in every possible strike price, which is obviously not the case. The strikes for the S&P 500 options, for instance, are only available in uneven intervals of integer dollar amounts; consecutive strike prices can differ by anywhere from $5 to $100. As a result of these gaps, the second differences of

---

anticipated mean will remain the same. Thus, if subjective probabilities are solely inferred from historical data, then it is important to use a shorter timeframe in order to better represent the overall shape, but if historical data is only used to approximate the mean of returns, then a longer time period is desired.

option prices mean something entirely different from and more complicated than the price of a single Arrow security. The following simplified example illustrates the idea behind the problem.

Suppose there are three options: c^{10}, c^{15}, and c^{20}, which have strike prices of \$10, \$15, and \$20, respectively. The payoffs of two new derivative contracts $c^{10'}$ and $c^{15'}$ are presented in Figures 3.1 and 3.2 respectively.

Figure 3.1. Payoff diagram for $c^{10'}$.

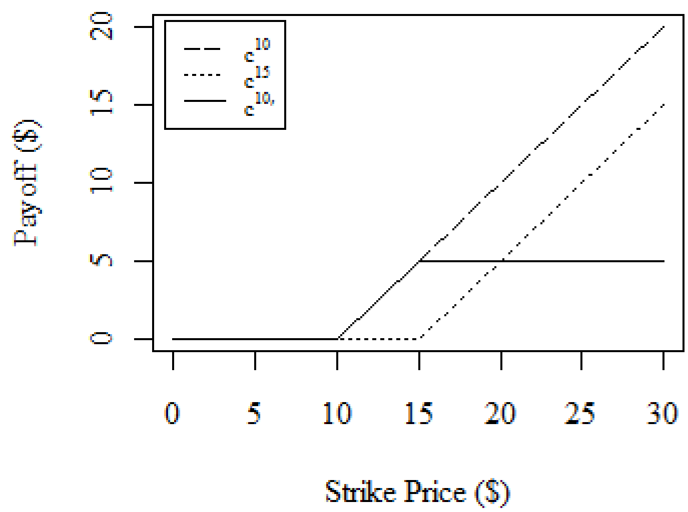

Figure 3.2. Payoff diagram for $c^{15'}$.

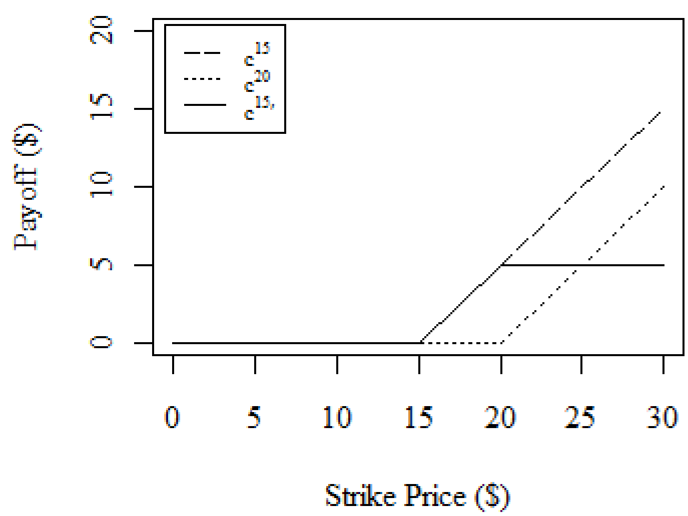

Finally, the payoff of $c^{10''}$ may be found as the difference between $c^{10'}$ and $c^{15'}$:

Figure 3.3. Payoff diagram for $c^{10''}$.

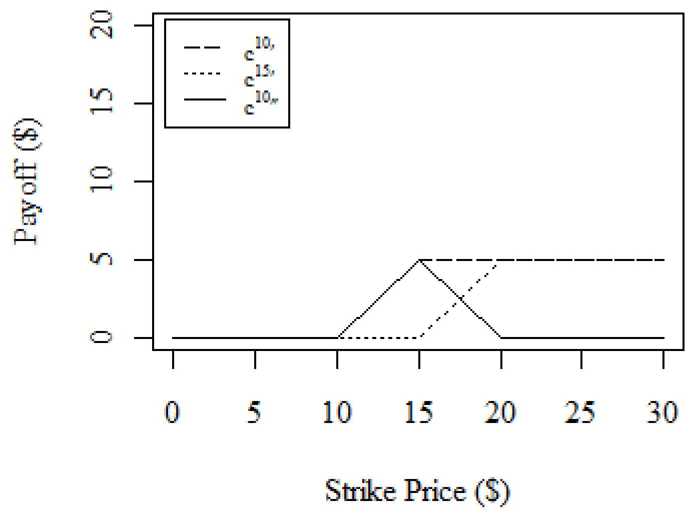

Figure 3.3 shows that the second difference of the option prices of c^{10}, c^{25}, and c^{20} will *not* be equal to the hypothetical price of a single Arrow security that corresponds to the state of the world where the stock price is $10. Instead, the second difference reflects the value of a basket of Arrow securities at mixed proportions, which cannot be easily translated into risk-neutral probabilities as a function of the state of the world (i.e., stock price) at expiration of the option.

Jackwerth (2000) proposes a promising solution to this problem: using the Black-Scholes option pricing formula and the available option prices to create a smooth function, effectively filling in the $5 to $100 gaps between the existing strike prices with synthetically-generated option prices. The intuition here is that the Black-Scholes formula, given the necessary parameters, will give the price of a call option as a continuous function of strike price. Taking the second derivative of the Black-Scholes formula with respect to strike price, then, is analogously equivalent to taking the second differences of the prices of options over different strikes. This method would produce a continuous function of Arrow security prices, which can then translate into a probability density function that represents the risk-neutral probabilities.

This method rests on the assumption that the Black-Scholes formula accurately reflects the call option prices from the risk-free rate r, time to expiration t, strike price K, spot price S, and volatility of the underlying security σ. Since σ is (the only parameter) not directly observable, it must first be estimated from available option data through the Black-Scholes formula. Jackwerth's entire process is roughly as follows:

1. Using the Black-Scholes formula and known values for the call option price, r, t, K, and S, calculate the implied volatility for each liquid option contract on the market. The rationale for using many different option prices instead of just one is that there will always be noise in market prices, and this noise can be reduced with a larger sample size.
2. Average the resulting implied volatilities. The justification for this step is that there should be only one value for the volatility of the underlying security; the observed differences found must thus be due to random noise.
3. Take the second derivative with respect to strike price of the Black-Scholes formula. This, as Breeden and Litzenberger (1978) suggest, is analogous to finding the second differences in the discrete example above.
4. Plug the final implied volatility back into the second derivative found. The resulting function is the risk-neutral probability distribution for the underlying, as implied by its option prices.

In an ideal world where the Black-Scholes formula perfectly describes option prices, the above procedure would produce the correct option-implied risk-neutral probability densities. Further, the method offered in the previous section for finding subjective probabilities from option prices would also give accurate results. This means that both probability distributions may be found solely with option prices and the Black-Scholes formula; one run-through of Jackwerth's proposed method would be sufficient to find risk aversion[12].

Unfortunately, this is not a Black-Scholes-perfect world, and there is one major problem in using the Black-Scholes formula. Empirically, when one finds the Black-Scholes implied volatilities of a series of strike prices for a particular option chain, the values for these volatilities are not, as is expected, the same (with random noise). Instead, the implied volatilities vary with "moneyness"; they tend to be lower for near-the-money strikes and higher the deeper the strike is in- or out-of-the-money. Graphically, these implied volatilities, plotted against their strike prices, create a curved shape often known as the *volatility smile* (Rubinstein (1994)). Figure 3.4 presents the implied volatilities found from the prices of June-07 S&P 500 call options on May 16, 2007.

This figure demonstrates an important characteristic of equity options: an asymmetric, skewed smile. It is extremely well documented (see e.g. Bates (2000) and Foresi and Wu (2005)) that, empirically, equity options typically exhibit only the left portion of the volatility smile. That is, implied volatility almost always has a negative

[12] It is interesting that Jackwerth himself does not use option prices to find the subjective probabilities. Perhaps this is an attempt to avoid the biases introduced by problems with the Black-Scholes formula that this paper will now describe, though Jackwerth does not appear to consider these difficulties.

slope with respect to strike price. This lopsided implied volatility function is sometimes affectionately termed the *volatility smirk*, a subcategory of the volatility smile.

Figure 3.4. SPX call option implied volatilities.

