THE TAILS OF OPTION-IMPLIED PROBABILITY DISTRIBUTIONS

By

JORDAN B. ZIMBELMAN

Bachelor of Science in Computer Science
Kansas State University
Manhattan, Kansas
2004

Master of Business Administration
Kansas State University
Manhattan, Kansas
2006

Submitted to the Faculty of the
Graduate College of
OKLAHOMA STATE UNIVERSITY
in partial fulfillment of
the requirements for
the Degree of
Doctor of Philosophy

May, 2020

THE TAILS OF OPTION-IMPLIED
PROBABILITY DISTRIBUTIONS

Dissertation Approved:

DR. BETTY SIMKINS

Thesis Adviser

DR. DAVID CARTER

DR. RAMESH RAO

DR. JOE BYERS

DR. LEE ADKINS

Name: JORDAN B. ZIMBELMAN

Date of Degree: MAY, 2020

Title of Study: THE TAILS OF OPTION-IMPLIED PROBABILITY DISTRIBUTIONS

Major Field: FINANCE

ABSTRACT

This dissertation builds on and contributes to work in the field of financial risk management, specifically option-implied probability distributions. Although a number of studies have examined estimating the middle portion of probability distributions, there has not been a strong focus on the tails of the distribution, which are of particular importance in a risk management setting. As such, this study provides additional insights about these tails, by horse-racing four different tail-fitting methods. This research differs from previous studies by introducing a new, non-parametric, heuristic tail-fitting method that is similar in methodology to the consensus, most-often used method to estimate the middle portion of the probability distribution; and, by identifying which tail fitting method produces the most stable estimate with the least tail-option pricing error. In short, the non-parameterized, heuristic method, similar to the fast and stable method most commonly used to estimate the middle portion of the probability distribution, is also stable, with the least option pricing bias in the tails of the distribution.

TABLE OF CONTENTS

Chapter Page

LIST OF TABLES

Table Page

LIST OF FIGURES

Figure Page

CHAPTER I

INTRODUCTION

Option pricing theory, along with increasingly available data and computational power, presents an opportunity to reveal valuable information about market expectations and risk preferences. Investors have long used derivatives to infer information, e.g., cost of carry rates and Black-Scholes implied volatilities. Beyond just this, with a complete set of options prices, an asset's entire risk-neutral probability distribution (RND) can be revealed. Given the well-documented non-normality of most assets' returns (i.e. the Black-Scholes volatility smile), the higher moments of the return distribution are of interest.

Foremost, knowledge of the RND is desirable because it enables the pricing of any derivative of the underlying asset with the same time to expiration, regardless if it be illiquid or 'non-vanilla.' Beyond this obvious application, there are several lines of research that use option-implied RNDs.

For example, central bankers use RNDs to assess market sentiment over future changes in interest rates, exchange rates and stock prices, and likewise to confirm the market's acceptance of major policy decisions. Of particular interest is the predictive power of RNDs prior to major economic events such as crash episodes, exchange rate/interest rate regime changes, wars, and elections. RNDs are fruitful for this type of analysis because unlike other market-based time series data, which in isolation capture only an expectation (and therefore represent only one expected economic scenario), RNDs capture the uncertainty that is fundamentally inherent in the marketplace -- that there are multiple future scenarios.

Another line of research uses RNDs to measure risk aversion. They separately estimate the risk neutral density from option prices and the objective density from historical returns of the underlying asset, and use these two separately derived functions to infer implied relative risk aversion. Knowledge of economy-wide risk aversion is of particular importance since it makes possible a number of forecasting applications in risk management, which require the RND to first be converted into its objective counterpart (the following section discusses risk-neutral versus objective probabilities).

In this thesis, I develop a particularly stable method for fitting the tails of the RND. The tails have proven difficult to estimate, due mainly to data limitations. Options are only written on a limited range of strike prices and those that are written far away-from-the-money tend to have little liquidity. However the tails are of prime importance when attempting to predict major market events, the most common application of RNDs. My model will be tested (horse-raced) against other methods that have been put forth: first, in a time series analysis of stability; then in a pricing error test using observed option prices.

CHAPTER II

REVIEW OF LITERATURE

On April 26th, 1973 the Chicago Board Options Exchange (CBOE) was spun-off from the Chicago Board of Trade (CBOT), originally to celebrate the Board of Trade's 125th anniversary. The idea for the Exchange is said to have come from the brothers "Eddie" and "Billy" O'Connor, prominent CBOT soybean traders, and Joseph Sullivan, then-assistant to the CBOT president, who would become the first president of the CBOE. The Exchange came about despite antagonism from many of the existing members of the CBOT, who felt options to be too specialized for exchange trading. Nevertheless, for the first two years of the CBOE's existence, the "smoking room" of the CBOT became the first trading floor for *standardized* options.

While this was the beginning of modern-day trading in *standardized* options, over-the-counter (OTC) options have existed since antiquity. In fact, some of the earliest-dated writings on record, cuneiform tablets

from circa 1750 BCE, are option contracts on real estate. Indeed, the Code of Hammurabi, pre-dating Moses' Law by roughly 200 years, gives legal foundation for option-contract rights on real property.

But 1973 wasn't really when organized options markets became a thing either. At that time, the CBOE only listed *call* options, no *puts* until 1977. As well, it wasn't until 1983 that they first listed *index* options, the most popular traded contracts today. To tell the truth, the date when the CBOE really came to become was Black Monday, October 19th, the 1987 crash! That being when the overall stock market lost over 20% during morning hours trading; all CBOE contracts nearly went to zero in the following 24 hours, only to survive by a single, miraculous trade in the CBOT's index futures market[1]; a date that the standard Black-Scholes Option Pricing Model completely broke down, and hasn't recovered since. *The problem*: mis-valued risk in the tails of assets' return distributions[2].

Option Pricing.

1973 was also the year when Black and Scholes (1973) and Merton (1973) standardized option pricing. Prior to, the valuing of options was

[1] See Vitale (2012).
[2] This could also be viewed as a *qualitatively* different *kind* of risk (i.e. an "on/off" risk, liquidity-risk), separate from the asset's return distribution.

somewhat 'chaotic' in-itself. Not just because the pricing of contingent claims was beyond the 'mathematic dossier' of most practitioners, but because the models depended on a correct valuation of the risk of the underlying asset; which is another, simpler, but also quasi-understood problem; too many *degrees-of-freedom* from a quant's perspective; and too complex overall from a practitioner's perspective. Sprenkle (1964) was the standard prior to Black-Scholes:

$$C = SN(d_1) - (1 - Z)KN(d_2)e^{-rt} \tag{1}$$

where,

$$d_1 = \frac{ln\left(\frac{S}{K}\right) + \left(r + \frac{\sigma^2}{2}\right)t}{\sigma\sqrt{t}}, \qquad d_2 = d_1 - \sigma\sqrt{t}$$

C = value of the call option

S = underlying share price

$N(.)$ = cumulative standard normal distribution function

Z = degree of risk-aversion

K = strike price

r = growth rate of the underlying share price

t = time to expiration

σ = standard deviation of share price returns.

Of the model's inputs *S*, *K*, and *t* are known, while *r*, *Z*, and *σ* must be assumed. The Black-Scholes model simplifies this by setting the growth rate of the underlying share price, *r*, equal to the risk-free interest rate, and risk-aversion, *Z*, equal to zero; leaving only one non-observable input, *σ*, the volatility of the underlying returns. The rationale for this simplification comes from a no-arbitrage argument. If the model holds true, then the call option's *delta*, the ratio of the change in the option's price to the change in the underlying's price, is known:

$$\Delta = \frac{dC}{dS} = N(d_1). \tag{2}$$

This *delta* can be replicated by taking a leveraged position in the underlying stock. Therefore, a risk-free portfolio can be created by combining a *long* position in the *call* with a dynamically hedged *short* position in the leveraged underlying stock; dynamic, because the *delta* changes as the underlying share price changes (this is referred to as the option's *gamma*). Essentially, in this model, risk is priced as if the derivative's required return falls along the capital market line of the Capital Asset Pricing Model (CAPM):

$$r_C = r_f + \Delta_C \beta_S (r_{mkt} - r_f), \tag{3}$$

where,

r_C = required return of the option

r_f = risk-free rate

Δ_C = *delta* of the option

β_S = *beta* of the underlying asset

r_{mkt} = required return of the market portfolio.

While equation (1) is un-wieldy at first glance, its intuition is relatively simple. This was shown by Cox, Ross, and Rubinstein's (1979) binomial option pricing model, which, while a simplification of Black-Scholes, is actually more flexible (less parameterized), and is used by many option practitioner-traders today.

The binomial option pricing formula fits the following model:

[Insert Figure (1) Here]

The rationale for *risk-neutral* pricing can be found if one considers the model in Figure (1). There are *three* securities: a risk-free bond, a stock, and a *call* option on said stock. But, there are only *two* possible states at time *t,* an *upward* move or a *downward* move. The prices and payouts of the three securities *must* be linked to each other, or else there are

arbitrage opportunities to lock-in risk-free returns greater than the risk-free rate (an assumed absurdity!). Who knew the game of musical chairs was preparing children to understand option pricing?

In essence, the *call* option only exists for convenience; its price and payouts can be replicated by combining the prices and payouts of the risk-free bond and the underlying stock. Therefore, no *new* risk (or additional *degree-of-freedom*) is added by introducing a *call* option to the market. Suddenly, there is *hope*; options aren't so mysterious. They are but an organized means of bringing risk-management to everyone! If people want to buy options, market makers can provide the market, *without* having to take the opposite risk of their clients.

The difference between the binomial option pricing model and Black-Scholes is analogous to the difference between the binomial probability distribution and the normal distribution. The binomial model can be applied in multiple, time-period steps. Figure (2) shows this with two time periods[3]. At the limit, with an infinite number of time periods between today and the option's expiration date, the Black-Scholes model is derived, where the underlying asset is assumed to follow a continuous-time Brownian motion.

[3] The left side of Figure (2) could also be viewed as a trinomial option pricing model, via Boyle (1986).

Problems with the Standard Model.

The 2008 real-estate, mortgage crash is more recent in our memories than the 1987 stock market crash, but the two events have many similarities. New technology, whether part of the real-asset portfolio, or the 'intangible' portfolio, demands an amount of capital be employed towards its undertaking. How much of our overall capital should be employed to these new endeavors is something of a guess. The 'norm,' in American business history has been to provide too much capital to new markets, wait until the 'bubble' bursts, and then feed capital at a more restrained rate to the surviving firms. Figure (3) paints this picture.