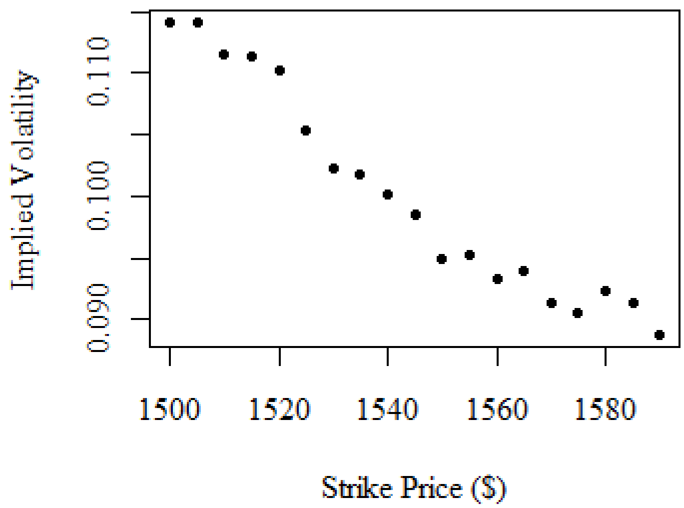

In any case, the volatility smile/smirk phenomenon indicates that, based on the Black-Scholes formula, there is not a single implied volatility to be found. Two related implications immediately follow. First, since theory requires that there is one value for the volatility of the underlying security, the Black-Scholes formula cannot be correct in its current form. Second, if the Black-Scholes formula is not correct, then no matter what value is used for the implied volatility, its second derivative with respect to strike price will not yield the price of Arrow securities.

3.2.2 Resolving the Volatility Smile

In principle, Jackwerth's procedure (and the method this paper proposes for inferring subjective probabilities from option prices) would work if implied volatilities were constant. Therefore, one way of improving on both methods would be to somehow resolve the volatility smile/smirk.

The existence of the volatility smile/smirk indicates that the normal distribution used by default in the Black-Scholes formula may not accurately approximate stock returns[13]. To be more precise, a volatility smile suggests that the normal distribution does not have enough excess *kurtosis*[14]. Further, the fact that the volatility smile is skewed into a smirk, as is the case with equity options, is evidence that the true distribution may be skewed as well.

The Black-Scholes formula is derived from the principle of no-arbitrage, and so its basic form is necessarily valid as long as no-arbitrage holds. All inputs besides volatility and the probability densities of the underlying security are directly observable, so the error must be related to one of these two areas. As the fatness of the tails of the subjective distribution increases, so does the expected value of an option with a strike price near those tails[15]. Further, the value of an option rises with the volatility of the underlying[16]. Since the normal distribution is built into the Black-Scholes formula, any "additional" value of an in- or out-of-the-money that is actually due to the true subjective distribution having positive excess kurtosis is misattributed to a higher volatility. Thus, in

[13] To clarify, stock returns here refer to the subjective distribution, not the objective distribution, which cannot be known.

[14] Kurtosis refers to the fatness of the tails of the distribution as well as the sharpness of the center peak. Higher kurtosis means fatter tails and, often but not necessarily, a sharper center peak for the distribution. A normal distribution has zero excess kurtosis by construct and is termed *mesokurtic*. A distribution with positive excess kurtosis in comparison with the normal distribution is *leptokurtic*, and one with negative excess kurtosis is *platykurtic*.

[15] This is because as the tails become fatter, the probability of an extreme move in the underlying increases. Increasing the probability of a positive extreme move increases the expected value of the option, but increasing the probability of a negative extreme move does not decrease the value of the option, because the holder of the option can choose to not exercise.

[16] Again, since an option confers the right, but not the obligation, to trade the underlying, larger stock moves mean bigger opportunities but no additional risk.

general, the appearance of a volatility smile is evidence that the distribution of security returns should be leptokurtic: relatively thicker at the tails than the normal distribution allows.

By similar logic, the downward-trending volatility smirk in equity options indicates that the left side of the probability density distribution should have relatively fatter tails compared to the right side. This observation indicates that the correct distribution should be asymmetric with a negative skew.

Since the problem with Jackwerth's use of the Black-Scholes formula appears to stem from the fact that the underlying subjective distribution is incorrect, this paper proposes to improve his procedure by substituting in a more suitable distribution in place of the normal distribution. Three possible candidate distributions are discussed below.

3.2.2.1 The Student's T Distribution

The Student's t distribution is a simple leptokurtic distribution that has, in the past, been empirically shown to be a good descriptor of stock returns (see e.g. Blattberg and Gonedes (1974) and Peiró (1994)). Further, Cassidy, Hamp, and Ouyed (2010) have also derived a modified version of the Black-Scholes formula that uses the t distribution in place of the normal distribution.

Unfortunately, unlike the normal distribution, the t distribution is not stable[17], so a t distribution that accurately describes 1-day returns cannot be easily scaled to describe n-

[17] A distribution is stable if and only if the product of two instances of the distribution is another instance of the same distribution. The product of two normal distributions, for example, is also normal.

day returns. In practice, this problem can be avoided by ensuring that the days to expiration for the options is equal to the number of days over which the underlying returns are found.

The normal distribution is easy to use with the Black-Scholes formula due to the fact that its form needs no further specification. In contrast, the t distribution requires a known value for the degrees of freedom *v*. Cassidy (2012) suggests that one could potentially approximate the degrees of freedom to be equal to N-1, where N is the number of days until expiration of the option. In practice, as Cassidy admits, this will likely be an overestimate of the degrees of freedom. This paper proposes an alternative method where the degrees of freedom will be chosen[18] to minimize the variance-to-mean ratio[19] in the resulting implied volatilities for all the different strike prices, following the intuition that the correct probability density distribution will eliminate the volatility smile and produce a single value for implied volatility.

Since there does not exist a closed-form for the modified Black-Scholes formula using the t distribution as the underlying distribution, the second partial derivative with respect to strike price will be found numerically[20].

[18] This will be accomplished numerically (see next footnote), to two significant figures.

[19] The variance-to-mean ratio is defined as $\frac{Var(x)}{\bar{x}}$.

[20] All numeric estimations are conducted with Java algorithms I have written. Particular attention is paid to ensuring a high level of accuracy and avoiding rounding errors.

3.2.2.2 The Non-central Student's T Distribution

While the t distribution does have the desired leptokurticity, it is symmetric, which means that it is unable to accommodate the asymmetries suggested by the existence of the lopsided volatility smirk.

De Jong and Huisman (2000) propose that a skewed t distribution may be a good approximation for underlying returns for European-style options. The authors find that, compared to other popular methods, using their proposed skewed t distribution in conjunction with the Black-Scholes model yields the lowest root mean squared error between the projected and actual option prices. A skewed t distribution can potentially be a more accurate descriptor of security returns, since its asymmetric properties allow it to accommodate the skew component of the volatility smirk. A further advantage of the skewed t distribution is that nests the standard t distribution, and thus it is guaranteed to be capable of providing a fit at least as good as the best fit with the t distribution.

There are several ways in which the standard t distribution may be "skewed" and no general consensus exists on which is the best. While De Jong and Huisman propose one particular formulation, this paper will use the non-central t distribution, which is skewed by a non-centrality parameter μ, where skewness increases with $|\mu|$ and a negative value for μ indicates a left skew. Since the non-central t distribution has mean $\mu\left(\sqrt{\frac{\nu}{\nu-2}}\right)\frac{\Gamma((\nu-1)/2)}{\Gamma(\nu/2)}$, it is manually re-centered by that amount to have zero mean when used in the Black-Scholes formula.

As with the symmetric t distribution, the parameters for the non-central t distribution will be chosen to minimize the variance-to-mean ratio in the implied volatilities produced. Further, with this underlying distribution, the second partial derivative of the Black-Scholes formula will also be calculated numerically.

3.2.2.3 Distribution of Historical Returns

A clear weakness of the above method lies in the fact that though the underlying returns have been empirically shown to more closely resemble a t distribution, skewed or not, than a normal distribution, the t still might not be the correct form. As a consequence, the imposition of this distribution – or, in fact, any general form – on the probability density function of underlying returns may provoke a worse fit and introduce biases into the subsequent analysis.

This problem may be circumvented with kernel density estimation (which, again, is able to produce a probability density function based on frequency data without presupposing a functional form for the distribution). Here, the idea would be to numerically substitute the kernel density estimate of historical returns found in Section 3.1.1 into the Black-Scholes formula in the place of the normal distribution. Again, this method's weakness lies in its reliance upon the accuracy of historical returns as an estimator of subjective views of future returns.

3.2.3 Kernel Density Estimation from Option Prices

Alternatively, one could conceivably circumvent the Black-Scholes formula and its related problems altogether by returning to the original method of taking the second differences of observed call option prices directly. The difficulty with the second differences approach arises from the fact that the strike prices for existing call options on the market are 1) not complete, since they are separated by \$5-\$100 gaps, and 2) unevenly spaced. If these two issues are addressed sufficiently, it may still be possible to use this method.

To understand the first problem, recall from Figure 3.3 that the second difference of option prices represents the price of a basket of Arrow securities. Call this basket X. Conceptually, this basket of Arrow securities is almost equivalent in value to X′, an alternate basket, which can be constructed by flipping and moving the triangular end-sections (A to A′ and B to B′) of the overall triangle shape in the payoff diagram to form a rectangle:

Figure 3.5. Rearranging X to form X′.