Portfolio insurance became popular in 1986 and 1987, leading to a 542% and 128% rise in average, daily option volume traded in those two years respectively. In the 2008 real-estate crash it was *collateralized debt obligations* (*CDOs*) and *credit default swaps* (*CDSs*) that overwhelmed the markets, but in the mid-to-late 80's it was options. *Portfolio insurance* is essentially a *put* option on a basket of underlying stocks. The market makers of *portfolio insurance* offset their position by taking *short* index futures at the CBOT. The CBOT market makers offset their position by *shorting* individual stocks at the New York Stock Exchange (NYSE).

The Friday prior to Black Monday, Oct. 16th, 1987 was an option and futures expiration date. That week had been one of the worst weeks in stock market history, with the market losing over 10%. Due to record volumes of trade, many market makers were unable to fully hedge their positions prior to the close of trade that Friday, leaving a back-log of *short* orders to be filled Monday morning. Monday morning opened on time, but in a free-fall, the stock market losing 20% in hours. Tuesday morning most of the specialists on the NYSE couldn't open at all, having run out of trading capital, with banks un-willing to lend in such a chaotic market.

Blair Hull was Chairman and chief executive officer of Hull Trading Company, a prominent market maker at the CBOE. At 12:15 pm ET, on Tuesday, the CBOE and NYSE have halted trading, with the NYSE in agreement to open in an hour at a level that would put the Dow Jones Industrial Average (DJIA) 500 points lower (roughly another 20% drop). Hull happened to also be a member of the nearby CBOT, the only one of his 285-employee operation, a membership he had never before utilized. He walked over to the index futures pit to find only one contract with sellers, there were *no* buyers in any contract. It was the thinly traded MMI contract, consisting of 20 stocks, all members of the DJIA.

He started buying, un-hedeged *long* positions, and quickly started a *short*-squeeze, panic buying by all traders of the MMI, who were suddenly all *short* in a rapidly rising market. The other index futures quickly followed suit, and by 12:38 pm members of the NYSE forced their exchange to re-open, as they were being *short*-squeezed in the underlying stocks. The DJIA sky-rocketed, closing the day with a record *gain* of 102 points! Many describe it as the greatest, group euphoria across all exchanges imaginable; having gone to the brink and survived. The next day the DJIA would rise a new record 126 points! Within a month, the stock market was setting new *all-time* highs, and would continue on for twelve more years of a nineteen-year bull market run.

Hull was relying on an old, trader adage: "be *long* the halt." It's a *courageous* strategy that relies on cooler heads, and lenders of last resort, to prevail during halted trading. But that doesn't always happen; that Tuesday morning most stocks on the NYSE didn't open until after 11:00 am, only to be halted again an hour later. Prior to normal morning trading hours, the Federal Reserve had issued a statement affirming its readiness to serve as a source of liquidity to support the economic and financial system. But the statement is all they would provide throughout the crisis. In fairness, politically the Federal Reserve wouldn't want to start a *short*-squeeze like Hull did, nor do we know what would have happened had Hull's trade not been so effective, so quickly.

Hull's company used a proprietary model separate from Black-Scholes to price options. Regardless, Black-Scholes, or the binomial model was the standard option pricing model in 1987. The Black-Scholes model assumes stock prices follow a *log-normal* distribution; and therefore their returns a *normal* distribution. This means *zero* skewness and kurtosis of *three*. Figure (4) shows the z-score of daily returns of the S&P 100 in the week before and of the Black Monday crash. Returns are scaled by the *implied-volatility* index of the S&P 100 (*VXO*), which is an average, *option-implied* volatility across all S&P 100 options.

As can be seen, in *four* of the *ten* days, the absolute value of the *z*-score is greater than *two*; that's 40% of the returns landing farther than *two*-standard deviations away from center. In a standard, normal distribution this should only happen *once* every 22 days; the odds of it happening *four* times in *ten* days: 0.0001%. Moreover, there are *two*, *three*-standard deviation events in-a-row. The odds of *one, three*-standard deviation event: 0.135%; the odds of two-in-a-row: basically *zero*, by any practical limits. The skewness of returns over this ten-day period was -1.56, and kurtosis 6.98!

It's not that *portfolio insurance* and *standardized*-options were a bad idea, precisely the opposite, but that market-makers hope in the Black-

Scholes model to hedge their trading-risk was unfounded. The idea of no-arbitrage pricing makes head-sense, but who in their heart really believes the 'invisible-hand' of the marketplace will stick to a *normal* distribution? Nor were *CDOs* and *CDSs* necessarily a bad idea during the 2000's real-estate boom. But liquidity is a blessing not to be taken for granted. It can't be stored up, predicted, or replicated. It exists for the moment; tailor-made; in-valuable.

The Volatility Smile

Ever since Black-Scholes it has been common to quote option prices in terms of their *implied-volatilities*. Since the underlying asset's standard deviation, σ, is the only non-observable input to the model, if one knows the option *price*, they can 'back-out' the σ that would cause the model to price the option correctly. This is theoretically similar to a Method-of-Moments estimate assuming a normal distribution of underlying returns. There is no closed-form solution for this inverse problem, however, Manaster and Koehler (1982) developed a Newton-Raphson numerical, iterative-search method that guarantees convergence to an *implied-volatility*.

The CBOE calculates and distributes many *implied-volatility* indices, but the *VIX* index is the unrivaled bellwether of U.S. large-cap stock market

volatility. It estimates the average, annualized 30-day *implied-volatility* across all strike prices of the S&P 500 index. Its methodology, outlined in CBOE (2015), is to create a theoretical portfolio of *calls* and *puts* that should return the underlying price minus its futures price, *squared*. Basically, the 'portfolio' engineers a parabolic return distribution, and gives us the square-root of the price of that payout-scheme. As long as we believe in no-arbitrage pricing, it all makes a lot of sense. Figure (5) shows the history of this bellwether.

A casual glance at Figure (5) reveals that the market-wide *implied-volatility* is time-varying, mean-reverting, and prone to large, upward spikes in down markets. The upward spikes are evidence of negatively skewed and leptokurtic returns. However, the business-cycle adjusted, average volatility has remained relatively constant, around 25%. Another way to view this is that the long-run, average volatility did not adjust higher after the extraordinarily large spike upwards on Black Monday. However, this does not tell the whole story. The *VIX* index is largely driven by *implied-volatilities* at strike prices around the mean of the S&P 500 index's expected return distribution, not by returns in the tails of the distribution.

Another way to view *implied-volatilities* is through the *volatility smile*, a graph of *implied-volatilities* across strike prices for a single date and

expiration date. It's called a 'smile' because *implied-volatilities* tend to be higher away-from-the-money, in other words, away from the underlying asset's price. Figure (6) illustrates this for several asset classes.

If the Brownian motion assumption of the Black-Scholes model holds true, then *implied-volatilities* should be the same for *all* strike prices, in other words, there should be no 'smile,' but a flat line with a level equal to the single volatility of the Brownian motion. If a 'smile' exists, it indicates that the underlying asset's expected return distribution has higher probabilities of returns in the tails of the distribution than would exist in a normal distribution, in other words, that the returns are leptokurtic, a kurtosis higher than *three*. If the smile is higher on one side than the other, oftentimes called a 'smirk,' then there exists skewness in returns to the higher side, which again, should not exist if returns are normally distributed.

Rubenstein (1985), published before Black Monday, found no *volatility smile* in index or individual equity options, finding any variance in *implied-volatilities* to be within option bid-ask spreads, which admittedly, were relatively large at the time. Bates (1991) and (2000) revisited the problem and explicitly found that Black Monday was the advent of the *volatility smile*. It didn't exist before, but has existed in S&P 100 and 500 index options ever since!

And the phenomenon is by no means unique to equity index options. Mayhew (1995) and Toft and Prucyk (1997) found *volatility smiles* for individual stocks, although surprisingly not as pronounced as those in index options. Campa, Chang, and Reider (1998) and Bollen and Rasiel (2002) found that foreign exchange options exhibit *volatility smiles*, although again, not near as pronounced as those in index or even individual equity options; however, currency *implied volatilities* are much lower than equities' in general. Jarrow, Li, and Zhao (2003) likewise found *volatility smiles* in the interest rate options market. Figure (6) also shows 'smiles' for gold, oil futures, and a 'smirk' for options on volatility itself. The positive skewness implied in the *VIX* index correlates with the large, positive spikes found in Figure (5).

Risk-Neutral Pricing

The fact that index options exhibit a larger 'smile' than individual stocks illustrates an important aspect of option pricing. A large tail event in the index is more likely due to a *liquidity*-risk event, whereas tail events in individual stocks could very well be unique to only that stock. Since the market-wide liquidity event is more 'important' to investors, its option-implied probability is higher than an otherwise equally-likely event with less 'importance.' While the reason for this is not immediately obvious, it

has to do with *risk-neutral* pricing, briefly discussed before in relation to Figure (1).

Note in Figure (1) that the underlying stock returns either $e^{rt+\sigma\sqrt{t}}$ or $e^{rt-\sigma\sqrt{t}}$. 'r' in this case represents the expected return of the stock. In turns out, in solving for the option's price, the choice of r is unimportant; it drops out of the final solution. Most textbooks set r equal to zero in the binomial model for simplicity, hence the nomenclature U and $1/U$ for $e^{\sigma\sqrt{t}}$ and $e^{-\sigma\sqrt{t}}$. However, if one really wants to maintain the binomial model's comparability with Black-Scholes, r should be set to the risk-free rate.

If one sets r equal to the risk-free rate, then both the bond and the stock have the same required return. Since these two securities 'complete' the market (since they are *two* unique securities in a *two*-state market), if their returns are the same, then *every* security in this market will have that same required return, in this case, the risk-free rate! The *call* option described by the middle tree in Figure (1) is another security in this market; therefore it too should return the risk-free rate. Equation (4) states this using the probabilities implied by the binomial model:

$$C_t = p_U\left(Se^{r_f t+\sigma\sqrt{t}} - K\right) + (1-p_U)\emptyset = C_o e^{r_f t}. \tag{4}$$

The value of the *call* at its expiration date, C_t, is equal to the probability of an *upward* move times its value in an *upward* move, plus the probability of a *downward* move times its value in a *downward* move, which equals the initial *call* price grossed up at the risk-free rate. In fact *any* security's value follows this same model:

$$X_t = p_U X_{t,U} + p_D X_{t,D} = X_o e^{r_f t}, \tag{5}$$

where X_t is the value of said security at time *t*, $X_{t,U}$ is its value in an *upward* move, and $X_{t,D}$ in a *downward* move. If there are more than *two* possible future states of the market, the future value of the security is a weighted average of the security in each state, weighted by the probability of each state:

$$X_t = \sum_i p_i X_{t,i} = X_o e^{r_f t}. \tag{6}$$

Note that while the Black-Scholes model assumes an *infinite* number of possible future states of the market (compared to only *two* for the binomial model), its Brownian motion assumption means the probabilities of those future states follow a normal distribution, a distribution with only *two* moments: a mean and a variance. Because there are only *two degrees-of-freedom* in defining a normal distribution,

the same two securities, a risk-free bond and an underlying stock returning the risk-free rate, are enough to 'complete' the market. That said, the future value of a security whose returns follow a continuous probability distribution function (*pdf*) can be described in the continuous state equivalent of equation (6):

$$X_t = \int X_t(i)\theta(i)di = X_o e^{r_f t}, \tag{7}$$

where $\theta(i)$ is the *pdf* of X_t.

One obvious concern should arise when considering *risk-neutral* pricing: in the *real*-world, neither the underlying stock nor the option should return the risk-free rate! Equations (4) through (7) do <u>not</u> represent pricing in the *real*-world; they represent pricing in what is often referred to as the *risk-neutral* world. It is an invented world created by assuming the stock returns the risk-free rate. This is purely for pricing convenience; in this invented world one never needs to know investor risk-aversion or risk-premia for risky assets.

However, this convenience comes at a cost. The probabilities implied by the binomial model, p_U and $(1 - p_U)$, are not *real*-world probabilities; they are referred to as *risk-neutral* probabilities. They are derived in such a way that forces the stock to return the risk-free rate. The *real*-world

probabilities would have the stock return the risk-free rate plus a risk premium, based on the co-variance of the stock with the market and the market-wide appetite for risk. For example Sprenkle's formula, Equation (1), prices the option in the *real*-world. All of the probabilities in Equations (4) through (7), including also those in the Black-Scholes formula, are *risk-neutral* probabilities.

One way to describe this is the following:

$$X_0 = E^P[X_t]e^{-rt} = E^Q[X_t]e^{-r_f t}, \tag{8}$$

where $E^P[.]$ is the expected value using *real*-world probabilities, often referred to as the *P* measure, and $E^Q[.]$ is the expected value using *risk-neutral* probabilities, referred to as the *Q* measure. '*r*' is the risk-adjusted required return *specific* to asset *X* in the *real*-world, while 'r_f' is the risk-free rate and the required return for *all* assets in the *risk-neutral* world.

Cochrane (2005) derives the relationship between *real*-world probabilities and *risk-neutral* probabilities:

$$p^Q(i) = p^P(i)\frac{U'(i)}{E^P[U'(.)]}, \tag{9}$$

where, $p^Q(i)$ is the *risk-neutral* probability of state i, $p^P(i)$ is the *real*-world probability of state i, $U'(i)$ is the market representative's marginal utility curve in state i, and $E^P[U'(.)]$ is the market representative's *expected* marginal utility, averaged across all states.

It turns out that risk preferences do show up in the Black-Scholes and binomial models (and any other risk-neutral pricing models) -- they just show up in the probabilities rather than the required return. Breaking down Equation (9), there are *two* components to *risk-neutral* probabilities: one due to rational expectations: the *real*-world probabilities; the second due to subjective preferences in the form of the marginal utility curve (scaled by average marginal utility), which represents the market-representative's subjective value of an additional dollar added to her returns (or wealth). *It is as if both the head and heart play a role in determining probabilities*.

The higher the marginal utility (the subjective value of returns), the higher the *risk-neutral* probability; in other words, states of the world deemed more 'important' to the representative investor get a higher *risk-neutral* probability than their corresponding *real*-world probability. Examples of people making decisions based on *risk-neutral* probabilities are common. Consider that people consistently over-estimate the

probability of a plane crash[4]. Consider that American's spend more money on the lottery [5] (an objectively negative net-present-value investment) than all other forms of entertainment combined. In an extreme example, consider that no one would take 5-to-6 odds on a game of Russian roulette!

Bernoulli (1738) used decreasing marginal utility to solve the St. Petersburg paradox. Consider a coin-flipping bet which pays-out based on the number of *heads* flipped in a row, doubling with each additional *heads* thrown (for example: it might payout $0 if a *tails* is flipped on the first toss; $100 if only *one heads* is thrown, followed by a *tails*; $200 if *two* heads are thrown; $400 for *three*, $800 for *four*, and doubling so on). One finds, using real-world probabilities of a fair coin, the expected payout of this bet is *infinite*! But who would pay more than a few hundred dollars for it?

The solution to the paradox is that people's marginal utility of returns decreases the richer they get. A marginal dollar is more 'important' to a poor person than to a rich person. Thinking in terms of the binomial model, if there is an *upward* move, the representative investor is better off, and therefore values the marginal dollar less than she would have in a *downward* state. Summing up, *risk-neutral* probabilities of a

[4] See Clark and Rock (2016).
[5] $73B in 2017 according to Gallup, with an expected loss of $0.40 on the dollar.

downward move are relatively higher than probabilities of an *upward* move, compared to their *real*-world counterparts. This causes the asset-specific, risk-adjusted expected returns of *all* risky assets to decrease exactly to the risk-free rate when using *risk*-neutral probabilities.

APPLICATIONS

Before delving into the mechanics of estimating option-implied *pdf*s, it behooves us to look into some *applications* of such *pdf*s. Foremost, knowledge of the *pdf* is desirable because it enables the pricing of *any* derivative of the underlying asset with the same time to expiration, regardless if it be illiquid or 'non-vanilla.' Beyond this obvious application, Table (3) summarizes three fronts of applications of option-implied *pdf*s.

For example, central bankers use option-implied *pdf*s to assess market sentiment over future changes in interest rates, exchange rates and stock prices, and likewise to confirm the market's acceptance of major policy decisions. Of particular interest is the predictive power of options prior to major economic events such as crash episodes, exchange rate/interest rate regime changes, wars, and elections. Option-implied *pdf*s are fruitful for this type of analysis because unlike other market-based time series data, which in isolation capture only an expectation

(and therefore represent only one expected economic scenario), *pdf*s capture the uncertainty that is fundamentally inherent in the marketplace -- that there are multiple future scenarios.

A third line of research uses option-implied *pdf*s to measure risk aversion. The risk-neutral density is estimated from option prices, while the objective density from historical returns; together these two infer implied-relative risk aversion. Knowledge of economy-wide risk aversion makes possible value-at-risk (VAR) and other forecasting applications in risk management.

THEORY

The volatility smile may invalidate the Black-Scholes model, but not risk-neutral pricing in general. *Real*-world probabilities and risk-adjusted required rates of return cannot be solved for without having knowledge of market-wide utility preferences. The subjective nature of these preferences makes them difficult to estimate (or to even grasp their full nature), which is therefore the reason investors found such *hope* in Black-Scholes and the advent of risk-neutral pricing.

However, the volatility smile is *full* of information about the underlying asset's risk-neutral *pdf*. *Any* asset with 'enough' option-prices (to

'*dynamically*-complete' the market; i.e. enough options at different strike prices), should have an estimable underlying risk-neutral *pdf*; and therefore, the price of *any* 'non-unique' security (i.e. a derivative) should also be estimable. This endeavor has led to research in the inverse problem to option-pricing: using known, market-set option-prices to estimate the underlying asset's *pdf*.

The Inverse Problem

Arrow (1964) showed that *risk-neutral* probabilities are unique and can be solved for in any 'complete' market model; i.e. a model with the same number of unique securities as states of the future world (such as a stock and a risk-free bond in a world with *two* states, as is the case with the binomial model). Merton (1971) showed that risk-neutral probabilities are unique and can be solved for in any '*dynamically*-compete' market model (such as a stock and a risk-free bond in a world with *two degrees-of-freedom*, as is the case with Black-Scholes and its normally-distributed returns).

Figure (7) gives an example of this inverse problem of using option prices to estimate *risk-neutral* probabilities under the binomial model.

The problem with Black-Scholes is its assumption of normally-distributed returns, which is too simple to account for liquidity-risk, time-varying volatility, volatility spikes, and so forth. The lack of an historical option-price database, and relatively poor computing power, meant the inverse option-pricing problem, of using option-prices to estimate the asset's implied risk-neutral *pdf*, unfeasible until the mid-1990s.

There are *two* polar approaches to estimating the *pdf* from option prices. One is based on relatively strict assumptions, a model, which the observed option-prices must to adhere. The opposite pole assumes no model, but a 'sort of' open-canvas for the option-prices to reveal their nature. There are trade-offs between accuracy, bias, stability, and so-forth, and as is usually the case, the best solution lies somewhere in-between the *two* poles.

Model-Based Methods.

How many option-prices are needed to '*dynamically*-complete' the market? According to Fama and French (1993) there are only *three*-factors driving market returns (meaning a stock, the risk-free bond, and *one* option are enough to '*dynamically*-complete' the *three degrees-of-freedom*). At most, a handful, *five*-factors, are cited in a single model by

researchers (meaning *three* options are needed). In other words, there aren't *that* many *degrees-of-freedom*. The Black-Scholes *two* aren't enough, but a *slightly* more sophisticated model may be enough.

Generalized distribution methods.

The *beauty* of a *model* is in its simplicity, understandability, and manipulability. For example, the analogy of a coin-flipping exercise to the normal distribution is easy for most to grasp. However, the cost is the normal distribution's naivety; it falls short in fully describing the reality of market returns. There are many probability distributions with more than *two* parameters; Table (1) gives those used to model option-implied *pdf*s. However, the reasoning for *why* these more robust models should be used is seldom explained. They are chosen simply because they allow for more *degrees-of-freedom*. The beauty of having only a few *degrees-of-freedom* is lost if no one understands the nature of these few, new *degrees*. Why not just build a model with *many degrees-of-freedom*, if that is the goal?

Alternative Return-Generating Processes.

Another method used to add robustness to the standard Black-Scholes model is to assume a return-generating process more robust than Brownian motion, then find a way to derive a closed-form solution

similar to the Black-Scholes equation but with more *degrees-of-freedom*. Factors that Black-Scholes doesn't address, such as liquidity-risk, time-varying volatility, and volatility spikes, can correspondingly be modelled with time-series components such as stochastic interest rates, stochastic volatility, and the Poisson-jump process.

The standard Black-Scholes' Brownian motion is based on the following stochastic differential equation:

$$\frac{dS_t}{S_t} = r_f dt + \sigma dW_t, \tag{10}$$

where S_t is the underlying security's price, r_f is the *constant* risk-free rate, σ is the volatility, and dW_t is a Wiener process (Brownian motion) representing a standard normal distribution (mean of *zero* and standard deviation of *one*). This has a closed-form solution of:

$$\ln(\frac{S_t}{S_0}) = \left(r_f - \frac{\sigma^2}{2}\right) t + N(0, \sigma\sqrt{t}), \tag{11}$$

where $N(0, \sigma\sqrt{t})$ is the normal distribution with mean *zero* and standard deviation $\sigma\sqrt{t}$. The correction term, $-\frac{\sigma^2}{2}$, effectively accounts for the difference in compounding between a continuous process and a periodic process.