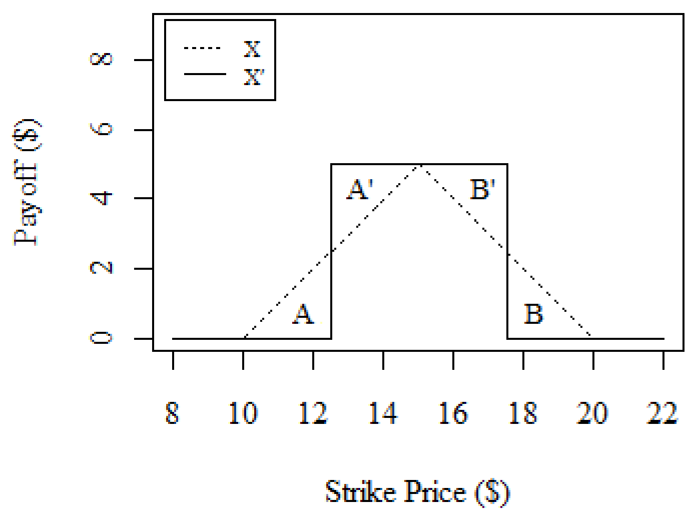

Of course, the values of X and X′ are not exactly the same. The risk-adjustment factor is greater for states of the world in which overall wealth is lower[21]. Therefore, the price of triangle A should have a higher risk premium than that of A′. Conversely, the price of triangle B contains a lower risk premium than price B′. The two effects, since

[21] This follows directly from the idea that the marginal utility of wealth is greater for states of the world where level of wealth is lower.

they act in opposite ways, somewhat offset each other and reduce the amount of bias introduced[22]. Since the difference between the states of the world that correlate to the Arrow securities in A and B are relatively close together in terms of wealth, they should exhibit similar degrees of realized risk adjustment, and so this paper assumes that the bias is small enough that it will not affect the validity of the results.

A single second difference of option prices thus approximately represents the price of a range of Arrow securities of equal proportions. Using the frequency form of this constructed data, a smoother shape for the overall probability density function can be then found through kernel density estimation.

Figure 3.6. Payoff diagram for $c^{10''}$.

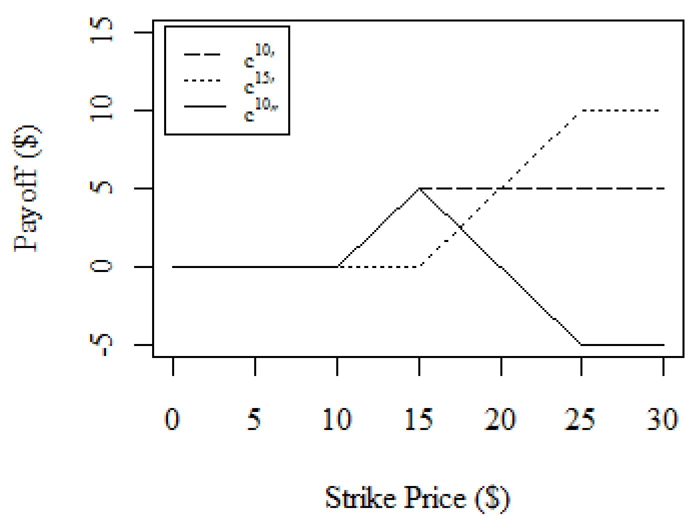

The uneven distances between strike prices are a problem for the second differences of the option prices because they cause the first differences to be “mismatched”. Consider Figure 3.6, which shows the first and second differences of

[22] There is bias, not just noise, introduced. Under the assumption that marginal utility function, which is proportional to the risk-adjustment factor, is positive-valued, strictly decreasing with respect to wealth, and concave-up, the absolute value of the difference between A and A′ will always be larger than that of B and B′, since A < B.

option prices in a case where three consecutive strike prices are $10, $20, and $25. Clearly, if the differences between strikes are inconsistent, the second differences will not represent the value of the correct bundle of Arrow securities.

This can be easily fixed by normalizing the first differences accordingly before finding the kernel density estimate. To be precise, if the range of strike prices represented by one first difference is *n* times that of another, then the first can be adjusted by dividing by *n*. Since the end result will be normalized such that the area under the curve sums to one, the absolute values do not matter as long as they are properly scaled relative to each other. Of course, this will also introduce some error into the data, as the risk-adjustment does not scale linearly[23]. However, bias is not necessarily added, especially when the entire probability density function is smoothed via kernel density estimation, since any particular first difference will serve as the minuend and then the subtrahend in the calculation of neighboring second differences[24].

3.3 Summary of Methods Used

First, the subjective probabilities will be estimated through:

1) the kernel density estimate of historical returns
2) the Black-Scholes formula, using an optimized non-central t distribution[25]

[23] This is to say that the marginal risk-adjustment factor, which is the second derivative of the utility function, is not constant.

[24] Errors will affect consecutive second differences in opposite directions. Globally, these differences should be "averaged out" and should not compromise the overall distribution shape.

[25] Note that the standard t distribution is not used in this paper because it is nested by the non-central t.

Next, the risk-neutral probabilities will be inferred via:

1) the Black-Scholes formula, using the same optimized non-central t distribution
2) the Black-Scholes formula, using the kernel density estimate of historical returns
3) the kernel density estimate of second differences of actual option prices

Once these probabilities are found, the risk-neutral probabilities will be divided by the subjective ones to obtain risk aversion (in the form of the risk-adjustment factor).

4 Data

The data used in this paper is the S&P 500 Index and its corresponding European options (the SPX). The S&P 500 is chosen because, as a capital-weighted index of 500 large-cap stocks traded on the NYSE or NASDAQ, it offers a good representation of overall market conditions and is well-protected against turbulence due to company-specific factors.

Since this is essentially a cross-sectional study, two dates are selected for analysis, one before and one after the financial crisis. May 16, 2007 is chosen as the pre-crash date, and May 20, 2009 as the post-crash. The pre-crash date is selected to slightly precede the beginning of the crash, marked by the collapse of two Bear Stearns hedge funds in July of 2007, when the VIX, the S&P volatility index, showed a substantial increase. Likewise, the post-crash date is set to coincide with the return of VIX to volatile pre-crash levels and before the Greek debt crisis in April of 2010. Both dates are roughly in the same time of year so as to control for possible seasonal fluctuations in risk aversion. The specific dates, the 16th and the 20th, are selected because they are exactly one month before the expiration of June options[26] for their respective years.

To infer the option-implied subjective and risk-neutral probabilities through the Black-Scholes formula, only the strike prices with the more liquid[27] contracts are used to reduce noise from pricing errors. These contracts are chosen as the longest possible series of consecutive strikes around the current spot price where the volume moved on options of each strike is at least 100[28] on the date in question. To further reduce the effect of possible mispricings, the option price used is taken as the average between the bid and ask prices at market close.

To estimate risk-neutral probabilities directly from the second derivative of option prices, on the other hand, requires having as complete a set of strike prices as possible in order to capture the full range of returns. Thus, call options with all publicly traded strike prices for the SPX on the examined dates are used. Without the liquidity criterion described above, it is possible that some option prices are noisy due to mispricing. These pricing errors sometimes lead to apparent arbitrage opportunities and, sometimes, a negative second difference. To combat this problem, this paper follows Jackwerth's example of removing certain strikes to eliminate arbitrage opportunities while

[26] June, as it is both a monthly and a quarterly expiration date, is one of the more popular contracts to trade and thus would offer relatively more liquidity.

[27] Since there is no single clear measure of liquidity, volume traded is used as proxy.

[28] A single option contract entails rights to 100 shares of the underlying, so a volume of 100 is equivalent to 10,000 "shares" of the S&P 500 Index. While this cutoff may seem arbitrary, it does seem to give a sufficiently large range of usable strike prices and eliminate the contracts that show obvious signs of inconsistent pricing.

maximizing the number of strike prices retained. Finally, the mid-market price is still used as the option price.

For the kernel density estimate of historical returns, the S&P 500 Index returns are taken – as they were in Jackwerth's study – over a period of four years before each of the two aforementioned dates. While this seems to be a rather arbitrary length of time, Jackwerth points out that changing the time horizon does not drastically alter the results obtained. Monthly returns are taken in order to match the time to expiration in the options studied. Four years of historical data gives 48 non-overlapping monthly returns, which is certainly not an abundance of data points and may cause some irregularities in the results.

The mean of historical returns used to center the option-implied subjective distribution is taken from the geometric mean of large company stock returns from 1926 to 2010 (Ibbotson (2012)). During this period, the geometric mean of annualized returns is 9.9%, which implies 1-month returns of 2.86%[29].

The risk-free rate is assumed to equal the three-month treasury rate on each date, which is 4.75% pre-crash and 0.18% post-crash.

5 Subjective Probabilities

5.1 Historical Realized Returns

Figure 5.1 presents the probability density function derived from monthly historical returns for the S&P 500 during the four years prior the pre- and post-crash dates. The distributions shown here are largely consistent with expectations.

[29] This calculation assumes that the returns follow a stable distribution. However, even if this assumption does not hold, the error introduced should not be significant.

Figure 5.1. Subjective probabilities derived from historical S&P500 returns.

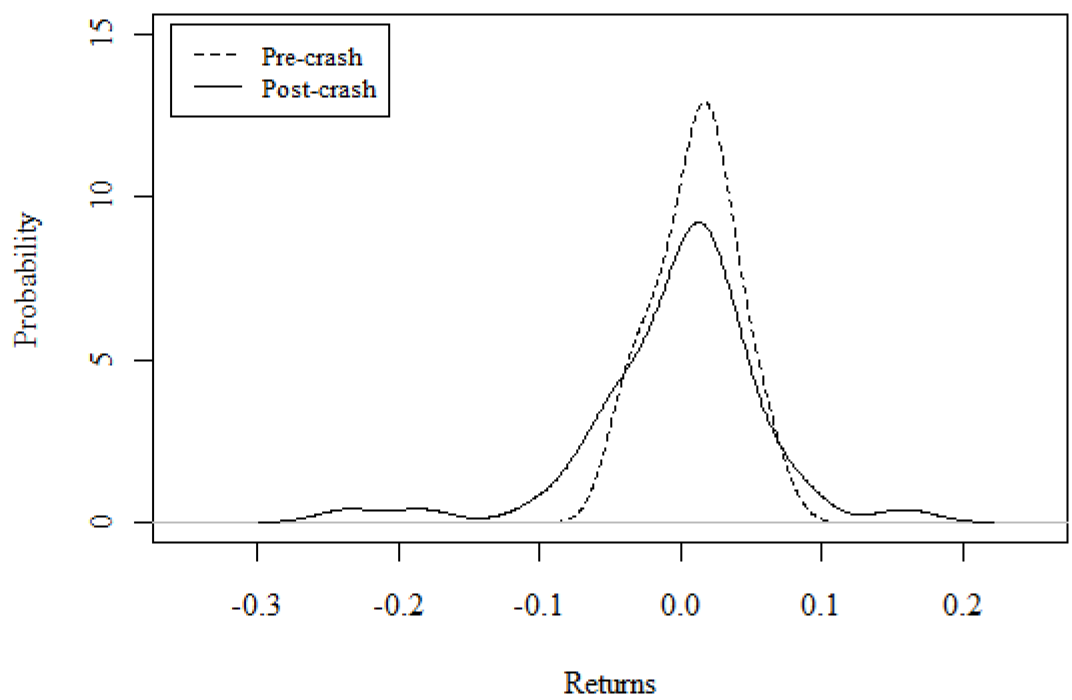

Pre-crash, returns appear to be almost normally-distributed, with relatively low variance (σ = 0.029) and kurtosis (2.415[30]). By comparison, the distribution of post-crash returns exhibits higher variance (σ = 0.063) and kurtosis (6.932). The increased variance is evident both in the way the distribution covers a wider range of values for returns and in the fact that the peak of the distribution reaches a much lower density. The greater kurtosis can be seen in the thickened – though uneven[31] – tails of the post-crash distribution. These differences between the pre- and post-crash distributions indicate that S&P 500 returns were more volatile during the four years prior to the post-crash date than

[30] Since the normal distribution has a kurtosis of 3, this distribution, with its kurtosis of 2.415, is actually platykurtic. This counterintuitive result may be a consequence of the small sample size used.

[31] This jaggedness in the kernel density estimate is due to the unfortunate, yet unavoidable, small sample size.

during the time frame before the pre-crash date. This is entirely unsurprising, as a financial crisis is characterized by large movements in market prices.

Further, there are different degrees of negative skew apparent in both subjective distributions (-0.149 and -1.196 for pre- and post-crash, respectively), which is consistent with general empirical findings[32]. More importantly, this result is consistent with the existence of the skewed volatility smirk when using the normal distribution, which has zero skew, as the underlying subjective distribution in the Black-Scholes formula (recall from Section 3.2.2 that the left-sided smirk indicates that the proper distribution should have a heavier tail on the left side).

There is some apparent multimodality in the post-crash distribution, which is unexpected as returns are generally thought to be unimodal. This anomaly most likely results from the fact that with a sample size of only 48, it is quite possible for the sample to be not perfectly representative of the population. The kernel density estimate, of course, becomes unimodal when a sufficiently large bandwidth is used in the estimation.

5.2 Option-Implied Distribution

5.2.1 Optimal Parameters

The shape of the non-central t distribution is controlled by two parameters: degrees of freedom v and non-centrality μ. As mentioned in Section 3.2, there is no way to directly calculate appropriate values for v and μ, so a pair of "optimal" v and μ is chosen to minimize the variance-to-mean ratio of the implied volatilities that result from

[32] Though there exists a large volume of empirical literature directed at attempting to explain this phenomenon (see e.g. Pindyck (1984) and Hong and Stein (2003)), no consensus appears to have been reached.

using the non-central t distribution defined by those parameters as the subjective probability distribution in the modified Black-Scholes formula.

The following table shows the variance-to-mean ratios (VTM) of the resulting implied volatilities when the specified distribution is used in the Black-Scholes formula. Three distributions are listed below for both pre-crash and post-crash: 1) the optimal non-central t distribution for that time period, 2) the optimal non-central t distribution for the other date, and 3) the normal distribution. The latter two distributions are included to provide contrast and a basis for evaluating the optimal distribution.

Table 5.1. Variance-to-mean ratios of potential subjective distributions.

	Distribution	***ν***	***μ***	**VTM**
Pre	non-central t (optimal)	6.00	-2.05	1.724×10^{-5}
	non-central t (post-crash)	5.08	-1.17	2.404×10^{-4}
	normal	∞	0	7.502×10^{-4}
Post	non-central t (optimal)	5.08	-1.17	1.262×10^{-4}
	non-central t (pre-crash)	6.00	-2.05	2.573×10^{-3}
	normal	∞	0	1.093×10^{-3}

There are several important observations to be made here. First, the optimized non-central t distribution always gives implied volatilities with variance-to-mean ratios substantially lower than those from the normal distribution. This is a good indication that the non-central t does, in fact, provide a better fit for underlying returns. Next, the variance-to-mean ratios in the post-crash period are always higher than their pre-crash counterparts. This phenomenon likely results from the post-crash option prices containing more noise (variation)[33]. Finally, a non-optimal non-central t distribution could provide a considerably worse fit. During the post-crash period, the optimal pre-crash non-central t managed to generate implied volatilities with a variance-to-mean ratio more than double that of implied volatilities found by using a normal distribution.

5.2.2 Volatility Smile

Figure 5.2. Implied volatilities from normal and non-central t distributions.

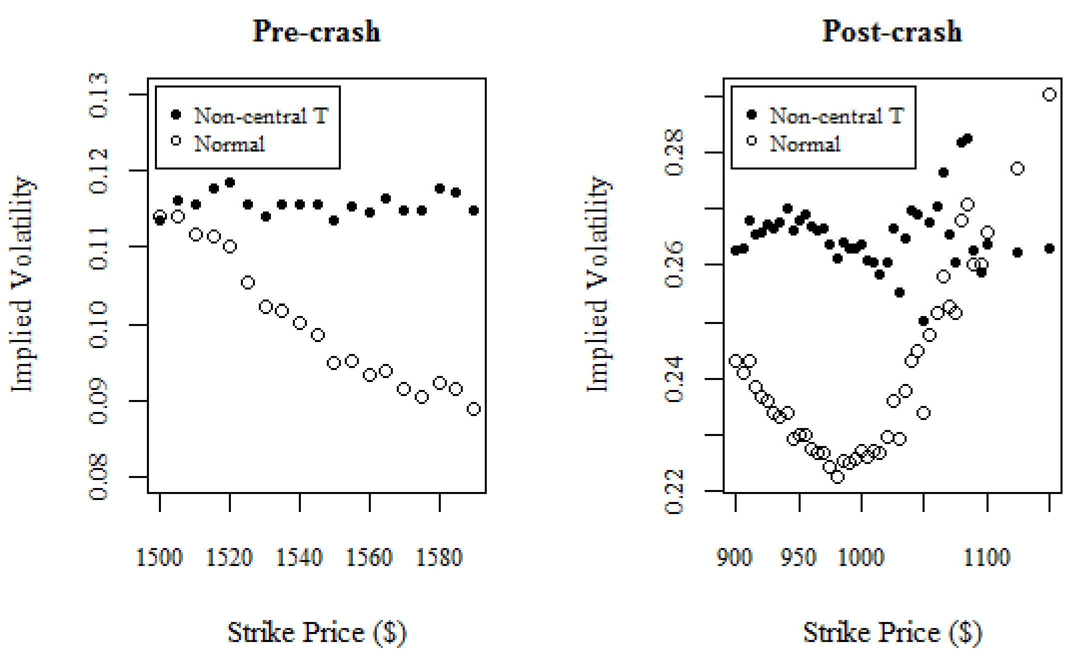

Figure 5.2 compares the implied volatilities found through using the optimized non-central t distribution versus the normal distribution in the Black-Scholes formula. As

[33] A higher variance-to-mean ratio could also signify that the same distributions from the pre-crash period are no longer good approximations of subjective probabilities. However, as Figure 5.2 on the next page shows, the optimal non-central t distribution does remove the volatility smile, which shows that it is an appropriate distribution to use.

is clearly evident, using the non-central t distribution in the Black-Scholes formula greatly reduces, if not completely eliminates, the volatility smirk that results from using the normal distribution. This fact is a good indication that the subjective and risk-neutral probabilities derived from these non-central t distributions will be much more accurate than those found using the normal distribution.