Bakshi et al. (1997) introduce a model where the interest rate is not constant, but instead has its own stochastic differential equation which follows a Vasichek (1977) mean-reverting process:

$$dr = \theta_r(\omega_r - r_t)dt + \varepsilon_r dB_{r,t}, \tag{12}$$

where ω_r is the mean long-term risk-free rate; θ_r the rate of mean-reversion; ε_r the standard deviation of the risk-free rate process, and $dB_{r,t}$, like dW_t, a normal distribution with mean *zero* and standard deviation ε_r. There is also an assumed constant correlation of $\rho_{S,r}$ between dB_t and dW_t. Equation (12) has a closed-form solution of:

$$r_t = r_0 + \theta_r(\omega_r - r_0)t + N(0, \varepsilon_{r,0}\sqrt{t}), \tag{13}$$

There are *three, new degrees-of-freedom* in this model, compared to the standard Brownian motion: θ_r, ε_r, and $\rho_{S,r}$. Each period's risk-free rate is a weighted average of the long-run risk-free rate, ω_r, and the prior period's risk-free rate:

$$\mathrm{E}\left[r_{f,t}\right] = (\theta t)\omega + (1 - \theta t)r_{f,0}, \tag{14}$$

Similar to stochastic interest rates, a number of stochastic *volatility* models have been developed; Heston (1993) being the most popular. In the Heston model volatility is not constant, but instead, the underlying's variance has its own stochastic differential equation which follows a Cox, Ingersoll, and Ross (*CIR*) (1985) mean-reverting process:

$$dv_t^2 = \theta_v(\omega_v - v_t^2)dt + \varepsilon_v v_t dB_{v,t}, \tag{15}$$

where ω_v is the mean long-term variance, θ_v is the rate of variance mean-reversion, ε_v is the standard deviation of the variance process, and $dB_{v,t}$, like dW_t, represents a normal distribution with mean *zero* and standard deviation v_t. There is also an assumed constant correlation of $\rho_{S,v}$ between $dB_{v,t}$ and dW_t. The difference between the *CIR* mean-reverting process and the Vasicek process is that the standard deviation of the *CIR* process is dependent on the most recent volatility of the Brownian motion. Equation (15) has a closed-form solution of:

$$\sigma_t^2 = \sigma_0^2 + \theta_v(\omega_v - \sigma_0^2)t + N(0, \varepsilon_v \sigma_0 \sqrt{t}), \tag{16}$$

Bates (1996, 2000, 2001) develops a stochastic model which includes a Merton (1976) Poisson-jump process:

$$\frac{dS}{S} = r_f dt + \sigma dW_t + kdq, \tag{17}$$

where k represents the size of jumps and dq a Poisson-counter. The jump-diffusion *call* option price, $Call_{J-D}$, is an adjustment of the standard Black-Scholes *call* option price, $Call_{B-S}$:

$$Call_{J-D} = \sum_{k=0}^{\infty} \frac{e^{-m\lambda t}(m\lambda t)^k}{k!} Call_{B-S}, \quad (18)$$

where m is the average jump size, and λ is the average number of jumps per time.

Three degrees-of-freedom can be added with stochastic interest rates, another *three* with stochastic volatility, and *two* (or *more*) with jumps; however, it is a challenge to get *any* of these time-series models to consistently converge to statistically-stable parameters; let alone if they are combined into a general model. So, from a practical stand-point, it is uncertain which of these time-series components should be chosen to add a few, *new* degrees of freedom to the standard Black-Scholes model.

Expansion methods.

Expansion methods start with a simple, known probability distribution (usually the normal or log-normal) and then add approximating terms similar to a Taylor expansion. The distribution's cumulant-generating

function's derivatives determine the distribution's higher-order moments, similar to the way higher-order derivatives of a function are used in a Taylor expansion. This method by no-means guarantees a well-behaved distribution (i.e. one whose probabilities are strictly-positive and integrate to *one*). Table (2) gives the various expansion methods used to model option-implied *pdf*s.

Mixture methods.

Yet another method that adds complexity to simple probability distributions is to create a weighted-average of *two* or more simple distributions. For example, Ritchey (1990) averages *two* normal distributions, which results in *five degrees-of-freedom*: *two* means, *two* standard deviations, and a single weighting parameter. Note that the result is a complex distribution with a single mean, variance, skewness, kurtosis, and a fifth moment. Melick and Thomas (1997) mix *three* log-normal distributions, which results in *eight degrees-of-freedom*.

Of all the model-based (parametric) methods, the *mixture methods* are: the simplest; easiest-to-understand; and, most-stable and consistently-statistically estimable; *yet*, *for whatever reason*, they are the most *seldom* used methods. One might even find *beauty* in the Ritchie model, a relatively simple distribution with *two*, separate spirits of uncertainty,

which is surprisingly robust to model samples that would otherwise appear only chaotic[6].

Non-Parametric Methods.

Non-parametric methods embrace the *mystery* that shrouds option-implied *pdf*s. Instead of treating asset returns as something to be modeled with as few *degrees-of-freedom* as possible, the returns are treated as a 'black-box,' *un-model-able* by nature. The number of *degrees-of-freedom* is not important, so long as the data is not 'over-fit;' that market micro-structure noise such as bid-ask spreads, sparse trading, minimum tick sizes, and the like don't cause unstable *pdf* estimates.

Table (4) summarizes papers from three non-parametric methods used to estimate option-implied *pdf*s. The first method, maximum entropy, directly embraces the idea of chaos (or a 'black box') being a natural part of a risk model. The next *two* methods fit a curve to the *IV* smile and back out probabilities based on Ross (1976), Breeden and Litzenberger (1978), and Banz and Miller (1978), who showed how to extract a *pdf* from a smooth set of option prices.

[6] Many have claimed this model graces the heavens as a model of mankind, in the form of *two* fishes (spirits) tied together at the tails -- Pisces -- in the chaotic, wintery quarter of the Zodiac, exact opposite of the pure, un-adulterated truth of Virgo.

Maximum entropy methods.

The term entropy was created to describe disorder in thermodynamics. It was later adapted to describe the unknown in information science. To *maximize* entropy is to *minimize* assumptions about the unknown. It is usually assumed that if one knows nothing about a probability distribution, then all probabilities should be equal (a uniform distribution; something rarely found in naturally created distributions). The idea is to know what you don't know, an admirable ambition, and pick the distribution that reflects that.

One usually does assume a prior distribution (normal or lognormal), then fits the data as close to it as possible based on minimizing the following error function:

$$\varepsilon = \sum p_i ln\left(\frac{p_i}{a_i}\right), \tag{19}$$

where p_i is the estimated probability and a_i is the assumed probability. It basically amounts to minimizing a weighted average of the error $ln(p_i/a_i)$.

Minimizing this function is unstable due to its non-linearity, and the fact that the logarithm can balloon to very large positive and/or negative numbers, which in turn, over-whelm and drive the estimation process.

Kernel methods.

Kernel methods sound complicated, but they are basically moving average curve-fitting methods. At each observation along the *implied-volatility* smile the current observation is assumed to be the mean of a normal distribution that fits the observations around it based on a chosen 'bandwidth,' which is the standard deviation of that normal distribution. The wider the bandwidth the 'smoother' the curve is allowed to be. A small bandwidth creates a 'tight' curve that comes close to all the points but has much 'jaggedness.' A large bandwidth is allowed to miss the mean (individual observation) for the sake of 'smoothness' of the overall curve.

The kernel method is the first method we have inspected that truly abandons *any* assumption of a prior distribution. The difference between it and the next methods we will examine, is that the kernel method fits each localized point -- it actually doesn't need data far away from it to estimate a localized kernel.

Curve-fitting methods.

If the kernel methods are a bottom-up approach (from each observation), the curve fitting methods are a top-down approach (one function to fit all the observations). Before we describe them, it behooves us to show why having a smooth set of *implied-volatilities* is desirable. The *implied-volatilities* are just a *transformation* of the option prices, transformed such that they do not tail off to *zero*-value when out-of-the-money, but are comparable across strike prices.

Ross (1976), Breeden and Litzenberger (1978), and Banz and Miller (1978) showed the relation between an asset's risk-neutral *pdf* and its option prices. Given the value of a call option,

$$Call = e^{-rt} \int_K^{\infty} (S_t - K) \phi(S_t)\, dS_t, \tag{20}$$

where K is the strike price, S_t is the underlying price, r is the risk-free rate, t is the time-to-expiration, and $\phi(S_t)$ is the risk-neutral *pdf*;

taking the partial derivative of (20) with respect to strike K, gives,

$$\frac{\partial Call}{\partial K} = e^{-rt}[1 - \Phi(K)], \tag{21}$$

where $\Phi(K)$ is the risk-neutral cumulative distribution function (*cdf*). Solving for $\Phi(K)$,

$$\Phi(K) = e^{rt}\frac{\partial Call}{\partial K} + 1, \tag{22}$$

and taking the partial derivative with respect to strike K a second time yields the risk-neutral *pdf*,

$$\phi(K) = e^{rt}\frac{\partial^2 Call}{\partial K^2}. \tag{23}$$

A similar derivation (or *put-call* parity) can be used to find the risk-neutral *cdf* and *pdf* from *put* prices,

$$\Phi(K) = e^{rt}\frac{\partial Put}{\partial K}, \tag{24}$$

$$\phi(K) = e^{rt}\frac{\partial^2 Put}{\partial K^2}. \tag{25}$$

Curve-fitting methods are necessary because options only come at discrete strike prices. They interpolate across the discrete strike prices, to create a 'smooth' set of *continuous* option-prices, whereby equations (22-25) can be used to calculate the *cdf* and *pdf*. Panel (C) of Table (4) summarizes some of the curve-fitting methods used to estimate option-

implied *pdf*s. These curve-fitting methods are fast and stable and have become a staple in research using option-implied *pdf*s[7].

THE TAILS OF THE DISTRIBUTION

Of the before-mentioned methodologies for estimating risk-neutral *cdf*s and *pdf*s, we find rest with *two* of the methodologies. If one desires a *model*-based method, the *mixture methods* are: the simplest; easiest-to-understand; and, most-stable and consistently-statistically estimable. If one allows for a non-parametric method, the *implied-volatility* curve-fitting methods are the fastest and most stable.