Further, the implied volatilities on the pre-crash date seem to be much less "scattered" in comparison with their post-crash counterparts. This observation is also consistent with the fact that the variance-to-mean ratio for the pre-crash implied volatilities is significantly lower than that of the post-crash ones. The fact that there seems to be much more noise in option pricing after the crash could be a symptom of general crash-induced confusion – perhaps market participants had very different subjective probabilities in mind, and those probabilities were not aggregated coherently on the market.

Finally, it is slightly disconcerting to see that the normal distribution-induced volatility smile curves upward to such an extent. This defies general empirical observations of equity volatility smiles. However, it is quite possible that the liquidity criteria for the data has managed to cut off the left side of the smirk – the underlying price on the post-crash date is $903.47, which is at the left edge of the graph, and options with strikes lower than $900 have not been used due to surprisingly low volume[34].

[34] Call options with a strike of 895 – one level down from 900 – traded a total of 0 contracts that day, compared to 11679 contracts traded of the 900 call. This is actually not too unexpected; it is entirely conceivable that during the recovery period after a crash, market participants anticipate very high returns and so strike prices of less than the spot price become largely irrelevant.

Alternatively, it may be that after the crash, market participants expected a strong recovery, and thus their subjective probability distribution has a large positive skew.

5.2.3 Distribution

Using the parameters obtained for v and μ through the aforementioned optimization process in conjunction with the average value of the implied volatilities from Section 5.2.2 (0.116 and 0.265 for pre- and post-crash, respectively), the following subjective distributions are derived:

Figure 5.3. Subjective distributions inferred from option prices.

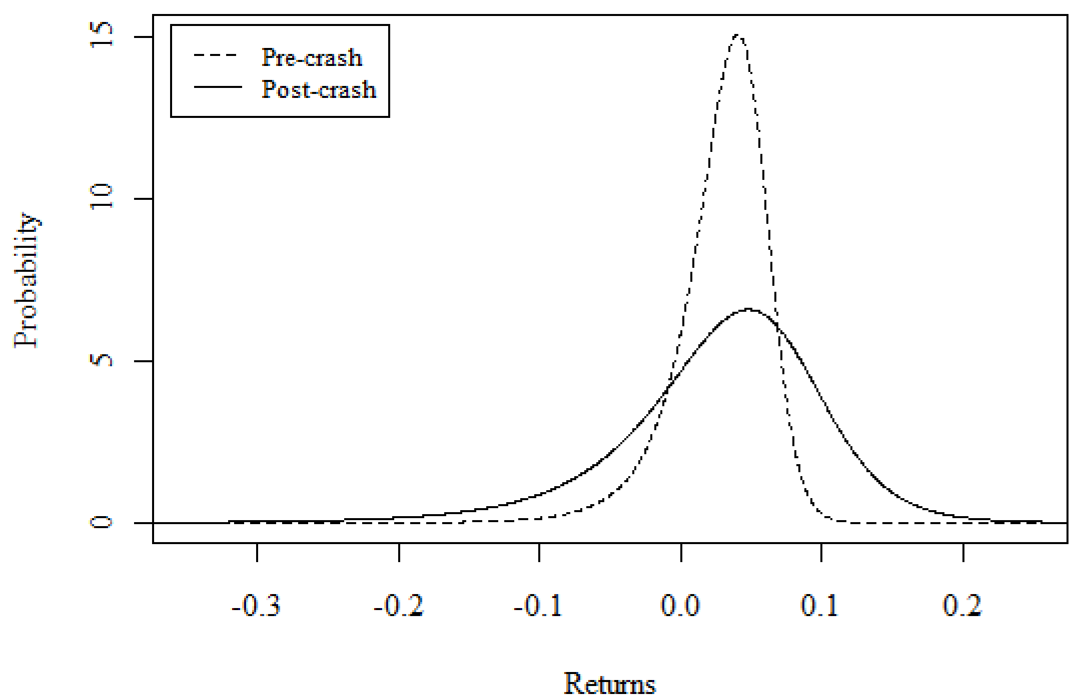

The standard deviation, skew, and kurtosis of these two distributions and the two distributions found from historical returns are summarized together in the following table for comparison.

Table 5.2. Moments of historical and option-implied subjective distributions.

	Method	st. dev.	skew	kurtosis
Pre	historical	0.029	-0.149	2.415
	option-implied	0.033	-1.333	7.503
Post	historical	0.063	-1.196	6.932
	option-implied	0.077	-1.179	7.828

It is interesting to see that the option-implied and historical subjective distributions are much more similar in the post-crash than the pre-crash. Perhaps this difference is due to small-sample bias affecting the kernel density estimate of the pre-crash historical returns more than the post-crash ones. This explanation would be consistent with the observation that the pre-crash historical distribution is unexpectedly platykurtic. The subjective distributions derived from option prices and from four-year historical returns are illustrated together in Figure 5.4.

As Figure 5.4 illustrates, while the shapes of the two subjective distributions are quite similar for both pre- and post-crash data, there is also a definite degree of mismatch between them, consistent with the fact that the distributions have different variance, skew, and kurtosis[35]. Most notably, the peak of the option-implied distribution is higher than that of the historical returns pre-crash but is lower post-crash. These differences, while substantial, are not necessarily cause for alarm. One possible reason for the discrepancy is that it is impossible to fit a non-central t distribution (or any formal distribution, for that matter) perfectly to the idiosyncrasies in the historical data. For instance, the bumps

[35] Undoubtedly, they are also centered on different means, since a 4-year historical mean is certainly not the same as an 84-year mean. The difference is especially pronounced for the post-crash distributions, since much of the 4 years sampled included the financial market collapse.

in the tails of kernel density estimates of the post-crash historical returns make the overall distribution distinctly multimodal, which is a characteristic that the unimodal non-central t cannot accommodate. The differences in the two sets of distributions may also be a symptom of historical performance being an imperfect predictor of future expectations, especially in regards to the mean of returns. In this case, the option-implied subjective probabilities are likely to be more accurate than the historical ones.

Figure 5.4. Subjective probabilities from historical returns vs. option prices.

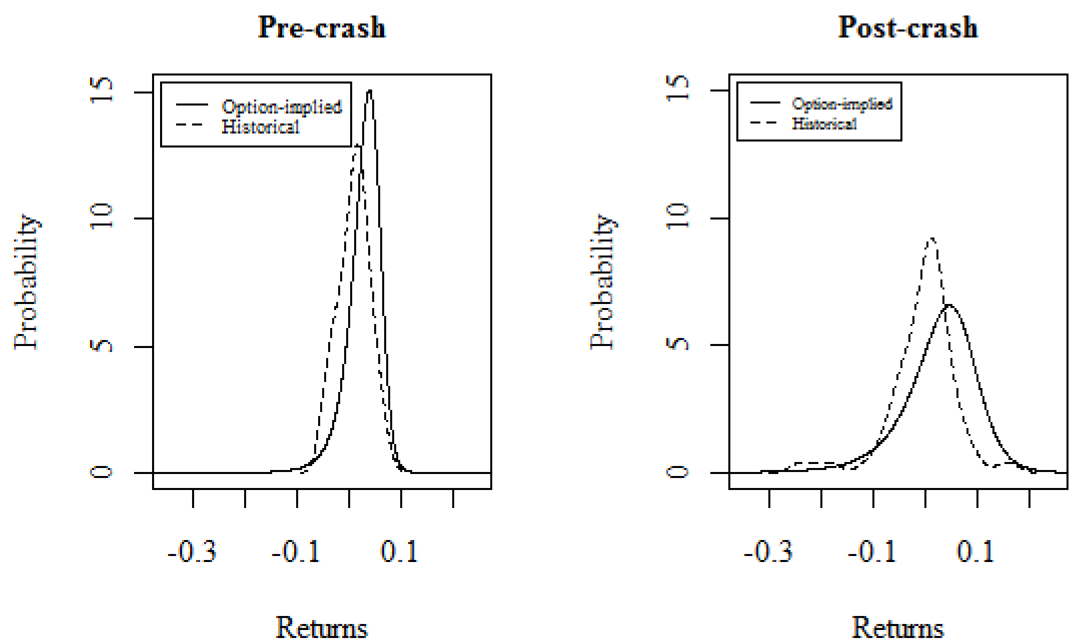

6 Risk-Neutral Probabilities

6.1 Non-Central T Distribution

By using the non-central t distribution with parameters ν and μ as the subjective distribution of underlying returns in the Black-Scholes formula (as discussed in the previous section), it is now possible to infer the risk-neutral probabilities from option prices. These risk-neutral probability densities for both the pre- and post-crash dates are shown below.

Figure 6.1. Risk-neutral distributions inferred from option prices (non-central t).

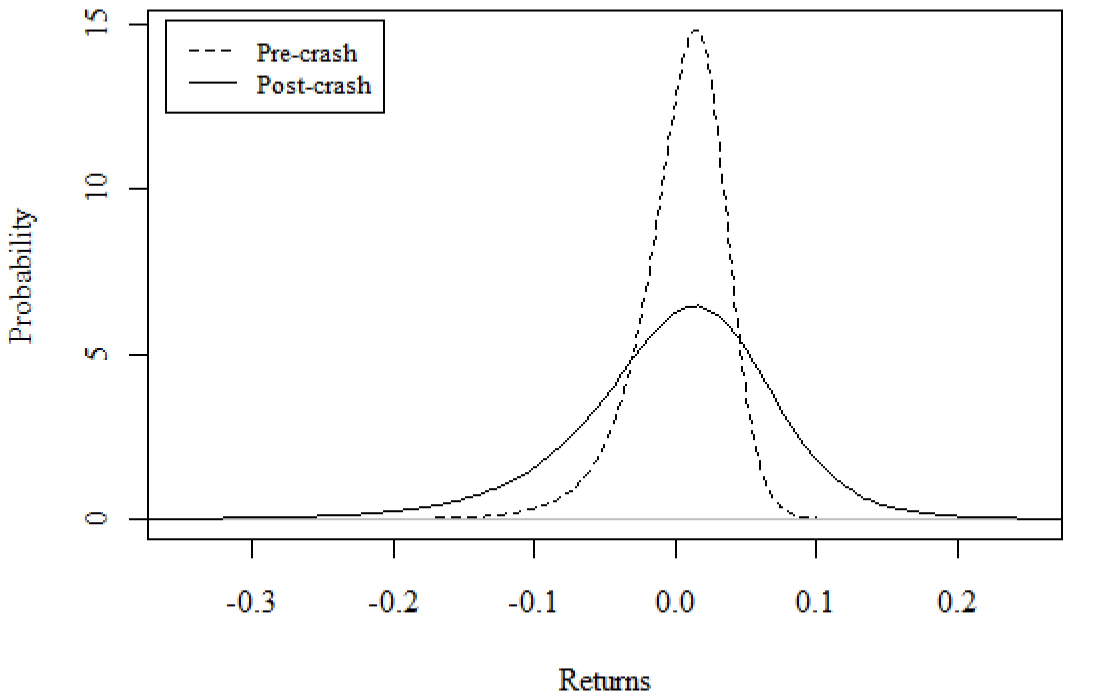

Table 6.1. Moments of subjective and risk-neutral distributions.

	Distribution	st. dev.	skew	kurtosis
Pre	risk-neutral	0.033	-1.169	6.857
	subjective (option-implied)	0.033	-1.333	7.503
	subjective (historical)	0.029	-0.149	2.415
Post	risk-neutral	0.074	-0.651	5.924
	subjective (option-implied)	0.077	-1.179	7.828
	subjective (historical)	0.063	-1.196	6.932

As Table 6.1 shows, the standard deviation, skew, and kurtosis of the option-implied risk-neutral distributions are much more similar to the option-implied subjective distributions than the historical-derived ones. This could be seen as evidence that the historical returns are not good estimators of the subjective distribution. However, it is important to remember that nothing requires that the risk-neutral and subjective distributions be similar.

6.2 Kernel Density Estimate of Realized Returns

A second method for inferring risk-neutral probabilities still uses the Black-Scholes formula but replaces the non-central t distribution from Section 6.1 with the kernel density estimate of the historical returns found in Section 5. In order to use the kernel density estimate in the formula, it is first normalized to have a standard deviation of one (while ensuring that the area under the curve is still one). This step ensures that the σ in the Black-Scholes formula still represents the volatility of the underlying[36].

The implied volatilities found through this method have substantially higher variance-to-mean ratios (3.155×10^{-4} for pre-crash and 1.663×10^{-3} for post-crash) than those found using the optimal non-central t distribution. In fact, in both pre- and post-crash, these variance-to-mean ratios fall in between those generated by using the normal distribution and the non-optimal non-central t distribution (see Table 5.1). The fact that these variance-to-mean ratios are so comparatively high seems to indicate that historical returns are not accurate estimators of subjective probabilities.

[36] σ is actually the scaling factor between the standard deviation of the underlying and that of the subjective probability distribution used in the formula. Thus, σ is only equal to the volatility of the underlying when the latter is 1.

This hypothesis is corroborated by Figure 6.2, which plots the implied volatilities against their strike prices and reveals a volatility smile. The existence of the volatility smiles in both pre- and post-crash data is further evidence that the historical distributions are extremely poor approximations for the subjective distribution. As such, this method is not used to infer risk-neutral probabilities, since any results obtained via this method are likely to be incorrect[37].

Figure 6.2. Implied volatilities, using historical returns as subjective distribution.

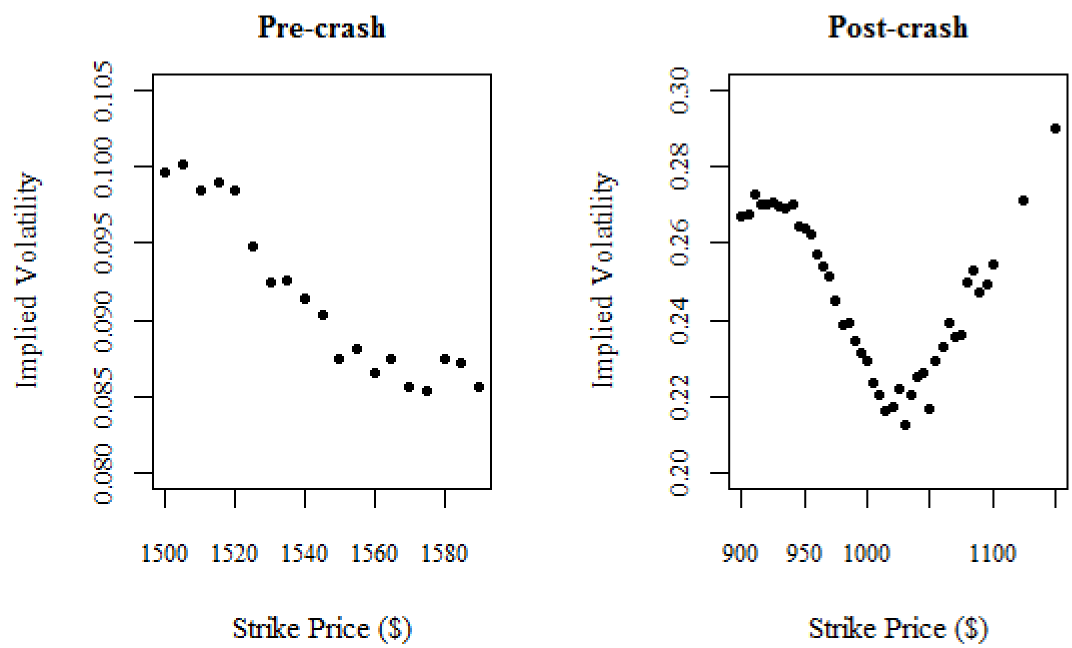

6.3 Direct Kernel Density Estimation with Option Prices

The third and final method of inferring risk-neutral probabilities involves finding the kernel density estimate of the second differences of option prices directly, after removing the contracts that provide arbitrage opportunities. The results are shown below.

[37] At the very least, it will be difficult to justify the selection of any one value as the implied volatility, since there is such clear evidence that implied volatility varies with strike price.

While the inferred post-crash distribution is smooth and unimodal as anticipated, the pre-crash risk-neutral distribution shows artifacts of pricing noise. The extremely bumpy left tail and the additional peak in the middle of the distribution are all results of inconsistent pricing resulting in great variance between consecutive second derivatives. In its current state, the kernel density estimate of second differences is clearly not a good approximation of the actual risk-neutral distribution. Of course, increasing the smoothing parameter will reduce the unevenness of the kernel density estimate, but it will also distort the overall shape and compromise the accuracy of the result.

Figure 6.3. Risk-neutral distributions from second differences of option prices.

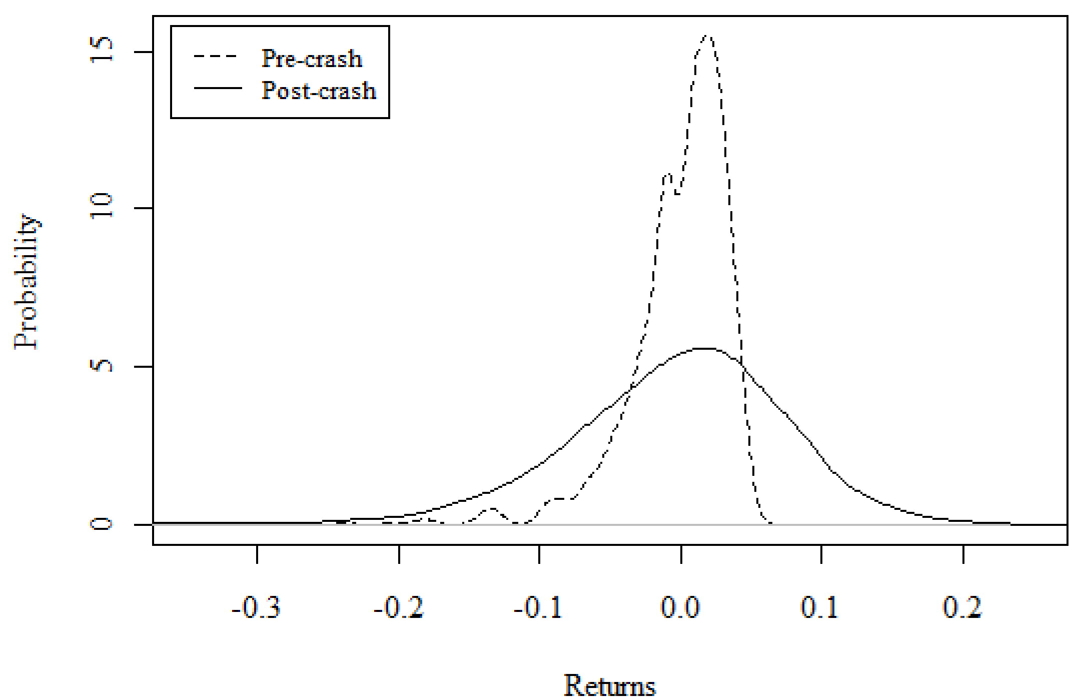

The bumpiness of the pre-crash distribution does go towards justifying methods where option prices are "smoothed" via the Black-Scholes model. Those methods will at least give a smooth, unimodal curve that can be used intelligibly to find risk aversion.