However, neither of these methods are particularly good at estimating probabilities in the tails of the *cdf*s and *pdf*s. In the case of *mixture methods*, the tails usually revert to one of the simple, inputted distributions, which amounts to a normal or log-normal tail; which is known, from the Black Scholes volatility-smile, to *not* be the correct shape for most assets' tails (see Figure (6)). Curve-fitting methods interpolate well *between* strike prices, but are *not* intended to extrapolate into the tails of the distribution, where there is no tradable, strike-price data.

[7] See Figlewski (2008) for an example.

The tails of the distribution are highly influential on estimates of higher-order moments of the *pdf*, such as skewness and kurtosis; and on estimates of tail-probabilities, such as those used in value-at-risk (*VAR*) applications.

As a base case, the simplest means to account for the tails is to simply truncate them. The probability of an event occurring outside the available range of strikes can be assumed to be *zero*. The remaining interior distribution can then be re-weighted to force the *cdf* to sum to *one*. See Anagnou et al. (2002) for an example of this. This method is a worst-case scenario, i.e. a floor to measure the relative improvement of other methods.

The most *common* means to account for the tails is to assume they are normal or log-normal (see Jiang and Tian (2005) for one example). This means a constant Black-Scholes *IV* for all strikes below the lowest observable strike, and a separate constant *IV* for all strikes above the highest observable strike. But again, the tails are known to be non-normal.

Theoretically, the Pareto distribution is the correct distribution to model the tails of an *unknown* distribution. However, it is overly conservative; using it often leads to *infinite* estimates of *implied-volatility*, or at the very

least, wildly unstable estimates. The methodology used to calculate the *VIX* index shows that empirically this shouldn't be the case. Probabilities in the tails of the implied-*pdf* 'fall' much faster than the *z*-scores of the underlying's payoffs 'rise.' In other words, the tails aren't even important in calculating *implied-volatility*; *implied-volatility* is driven primarily by option-prices near-the-money.

Therefore, the Pareto distribution's cousin, the extreme value distribution (*EVD*) is more commonly used to model the *pdf*'s tails. Technically, the *EVD* models the single, most extreme observation of the distribution, not the entire tail. However, it has *three degrees-of-freedom*, which when estimated separately for the low and high tails, is a total of *six* -- much more robust than the Black-Scholes' *two degrees-of-freedom*[8].

Finally, a new tail-fitting methodology could be introduced based upon the same curve-fitting logic used to interpolate between option strikes. However instead of interpolating between strikes, the *IV* smile could be used to extrapolate beyond the available range of strikes. Heuristically, the *IV* smile could be extended to the left and right based on its linear trend. This would provide a non-parametric tail equivalent to the fast

[8] Theoretically, there should only be *five degrees of freedom* if one knows the forward-price of the underlying; only *four* if one also knows its *implied-volatility*. Both these conditions link the lower tail to the upper tail. However, there is no known closed-form solution or numerical method to assure these conditions.

and stable curve-fitting methodology used to model the middle of the distribution.

CHAPTER III

DATA AND METHODOLOGY

S&P 500 index option bid-ask quotes with monthly expirations are obtained from Optionmetrics through the WRDS system. Data is available from January 1, 2003 thru December 29, 2017. The risk-free rate is interpolated from the zero curve as obtained from Optionmetrics which is derived from BBA LIBOR rates and settlement prices of CME Eurodollar futures. The forward price is interpolated from option prices (see Appendix 1).

Summary statistics are given in Table (5). The statistics suggest that the spline fit the data well. At a minimum, five strikes are needed to fit a fourth order spline with one knot. The fewest strikes available were eleven. The average number of strikes available was 68.4 and the maximum was 164. The large number of strikes correlated with a wide range of strikes. The minimum of the cumulative distribution function, as obtained only from the spline fitting process (no affixation of tails)

averaged 0.0090 with a third quartile of 0.0085. However there does appear to be some outlier days where there were too few strike prices available in the lower tail. The same pattern of outliers appears for the range of strikes in the upper tail. The statistical fit of the spline-fitting process was very good: the daily *r*-squared of the spline regression averaged 0.9990 with a minimum of 0.9742.

ILLUSTRATIVE EXAMPLE

To illustrate the before-mentioned methodologies, this section steps through the RND estimation process for a single day.

Step 1: Sort raw data.

Table (6) provides the raw option price data used to estimate the RND of the S&P 500 index on January 31, 2012. One expiration date is chosen to estimate the RND, which for Table 1 is March 17, 2012, which is 46 days out. The nearest to expiration date is chosen that is greater than two weeks. The two week buffer follows Figlewski (2009) and is a choice of balance: as options approach expiration, the prices of away-from-the-money options approach zero and fall out of the available range of strike prices, making estimation of the tails more difficult; however, longer maturity options tend to have less liquidity.

All options used in the analysis are required to have a minimum bid of 0.05 and the maximum distance allowed between strike prices is 25 points (if there is a jump in strike prices greater than 25, all options further away from at-the-money are dropped). This eliminates some erroneous quotes in the extreme tails of the distribution.

Step 2: Transform option prices into implied volatilities.

It is easier computationally to work with Black-Scholes IVs rather than actual options prices. Table 1 gives the IV that corresponds to the bid-ask midpoint based on the Black-Scholes model given the forward price (see Black (1976))

$$C = [FN(d_1) - KN(d_2)]e^{-r_f t} \tag{26}$$

$$d_1 = \frac{\ln\left(\frac{F}{K}\right) + (\sigma^2/2)t}{\sigma\sqrt{t}}, \quad d_2 = d_1 - \sigma\sqrt{t} \tag{27}$$

$$P = [KN(-d_2) - FN(-d_1)]e^{-r_f t} \tag{28}$$

where F is the forward price as obtained from options prices. The forward price is the strike where the call price equals the put price (see

Appendix 1 for the details of obtaining the forward price from option prices). Using the forward price has advantages over the traditional Black-Scholes model in that the underlying price and its dividend yield are not required. Daily closing prices of the underlying do not necessarily line up at the same exact time with closing prices of its options, and the dividend yield is not directly observable. The only data besides option prices required in equations (26-28) is the risk-free rate, which is interpolated from the LIBOR zero curve.

Step 3: Fit spline to implied volatility smile.

To interpolate between observable strike prices, a fourth-degree spline with a single knot at-the-money is fit to the IV smile. Figure (8) shows the IV smile with both call and put prices. Following Figlewski (2009), put IVs are used for below-the-money strikes and call IVs for above-the-money, while blending put and call IVs within 3% of at-the-money. The spline is fit to the midpoint of the bid-ask volatility spread. The bid-ask midpoint is preferred to transaction data which is often sparse in options markets. Figure (9) shows the spline fitted to these IVs. Note that the fourth degree spline models the complex shape of the IV smile without over-fitting noise in options prices.

Step 4: Estimate risk-neutral density from interpolated prices.

Next, the spline is converted back into a dense set of options prices via equations (26-28). Following Figlewski (2009), I build a set of options prices at one cent strike intervals. Then, using equations (22-25), the risk-free cumulative distribution and density function can be numerically obtained at one cent intervals. Figure (10) gives the resulting RND between the 2nd percentile and 98th percentile. Note the obvious presence of negative skewness.

Step 5: Affix tails to the distribution.

The above method approximates the RND within the range of available strike prices. However a separate method must be chosen to estimate the tails of the distribution, which lie outside of the range of available strike prices. The tails of the distribution are highly influential on estimates of higher moments of the RND, such as skewness and kurtosis; and on estimates of tail probabilities such as those used in value-at-risk applications. Four potential tail estimation methods are described below, the last being a new methodology developed for this thesis.

Truncated tails.

As a base case, the simplest tail methodology is to assume a truncated distribution. The probability of an event occurring outside the available range of strikes is assumed to be zero. The remaining interior distribution is then reweighted to force the cumulative distribution to sum to one. See Anagnou et al. (2002) for an example of this. This method is included to provide a worst-case scenario, i.e. a floor to measure the relative improvement of all other methods.

One assumes there is zero probability of an event occurring outside some range, which I choose to be the 2nd and 98th percentiles. The truncated RND then becomes

$$f_{truncated}(K) = f_{non-truncated}(K)/0.96 \tag{29}$$

The denominator has to be adjusted if strike prices are not available over the entire 2nd to 98th percentile range. This method is crude, but provides a worst-case scenario to compare more complex methods against. Figure 4 displays the difference between the spline-estimated RND and the truncated RND.

Lognormal tails.

The most common method for attaching tails to the distribution is to assume they come from the tails of a lognormal distribution (see Jiang and Tian (2005) for one example). This amounts to choosing a constant Black-Scholes IV for all strikes below the lowest observable strike, and a separate constant IV for all strikes above the highest observable strike. These IVs are usually chosen to match the observed IVs at the barriers of the IV smile, thus keeping the option price function smooth. However the first and second derivatives of the option price function (representing the cumulative probability distribution and density function) usually jump at these barriers. Figure (12) shows an example of this.

Generalized extreme value tails.

Figlewski (2009) attaches tails taken from a Generalized Extreme Value distribution (GEV). Similar to the way the Central Limit Theorem makes the normal distribution the natural candidate for modeling a distribution's average, the GEV is the natural candidate for modeling the largest value of an unknown distribution. It has three parameters: a location parameter, a scale parameter, and a shape parameter which determines the amount of weight placed within the right tail.

Figlewski attaches GEV tails to the density at the 5th and 95th percentiles. The parameters of the GEV tails are chosen such that the GEV and RND have the same cumulative distribution and density function at the 5th and 95th percentiles, and also so that the density functions match at the 2nd and 98th percentiles. Namely, for the upper tail,

$$F_{GEV}(K_{0.95}) = F_{RND}(K_{0.95}) \tag{30}$$

$$f_{GEV}(K_{0.95}) = f_{RND}(K_{0.95}) \tag{31}$$

$$f_{GEV}(K_{0.98}) = f_{RND}(K_{0.98}) \tag{32}$$

where K_p is the strike price at the p^{th} percentile. This gives three degrees of freedom to each tail, to match the three parameters within each tail's GEV. Note that the GEV only models upper tails, and therefore when fitting the lower tail of the RND, the cumulative distribution and the strike prices must be converted into their upper tail equivalents, namely

$$F_{GEV}(-K_{0.05}) = 1 - F_{RND}(K_{0.05}) \tag{33}$$

$$f_{GEV}(-K_{0.05}) = f_{RND}(K_{0.05}) \tag{34}$$

$$f_{GEV}(-K_{0.02}) = f_{RND}(K_{0.02}) \quad (35)$$

If the range of available strike prices does not extend to the 2nd and/or the 98th percentile, the outermost available percentile is used for equation (35) and a percentile three less than that is used for equations (33) and (34). For example, if the highest observable percentile is the 95th, then the 95th and 92nd percentiles are used to match the distributions.