Further, if the correct distribution is used to describe underlying returns (which can be checked by looking for a volatility smile), then the results should not be biased either.

Incidentally, it is interesting that the post-crash risk-neutral distribution here does not suffer the same problem. This is surprising, since, as Figures 5.2 and 6.2 demonstrate, the implied volatilities found for the post-crash option data were much noisier than those for the pre-crash regardless of the underlying subjective distribution used.

7 Risk Aversion / Results

Finally, now that both subjective and risk-neutral probabilities have been estimated, it is possible to infer risk aversion. Since there are many indications to support the hypothesis that historical returns are, in fact, *not* a good indicator of future expectations, the subjective probabilities inferred through the Black-Scholes formula with the optimized non-central t distribution are used. The risk-neutral probabilities used are estimated through the same modified Black-Scholes formula. This choice is due to the fact that the pre-crash risk-neutral probabilities found directly from the second differences of option prices contained unrealistic bumps that could not be eliminated without introducing additional errors, and using the Black-Scholes formula with the historical distribution created such obvious trends in implied volatility that risk-neutral probabilities were not even calculated. Figures 7.1 and 7.2 show the non-central t option-implied subjective and risk-neutral probabilities for the pre-crash and post-crash data, respectively.

Figure 7.1. Pre-crash option-implied subjective and risk-neutral distributions.

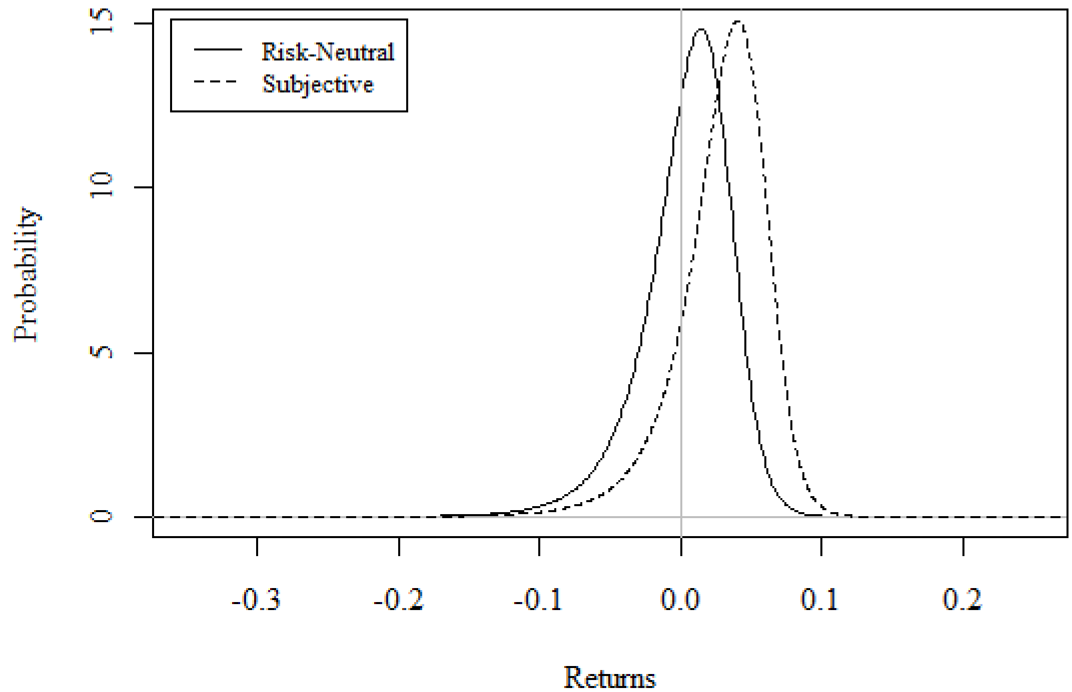

There are several important areas of note in Figure 7.1. First, the shapes of the subjective and risk-neutral probabilities are almost identical, just offset horizontally. This is not surprising, as the risk-neutral probabilities here are derived from option prices synthetically generated by assuming the very same subjective probabilities. Second, the risk-neutral distribution is higher than the subjective distribution for low returns and lower than it for high returns. This observation coincides with the fact that marginal utility – here as the ratio of risk-neutral to subjective probability – should be higher for lower levels of wealth, and vice versa. Finally, the risk-neutral probability associated with excess returns of 0 is higher than the related subjective probability. This makes sense, since individuals are thought to be willing to pay a risk premium to avoid a situation where there is risk (variance in returns) but no expected gain (no excess returns). All things considered, the situation portrayed by Figure 7.1 is consistent with theory.

Figure 7.2. Post-crash option-implied subjective and risk-neutral distributions.

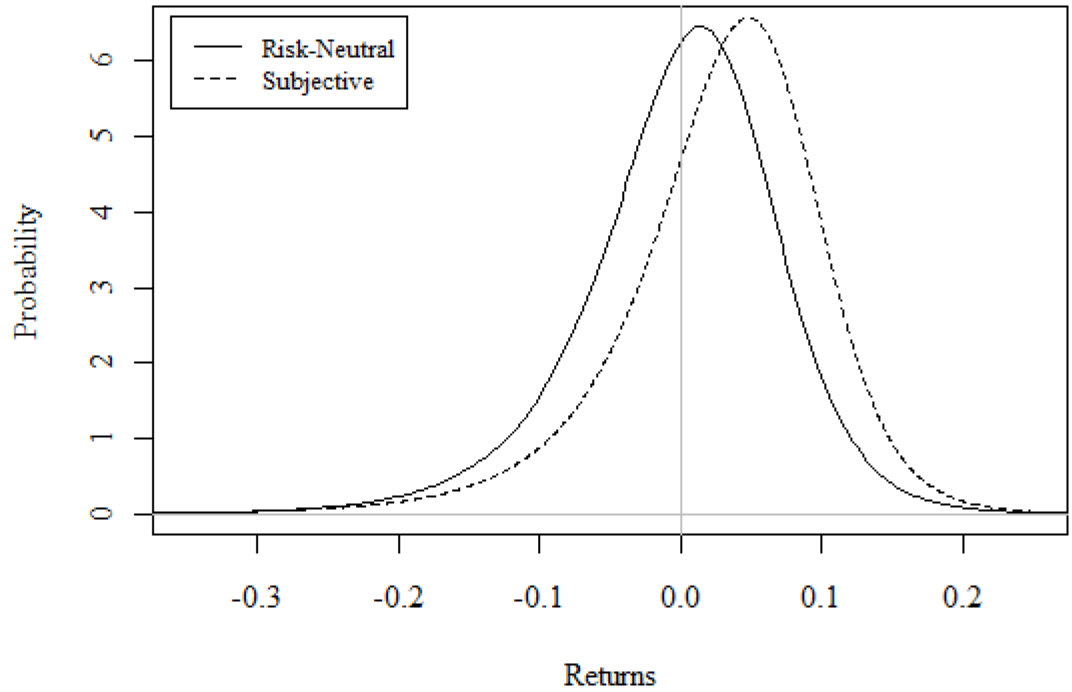

Figure 7.2 illustrates a situation for the post-crash data similar to that shown in Figure 7.1 for the pre-crash. This is particularly interesting, as Jackwerth finds that, while risk aversion seemed normal before the crash of 1987 (the focus of his study), market participants actually became risk-seeking following the crash. The main difference between the method used in this paper and that used by Jackwerth is that his subjective probabilities are inferred from historical returns, not from option prices (as they are here). Clearly, since post-crash historical returns incorporate a large portion of the downward momentum realized during the crash, Jackwerth's subjective distribution will likely be biased leftward. If the left-offset is large enough, it is conceivable that the subjective distribution could be placed to the left of the risk-neutral one, thus giving apparent evidence of risk-seeking tendencies in the market. Thus, it is quite possible that

Jackwerth's alarming findings regarding post-crash risk aversion are actually results of an oversight in methodology.

Dividing the risk-neutral probabilities by the subjective probabilities produces the inferred risk-adjustment factors, which, again, are directly correlated with marginal utility.

Figure 7.3. Risk-adjustment factor as a function of returns.

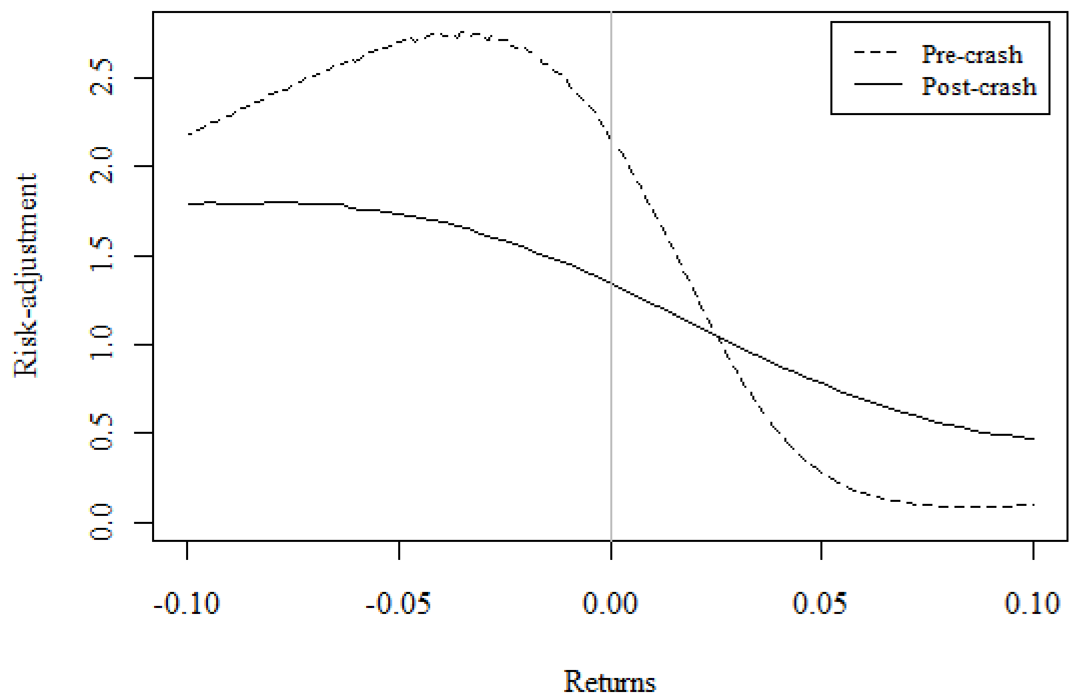

First, Figure 7.3 shows that, contrary to theory, marginal utility is not strictly decreasing with respect to wealth[38]. Second, in the range of x-values where the functions are strictly decreasing, the pre-crash exhibits steeper slope than does the post-crash. This indicates decreased risk aversion following the 2008 financial crisis, since a smaller risk premium is paid to avoid an equal loss of wealth. This apparent result is unexpected: it is

[38] If Figure 7.3 is cropped in the x-axis as Jackwerth does in his paper, then marginal utility will appear to be strictly decreasing.

contrary to both the general theoretical assumption that the utility function is constant over time and the hypothesis that risk aversion increases following a crash.

8 Concluding Remarks

Section 7 briefly hints at the tentative conclusion that, of the different methods outlined in the beginning of this paper, the subjective and risk-neutral probabilities are perhaps best[39] estimated through the use of the Black-Scholes formula modified to assume a non-central t distribution as the underlying subjective distribution.

The obvious advantage of using an option-implied method for finding the subjective distribution is that it reduces the reliance on the assumption that actualized returns are a good approximation of expectations of future returns. As the results above suggest, this advantage may be substantial, since the historical and option-implied subjective distributions differ to a great extent. The fact that the optimized non-central t distribution does not seem to be producing an obvious volatility smile increases confidence that it is, in fact, an appropriate distribution to use. Regarding the risk-neutral probabilities, using the non-central t distribution with the Black-Scholes formula is clearly superior to using the kernel density estimate of historical returns, since the latter produces a visible volatility smile. Additionally, the kernel density estimation of the second differences of option prices is not a good alternative, as it suffers from too much pricing noise. Thus, it is best to use the Black-Scholes formula with the non-central t distribution.

[39] "Best" here is used with consideration to both the accuracy of the estimate and the usability of result in the specific context of inferring risk-aversion.

Using these distributions, this paper then finds that risk aversion decreased between the pre- and post-crash dates. However, this result is subject to a significant objection: it is quite possible that, since 1) historical returns are far from guaranteed to produce an accurate estimate of future outlook and 2) the subjective distribution is centered on the historical mean, *the subjective distribution is not correctly placed*. If, for instance, market participants expect that the market will exhibit higher-than-average returns as part of the recovery process following the financial crisis, then the post-crash subjective distribution would actually be shifted farther to the right of its current location. If the magnitude of this move is sufficiently large, then it could result in greater risk aversion than is found in the pre-crash data. Unless there is a more accurate way to estimate the mean of the subjective distribution, any results derived from it are tenuous at best.

Though a definitive conclusion regarding changes in risk aversion as a result of the financial crisis is not reached, this paper does make several important contributions. First and most significantly, it has found better methods for inferring subjective and risk-neutral probabilities from option prices (though there still needs to be a more accurate way to determine the mean of the subjective distribution, the option-implied method is still an improvement over using historical returns). As part of this process, the paper finds that using a non-central t distribution instead of the normal distribution in the Black-Scholes formula can largely reduce the volatility smile problem. Finally, this paper shows that the apparent evidence of risk-seeking behavior that Jackwerth finds may be due to an error in his method.

References

Bates, David S. “Post-'87 Crash Fears in the S&P 500 Futures Option Market.” *Journal of Econometrics* 94, no. 1 (Jan., 2000): 181-238. http://www.sciencedirect.com/science/article/pii/S0304407699000214 (accessed April 17, 2012).

Blattberg, Robert C., and Nicholas J. Gonedes. “A Comparison of the Stable and Student Distributions as Statistical Models For Stock Prices.” *The Journal of Business* 47, no. 2 (Apr., 1974): 244-80. http://www.jstor.org/stable/2353383. (accessed April 17, 2012).

Breeden, Douglas T., and Robert H. Litzenberger. “Prices of State-Contingent Claims Implicit in Option Prices.” *The Journal of Business* 51, no. 4 (Oct., 1978): 621-51. http://www.jstor.org/stable/2352653 (accessed April 17, 2012).

Cassidy, Daniel T. E-mail message to author. February 22, 2012.

Cassidy, Daniel T., Michael J. Hamp, and Rashid Ouyed. “Pricing European Options with a Log Student's T-Distribution: A Gosset Formula.” *Physica A* 389 (Sept., 2010): 5736-48. http://epic.mcmaster.ca/foswiki/pub/Main/DanielCassidy/111_ Cassidy_t-distribution.pdf (accessed April 17, 2012).

De Jong, Cyriel and Ronald Huisman. “From Skews to a Skewed-T.” ERIM Report Series Reference No. ERS-2000-12-F&A. (2000). http://papers.ssrn.com/sol3/papers.cfm?abstract_id=370814 (accessed April 17, 2012).

Foresi, Silverio and Liuren Wu. “Crash-o-phobia: A Domestic Fear or Worldwide Concern?” *The Journal of Derivatives* 13, no. 2 (Winter 2005): 8-21. http://www.iijournals.com/doi/abs/10.3905/jod.2005.605352 (accessed April 17, 2012).

Gisiger, Nicolas. “Risk-Netural Probabilities Explained.” Working Paper (Oct., 2010). http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1395390 (accessed April 17, 2012).

Hong, Harrison, and Jeremy C. Stein. “Differences of Opinion, Short-Sales Constraints, and Market Crashes.” *The Review of Financial Studies* 16, no. 2 (Summer 2003): 487-525. http://rfs.oxfordjournals.org/content/16/2/487.full.pdf (accessed April 17, 2012).

Ibbotson, Robert G. “The Equity Risk Premium.” *The Research Foundation of CFA Institute* (Mar., 2012). http://www.cfapubs.org/doi/pdf/10.2470/rf.v2011.n4.8 (accessed April 17, 2012).

Jackwerth, Jens C. “Recovering Risk Aversion from Option Prices and Realized Returns.” *The Review of Financial Studies* 13, no. 2 (Summer 2000): 433-51. http://www.jstor.org/stable/2646032 (accessed April 17, 2012).

Peiró, Amado. “Distribution of Stock Returns: International Evidence.” *Applied Financial Economics* 4, no. 6 (1994): 431-39. http://www.tandfonline.com/doi/pdf/10.1080/758518675 (accessed April 17, 2012).

Pindyck, Robert S. “Risk, Inflation, and the Stock Market.” *The American Economic Review* 74, no. 3 (jun., 1984): 335-51. http://www.jstor.org/stable/1804011 (accessed April 17, 2012).

Rubinstein, Mark. “Implied Binomial Trees.” *Journal of Finance* 49, no. 3 (Jul., 1994): 771-818. http://www.haas.berkeley.edu/groups/finance/WP/rpf232.pdf (accessed April 17, 2012).

Scott, David W. *Multivariate Density Estimation*. New York: Wiley, 1992.

Silverman, Bernard W. *Density Estimation for Statistics and Data Analysis*. London: Chapman and Hall, 1986.