Smile-extrapolated tails.

The new tail-fitting methodology introduced in this thesis is based upon the same logic used to interpolate between option strikes. However instead of interpolating between strikes, the IV smile is used to extrapolate beyond the available range of strikes. Bliss and Panigirtzoglou (2004) use a similar methodology, however they forces the IV smile to become horizontal as strikes go towards zero or infinity. This is essentially equivalent to the lognormal tail methodology. I propose to extrapolate the IV smile based on its linear trend.

Specifically, the IV smile is extended to the left and right based on its linear trend between the 2nd and 5th percentiles (for the left tail), and 95th and 98th percentiles (for the right tail). Similar to Figlewski's method, if

the range of available strikes does not extend to the 2nd and/or the 98th percentile, the range of three percentiles farthest to the edge of the available strike prices is used to capture the linear trend. Once the IV is extrapolated, it is converted back into a dense set of option prices and the RND is obtained numerically via equations (22-25).

Figure (12) shows the IV smiles that result from the lognormal tails method and the smile-extrapolation method. Between the 5th and 95th percentiles, the IV smiles for both methods are based solely off the fitted spline. For the smile-extrapolation method, outside the 2nd and 98th percentiles, the smile is based solely off the high and low tail linear trends. Within the trend estimation zones, the smile is a blend of the fitted spline and the linear trend.

For the lognormal tails method, tails from a lognormal distribution are affixed at the 2nd and 98th percentiles such that the option price function remains smooth. In terms of the volatility smile this amounts to affixing a flat line that connects to the fitted spline at the 2nd and 98th percentiles, as can be seen in Figure 5. Note that the right tails are nearly identical between the lognormal tail method and the smile-extrapolation method. However, the left tails are drastically different as the lognormal tail forces the IV smile to take a sharp downward kink, which obviously does not reflect the observed option data. This problem

leads to instability when using the lognormal tail method to estimate higher moments of the distribution and/or tail probabilities. The exact location of the 2nd and 98th percentiles is affected by market imperfections such as the time to maturity of the options. The smile-extrapolation method is not nearly as dependent on the exact location of the 2nd and 98th percentiles.

Figures (13) and (14) compare the density function and cumulative distribution of lower tails that result from the GEV method and the smile-extrapolation method. Within Figure 6, note that the two methods are guaranteed to overlap at the 2nd and 5th percentiles, as a result of the GEV fitting process. In Figure (14) the two methods are guaranteed to overlap at the 5th percentile. However, outside of those three points the two methods could potentially vary from each other. That is not the case in this particular example though. The two tails are very similar to each other, suggesting that the two tail fitting methods should provide decent benchmarks against each other.

TIME SERIES STABILITY TEST

I undertake two empirical tests with aim to identify which method best estimates the tails of the distribution. The time series stability test measures the stability of each tail fitting method. It is important for

users of RNDs to know whether changes in moment estimates are due to underlying fundamentals or random noise. For example, central bankers have suggested using RNDs to monitor investor sentiment and assess the effectiveness of monetary policy actions. Noisy moment estimates could lead to false positive conclusions and poor decision making.

First, time series of the moments generated by each of the methods will be graphically depicted and viewed for general stability. Then the noise component of each estimated moment's time series will be estimated in an GARCH(1,1) model. Since each method estimates the interior of the RND the same, any differences in the noise components should be due to the tail fitting method. The most stable tail fitting method should generate a noise component will the smallest mean squared error (MSE).

Next, normal probability plots of the noise components will be graphed and probability plot correlation coefficients (PPCCs) calculated. If the residuals of the GARCH(1,1) model are truly noise, then the normal probability plots should be straight lines with a slope of one. The PPCC measures how close each of the noise components comes to this line. A higher PPCC should indicate a more stable noise component.

Last, cumulative distribution functions (CDFs) of the noise components will be graphed and compared against a normal distribution. If the

residuals of the GARCH(1,1) model are truly noise, then they should follow a normal distribution. Kolmogorov-Smirnov (K-S) tests will measures how close the CDFs of the noise components come to normal distribution CDFs. A lower K-S test statistic indicates a more stable noise component.

PRICING ERROR TEST

As a compliment to the stability test, I undertake a pricing error test comparing estimated option prices to actual option data. For a given day's estimated RND, the corresponding option price estimates are compared to actual option prices in terms of the root mean squared error (RMSE) and the root mean squared relative error (RMSRE).

$$\text{RMSE} = \sqrt{\frac{1}{n}\sum(IV_i - \widehat{IV}_i)^2} \tag{33}$$

$$\text{RMSRE} = \sqrt{\frac{1}{n}\sum\left(\frac{IV_i - \widehat{IV}_i}{\widehat{IV}_i}\right)^2} \tag{34}$$

It is common in practice to measure pricing error in terms of Black-Scholes implied volatilities rather than call or put prices. Note, that on an economic basic, *vega* (the option-greek measuring the sensitivity of the option's price relative to the option's implied volatility) gets smaller as

the strike price moves away from the money (*very* small in the tails as defined in this thesis). Therefore, in terms of the tail-option's *price*, it is difficult to compare the performance of each tail-fitting method. The implied-volatilities much better estimate each tail-fitting method's estimate of the future uncertainty of the underlying asset's return.

In order to test if the tail estimates are biased above or below the actual option prices, mean error (ME) and mean relative error (MRE) are also calculated.

$$\text{ME} = \frac{1}{n}\sum(IV_i - \widehat{IV_i}) \tag{36}$$

$$\text{MRE} = \frac{1}{n}\sum\frac{(IV_i - \widehat{IV_i})}{\widehat{IV_i}} \tag{37}$$

Since the subject of the thesis is the fitting of the distribution's tails, the RND will be estimated without using the entire range of available strike prices. The RMSE, RMSRE, ME, and MRE will be calculated over the outer range of strikes -- the strikes not used to estimate the RND. This will identify which estimation method best prices options outside of the range of observable strikes (i.e. in the tails of the distribution).

CHAPTER IV

FINDINGS

TIME SERIES STABILITY TEST

Figures (15-18) show time series of the implied volatility, implied skewness, and implied kurtosis estimated using the four separate tail fitting methods. At first glance, all four methods appear to generate stable estimates of implied volatility. However the GEV tail method produces time periods of highly unstable estimates of implied skewness and implied kurtosis. The truncated tails method also produces a few periods of unstable implied skewness and implied kurtosis, but appears to produce stable estimates most of the time.

To quantify the stability of each tail estimate, Tables (7-9) give parameter estimates of a GARCH(1,1) model for the implied volatility, implied skewness, and implied kurtosis times series, respectively. The MSE of the residuals is by far the highest using the GEV tails method for all

three implied moments. Furthermore, the r-squared is by far the lowest using the GEV method for all three implied moments. In terms of MSEs, the naïve methods (truncated tails and lognormal tails) produce the most stable implied moments, especially for higher order moments. However, this is largely due to the fact that they consistently underestimate the higher order moments. Using *r*-squared as a criteria instead, the smile-extrapolated tails method outperforms the naïve methods, especially for higher order moments.

Figures (19-21) give normal probability plots for the residuals of each GARCH(1,1) model. If the residuals truly represent random noise, as GARCH models assume, then the plots should form a straight line with a slope of one. The probability plot correlation coefficient (PPCC) measures the correlation of each probability plot with the ideal straight line. The GEV method generates by far the lowest PPCC for each implied moment. It can be seen from the plots that the GEV method produces a residual distribution far too heavy-tailed compared to a normal distribution. However, the heavy tails do not appear to be driven by outliers. There is some non-random factor causing the GEV method to estimate widely fluctuating implied moments. Surprisingly, for implied kurtosis, the naïve methods produce residuals closest to random noise.

Figures (22-24) give the CDF for the residuals of each GARCH(1,1) model. The Kolmogorov-Smirnov (K-S) test statistic compares each CDF to its ideal normal distribution CDF. The results are similar to those from the normal probability plots.

Note that GARCH models with lags greater than one were also estimated, but the results were comparable.

PRICING ERROR TEST

Table (10) gives the general results for the pricing error test. The baseline is set by the truncated tails method, with a ME of -0.3974 (a MRE of -100%) and a RMSE of 0.4179 (a RMSRE of -100%). Note that all error estimates use the estimated implied volatility versus the implied volatility observed from actual option prices. It is not surprising that estimates have a very large negative bias, as the truncated tail method assumes all tail options have a price (and implied volatility) of zero. All other methods greatly improve upon these results.

The normal tail method results in a ME of -0.1498 (ln[-0.1498/-0.3974] = a 98% improvement over the baseline) and a RMSE of 0.1721. The relative errors are a MRE of -31.27% and a RMSRE of 33.76%. Despite the massive improvement over the baseline, the normal tail method is

still a bit of a straw-man, as it is well-known that implied volatilities continue to rise as they move away from the mean of the distribution. -31.27% is most likely an unacceptably negative bias for any user of option-implied probability distributions.

The robust model-based method, the GEV tails method, results in a ME of -0.0152 (-4.54% MRE) and a RMSE of 0.0326 (7.81% RMSRE). This is a (*ln*[-0.0152/-0.3974]) 326% improvement over the baseline and a (*ln*[-0.0152/-0.1498]) 229% improvement over the normal tails method.

The robust non-parametric method, the spline-extrapolated method, gives the best pricing error results: a ME of -0.0042 (only -1.57%! MRE) and a RMSE of 0.0134 (4.42% RMSRE). This is a (*ln*[-0.0042/-0.3974]) 455% improvement over the baseline, a (*ln*[-0.0042/-0.1498]) 357% improvement over the normal tails method, and a (*ln*[-0.0042/-0.0152]) 129% improvement over the GEV method.

Tables (11-12) give results for just the lower and upper tail respectively. Generally, relative results are the same between methods as in Table (10). It is difficult to summarize general results between the lower and upper tail estimates. ME and RMSE are generally better for the upper

tail, but MRE and RMSRE are generally better for the lower tail. This is understandable, as option-implied probability distributions are typically negatively skewed, meaning higher implied volatilities (and higher ME) in the lower tail, and there is typically option strikes that go much farther below-the-money than above-the-money (again, implying higher ME).

CHAPTER V

CONCLUSION

This thesis builds on and contributes to work in the field of financial risk management, specifically option-implied probability distributions. Although a number of studies have examined estimating the middle portion of probability distributions, there has not been a strong focus on the tails of the distribution, which are of particular importance in a risk management setting. As such, this study provides additional insights about these tails, by horseracing four different tail-fitting methods. This research differs from previous studies by introducing a new, non-parametric, heuristic tail-fitting method that is similar in methodology to the consensus, most-often used method to estimate the middle portion of the probability distribution; and, by identifying which tail fitting method produces the most stable estimate with the least tail-option pricing error.

Four tail-fitting methods are examined. First (1), a base-case is laid down assuming truncated tails. The middle of the distribution is

estimated using the popular curve-fitting method, where a fourth-degree spline with a single knot at-the-money is fit to the implied-volatility smile. This same method is used to estimate the middle portion of the distribution for all four tail-fitting methods. In the truncated tails case, the distribution is then truncated at the 2nd and 98th percentiles, with the middle portion reweighted to sum to 100%. This method produces the maximum possible tail-option pricing error, as it assumes all options below the 2nd percentile and above the 98th percentile are worthless, which they never are in practice. This produces a mean error of -0.3974 (measured as the difference between the estimated implied volatility and the actual implied volatility from observed option prices) and a root mean squared error of 0.4179. All other methods will improve beyond these baselines. Despite its poor pricing ability, the truncated tails method does produce one of the most *stable* distribution estimates (as measured by a GARCH model of the time-series of the distribution's moments).

The second (2) tail-fitting method assumes normally distributed tails beyond the 2nd and 98th percentiles. This method is also a bit of a straw-man, as it is well-known that implied volatilities continue to rise as the strike price gets further away from the mean of the distribution. However, in practice, normally distributed tails are one of the most commonly used methods by researchers. This methodology produces

distributions on an order equal in *stability* to the truncated tails method, but also produces poor pricing error results. The mean error is -0.1489 (a mean relative error of -31.27%) and the root mean squared error is 0.1721 (a root mean squared relative error of 33.76%). These are massive improvements over the truncated tails method, but still relatively poor compared to the final two methods.

The third and fourth methods are not at all straw-men arguments and are designed to generate the best possible model-based estimate of the tails, and the best possible non-parametric, heuristic estimate of the tails. The third (3) method assumes generalized extreme value (GEV) distributed tails beyond the 2nd and 98th percentiles. Generalized Pareto distributed tails were also examined and the results were similar to the GEV tails. This method greatly improves pricing error. It produced a mean error of -0.0152 (a mean relative error of -4.54%) and a root mean squared error of 0.0326 (a root mean squared relative error of 7.81%). However, the estimated distributions are much less *stable* than the two naïve methods, on an order of magnitude of three to four times less stable. This means, that while the average estimate of the distribution is closer to the actual, observed option-implied distribution, the estimation varies widely from day-to-day, giving users of the estimated distribution little confidence in its immediate accuracy.

The fourth (4) and last method uses an implied-volatility curve-fitting method to fit the tails beyond the 2nd and 98th percentiles. This method assumes no pre-determined parameterization of the tails, and is similar to the fast and stable curve-fitting methodology used to fit the middle portion of the distribution. Instead of interpolating between observable strike prices, the implied volatility smile is used to extrapolate beyond observable strike prices, into the tails of the distribution. It produces the best pricing error results: a mean error of -0.0042 (a mean relative error of only -1.57%) and a root mean squared error of 0.0134 (a root mean squared relative error of 4.52%). Also, this method is on an order of magnitude equal in *stability* to the two naïve methods.

In short, the non-parameterized, heuristic method, similar to the fast and stable method most commonly used to estimate the *middle* portion of the probability distribution, is also *stable*, with the least option pricing bias in the *tails* of the distribution. Researchers and other users of option-implied probability distributions would be wise to use the spline-extrapolation method used in this thesis, or other non-parameterized, heuristic methods.

This thesis focused on determining whether a model-based or a non-parameterized method is best at estimating the tails of the option-implied probability distribution. Results confidently favor the latter. Future research might focus on determining the best possible non-parameterized method, or use a non-parameterized method to continue research along the three lines of applications outlined in Table (1).

Option-implied probability distributions are still a burgeoning area of research and application in financial risk management. I am confident of their utility and continued growth in use, and rest assured this thesis accomplished its goal. Thank you for your interest and time. God bless!

REFERENCES

Abken, P., D. Madan, and S. Ramamurtie. (1996a). "Estimation of Risk-Neutral and Statistical Densities by Hermite Polynomial Approximation: With an Application to Eurodollar Futures Options." Working paper, Federal Reserve Bank of Atlanta.

Abken, P., D. Madan, and S. Ramamurtie. (1996b). "Pricing S&P 500 Index Options Using a Hilbert Space Basis." Working paper, Federal Reserve Bank of Atlanta.

Aït-Sahalia, Y., and A. Lo. (1998). "Nonparametric Estimation of State-Price Densities Implicit in Financial Asset Prices." *Journal of Finance*, vol. 53, no. 2 (April):499–547.

Aït-Sahalia, Y., and A. Lo. (2000). "Nonparametric Risk Management and Implied Risk Aversion." *Journal of Econometrics*, vol. 94, nos. 1–2 (January/February): 9–51.

Anagnou, I., M. Bedendo, S. Hodges, and R. Tompkins. (2002). "The Relation between Implied and Realized Probability Density Functions." Working paper, University of Warwick.

Andersen, A., and T. Wagener. (2002). "Extracting Risk Neutral Probability Densities by Fitting Implied Volatility Smiles: Some Methodological Points and an Application to the 3M EURIBOR Futures Option Prices." Working paper, European Central Bank.

Aparicio, S., and S. Hodges. (1998). "Implied Risk-Neutral Distribution: A Comparison of Estimation Methods." Working paper, Warwick University.

Arrow, K. (1964). "The Role of Securities in the Optimal Allocation of Risk-bearing." *Review of Economic Studies*, vol. 31, no. 2, 91-96.

Bahra, B. (1997). "Implied Risk-Neutral Probability Density Functions from Option Prices: Theory and Application." Working paper, Bank of England.

Bakshi, G., C. Cao, and Z. Chen. (1997). "Empirical Performance of Alternative Option Pricing Models." *Journal of Finance*, vol. 52, no. 5 (December):2003–49.

Banz, R. and M. Miller. (1978). "Prices for State-Contingent Claims: Some Estimates and Applications," *Journal of Business,* vol. 51: 653-672.

Bates, D. (1991). "The Crash of '87: Was It Expected? The Evidence from Options Markets." *Journal of Finance*, vol. 46, no. 3 (July):1009–44.

Bates, D. (1996). "Jumps and Stochastic Volatility: Exchange Rate Processes Implicit in Deutsche Mark Options." *Review of Financial Studies*, vol. 9, no. 1 (Spring):69–107.

Bates, D. (2000). "Post-'87 Crash Fears in S&P 500 Futures Options." *Journal of Econometrics*, vol. 94, nos. 1–2 (January/February):181–238.

Bernoulli, D. (1738). "Exposition of a New Theory on the Measurement of Risk" [translated by Sommer, L. (1954)]. *Econometrica*, vol. 22, no. 1 (January): 22–36.

Black, F., and M. Scholes. (1973). "Pricing of Options and Corporate Liabilities." *Journal of Political Economy*, vol. 81, no. 3 (May/June):637–659.

Bliss, R., and N. Panigirtzoglou. (2001). “Testing the Stability of Implied Probability Density Functions.” *Journal of Banking and Finance*, vol. 26, no. 2–3 (March):381–422.

Bliss, R., and N. Panigirtzoglou. (2002). “Option-Implied Risk Aversion Estimates: Robustness and Patterns.” Previously titled: “Recovering Risk Aversion from Options.” Working paper, Federal Reserve Bank of Chicago.

Bollen, N., and E. Rasiel. (2002). “The Performance of Alternative Valuation Models in the OTC Currency Options Market.” Working paper, Vanderbilt University.

Bondarenko, O. (2000). “Recovering Risk-Neutral Densities: A New Nonparametric Approach.” Working paper, University of Illinois.

Boyle, P. (1986). "Option Valuation Using a Three-Jump Process", *International Options Journal,* vol. 3: 7-12.

Branger, N. (2002). “Pricing Derivative Securities using Cross-Entropy: An Economic Analysis.” Working paper, School of Business and Economics, Goethe University.

Breeden, D. and R. Litzenberger (1978). "Prices of State-Contingent Claims Implicit in Option Prices." *Journal of Business*, vol. 51: 621-652.

Brenner, M., and Y. Eom. (1997). "No-Arbitrage Option Pricing: New Evidence on the Validity of the Martingale Property." Working paper, New York University.

Brown, D., and J. Jackwerth. (2000). "The Information Content of the Volatility Smile." Working paper, University of Wisconsin at Madison.

Brown, G., and K. Toft. (1999). "Constructing Binomial Trees from Multiple Implied Probability Distributions." *Journal of Derivatives*, vol. 7, no. 2 (Winter): 83–100.

Buchen, P., and M. Kelly. (1996). "The Maximum Entropy Distribution of an Asset Inferred from Option Prices." *Journal of Financial and Quantitative Analysis*, vol. 31, no. 1 (March): 143–159.

Campa, J., K. Chang, and R. Reider. (1998). "Implied Exchange Rate Distributions: Evidence from OTC Option Markets." *Journal of International Money and Finance*, vol. 17, no. 1 (February):117–160.

Campa, J., K. Chang, and J. Refalo. (1999). “An Options-Based Analysis of Emerging Market Exchange Rate Expectations: Brazil’s Real Plan, 1994–1999.” Working paper, New York University.

CBOE (2015). "The CBOE Volatility Index – VIX,” [CBOE white paper]. Retrieved from www.cboe.com/micro/vix/vixwhite.pdf.

Clark, G., and Rock, A. (2016). "Processes Contributing to the Maintenance of Flying Phobia: A Narrative Review". *Frontiers in Psychology*, vol. 7: 754.

Cochrane, J. (2005). *Asset Pricing*. Princeton University Press; Revised edition.

Corrado, C. (2001). “Option Pricing Based on the Generalized Lambda Distribution.” *Journal of Futures Markets*, vol. 21, no. 3 (March):213–236.

Corrado, C., and T. Su. (1996). “Skewness and Kurtosis in S&P 500 Index Returns Implied by Option Prices.” *Journal of Financial Research*, vol. 19, no. 2 (Summer): 175–192.

Corrado, C., and T. Su. (1997). “Implied Volatility Skews and Stock Index Skewness and Kurtosis Implied by S&P 500 Index Option Prices.” *Journal of Derivatives*, vol. 4, no. 4 (Summer): 8–19.

Coutant, S. (2000). “Time-Varying Implied Risk Aversion in Options Prices Using Hermite Polynomials.” Working paper, Banque de France.

Coutant, S., E. Jondeau, and M. Rockinger. (2001). “Reading PIBOR Futures Options’ Smiles: The 1997 Snap Election.” *Journal of Banking and Finance*, vol. 25, no. 11 (November): 1957–87.

Cox, J., J. Ingersoll, and S. Ross (1985). "A Theory of the Term Structure of Interest Rates". *Econometrica*. 53: 385–407.

Cox, J., S. Ross, and M. Rubinstein. (1979). “Option Pricing: A Simplified Approach.” *Journal of Financial Economics*, vol. 7, no. 3 (September):229-263.

De Jong, C., and R. Huisman. (2000). “From Skews to a Skewed-*t*: Modelling Option-Implied Returns by a Skewed Student-*t*.” Working paper, Erasmus University.

Fama, E., French, K. (1993). "Common risk factors in the returns on stocks and bonds." *Journal of Financial Economics*, vol. 33: 3–56.

Figlewski, S. (2008). “Estimating the Implied Risk Neutral Density for the U.S. Market Portfolio”. *VOLATILITY AND TIME SERIES ECONOMETRICS: ESSAYS IN HONOR OF ROBERT F. ENGLE*, Tim Bollerslev, Jeffrey R. Russell and Mark Watson, eds., Oxford, UK: Oxford University Press, 2008.

Gemmill, G., and A. Saflekos. (2000). “How Useful Are Implied Distributions? Evidence from Stock-Index Options.” *Journal of Derivatives*, vol. 7, no. 3 (Spring): 83–98.

Giamouridis, D., and M. Tamvakis. (2001). “Can We Rely on Implied Distributions? Evidence from the Markets.” Working paper, City University Business School.

Giamouridis, D., and M. Tamvakis. (2002). “Asymptotic Distribution Expansions in Option Pricing.” *Journal of Derivatives*, vol. 9, no. 4 (Summer):33–44.

Härdle, W., and A. Yatchew. (2002). “Dynamic Nonparametric State Price Density Estimation Using Constrained Least Squares and the Bootstrap.” Working paper, Humboldt University Berlin.

Hartvig, N., J. Jensen, and J. Pedersen. (1999). “Risk Neutral Densities of the ‘Christmas Tree’ Type.” Working paper, Centre for Mathematical Physics and Stochastics, Aarhus University.

Hayes, S., and H. Shin. (2002). “Liquidity and Risk Appetite: Evidence from Equity Index Option Prices.” Working paper, London School of Economics.

Heston, S. (1993). “A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options.” *Review of Financial Studies*, vol. 6, no. 2 (Summer): 327–343.

Jackwerth, J. (2000). “Recovering Risk Aversion from Option Prices and Realized Returns.” *Review of Financial Studies*, vol. 13, no. 2 (Summer): 433–451.

Jarrow, R., and A. Rudd. (1982). “Approximate Valuation for Arbitrary Stochastic Processes.” *Journal of Financial Economics*, vol. 10, no. 3 (November): 347–369.

Jarrow, R., H. Li, and F. Zhao. (2003). “Interest Rate Caps ‘Smile’ Too! But Can the LIBOR Market Models Capture It?” Working paper, Cornell University.

Jiang, G., and Y. Tian. (2005). “The Model-Free Implied Volatility and Its Information Content,” *The Review of Financial Studies*, vol. 18, no. 4 (Winter): 1305–1342

Jondeau, E., and M. Rockinger. (2001). “Gram-Charlier Densities.” *Journal of Economic Dynamics and Control*, vol. 25, no. 10 (October): 1457–83.

Leahy, M., and C. Thomas. (1996). “The Sovereignty Option: The Quebec Referendum and Market Views on the Canadian Dollar.” Working paper, Federal Reserve Board, Washington, DC.

Longstaff, F. (1995). “Option Pricing and the Martingale Restriction.” *Review of Financial Studies*, vol. 8, no. 4 (Winter):1091–1124.

Malz, A. (1996). “Using Option Prices to Estimate Realignment Probabilities in the European Monetary System: The Case of Sterling–

Mark." *Journal of International Money and Finance*, vol. 15, no. 5 (October):717–748.

Malz, A. (1997). "Estimating the Probability Distribution of the Future Exchange Rate from Option Prices." *Journal of Derivatives*, vol. 5, no. 2 (Winter):18–36.

Manaster, S., and G. Koehler. (1982). "The Calculation of Implied Variances from the Black–Scholes Model: A Note." *Journal of Finance*, vol. 37, no. 1 (March):227–230.

Mayhew, S. (1995). "On Estimating the Risk-Neutral Probability Distribution Implied by Option Prices." Working paper, Purdue University.

McCauley, R., and W. Melick. (1996). "Risk Reversal Risk." *Risk*, vol. 9, no. 11 (November): 54–57.

McManus, D. and D. Watt. (1999). "Estimating One-Factor Models of Short-Term Interest Rates." Bank of Canada Working Paper. 99-18.

Melick, W., and C. Thomas. (1997). "Recovering an Asset's Implied PDF from Option Prices: An Application to Crude Oil during the Gulf Crisis."

Journal of Financial and Quantitative Analysis, vol. 32, no. 1 (March): 91–115.

Merton, R. (1973). “Theory of Rational Option Pricing.” *Bell Journal of Economics and Management Science*, vol. 4, no. 1 (Spring): 141–183.

Merton, R. (1974). “On the Pricing of Corporate Debt: The Risk Structure of Interest Rates.” *The Journal of Finance*, vol. 29, no. 2 (May): 449-470.

Merton, R. (1976). "Option pricing when underlying stock returns are discontinuous". *Journal of Financial Economics*. 3: 125–144.

Nakamura, H., and S. Shigenori. (1999). "Extracting Market Expectations from Option Prices: Case Studies in Japanese Option Markets," *Monetary and Economic Studies*, Institute for Monetary and Economic Studies, Bank of Japan, vol. 17(1), pages 1-43, May.

Neuhaus, H. (1995). “The Information Content of Derivatives for Monetary Policy: Implied Volatilities and Probabilities”. Deutsche Bundesbank Economic Research Group, discussion paper No. 3/95.

Perignon, C., and C. Villa. (2002). “Extracting Information from Options Markets: Smiles, State-Price Densities and Risk Aversion.” *European Financial Management*, vol. 8, no. 4 (December): 495–513.

Posner, S. and M. Milevsky. (1998). “Valuing Exotic Options by Approximating the SPD with Higher Moments.” *Journal of Financial Engineering*, vol. 7, no. 2
(June):109–125.

Potters, M., R. Cont, and J. Bouchaud. (1998). “Financial Markets as Adaptive Systems.” *Europhysics Letters*, vol. 41, no. 3 (February): 239–244.

Pritsker, M. (1998). “Nonparametric Density Estimation and Tests of Continuous Time Interest Rate Models.” *Review of Financial Studies*, vol. 11, no. 3 (Autumn): 449–487.

Ritchey, R. (1990). “Call Option Valuation for Discrete Normal Mixtures.” *Journal of Financial Research*, vol. 13, no. 4 (Winter): 285–295.

Rockinger, M., and E. Jondeau. (2002). “Entropy Densities with an Application to Autoregressive Conditional Skewness and Kurtosis.” *Journal of Econometrics*, vol. 106, no. 1 (January):119–142.

Rookley, C. (1997). “Fully Exploiting the Information Content of Intra Day Option Quotes: Applications in Option Pricing and Risk Management.” Working paper, University of Arizona.

Rosenberg, J. (1998). “Pricing Multivariate Contingent Claims using Estimated Risk-Neutral Density Functions.” *Journal of International Money and Finance*, vol. 17, no. 2 (April):229–247.

Rosenberg, J. (2003). “Nonparametric Pricing of Multivariate Contingent Claims.” *Journal of Derivatives*, vol. 10, no. 3 (Spring):9–26.

Rosenberg, J., and R. Engle. (2002). “Empirical Pricing Kernels.” *Journal of Financial Economics*, vol. 64, no. 3 (June): 341–372.

Ross, S. (1976). “Options and Efficiency.” *Quarterly Journal of Economics*, vol. 90: 75–89.

Rubinstein, M. (1985). “Nonparametric Tests of Alternative Option Pricing Models Using All Reported Trades and Quotes on the 30 Most Active CBOE Option Classes from August 23, 1976 through August 31, 1978.” *Journal of Finance*, vol. 40, no. 2 (June):455–480.

Rubinstein, M. (1994). “Implied Binomial Trees.” *Journal of Finance*, vol. 49, no. 3 (July): 771–818.

Rubinstein, M. (1998). “Edgeworth Binomial Trees.” *Journal of Derivatives*, vol. 5, no. 3 (Spring):20–27.

Sherrick, B., P. Garcia, and V. Tirupattur. (1995). “Recovering Probabilistic Information from Option Markets: Tests of Distributional Assumptions.” Working paper, University of Illinois at Urbana-Champaign.

Sherrick, B., S. Irwin, and D. Forster. (1992). “Option-Based Evidence of the Nonstationarity of Expected S&P 500 Futures Price Distributions.” *Journal of Futures Markets*, vol. 12, no. 3 (March):275–290.

Sherrick, B., S. Irwin, and D. Forster. (1996). “An Examination of Option-Implied S&P 500 Futures Price Distributions.” *Financial Review*, vol. 31, no. 3 (August):667–695.

Shimko, D. (1993). “Bounds on Probability.” *Risk*, vol. 6, no. 4 (April): 33–37.

Shiratsuka, S. (2001). “Information Content of Implied Probability Distributions: Empirical Studies of Japanese Stock Price Index Options,” *Monetary and Economic Studies*, vol. 19, no. 3, p. 143-170.

Söderlind, P. (2000). “Market Expectations in the UK before and after the ERM Crisis.” *Economica*, vol 67, no. 265 (February):1–18.

Sprenkle, C. (1964). “Warrant Prices as Indicators of Expectations and Preferences.” In *The Random Character of Stock Market Prices*. Edited by Paul H. Cootner. Cambridge, MA: MIT Press.

Stutzer, M. (1996). “A Simple Nonparametric Approach to Derivative Security Valuation.” *Journal of Finance*, vol. 51, no. 5 (December):1633–52.

Toft, K., and B. Prucyk. (1997). “Options on Leveraged Equity: Theory and Empirical Tests.” *Journal of Finance*, vol. 52, no. 3 (July):1151–80.

Vasicek, O. (1977). "An equilibrium characterization of the term structure". *Journal of Financial Economics*. 5 (2): 177–188.

Vitale, T. (2012). *Cancel Crash* [documentary]. United States: tastytrade Network.

Weinberg, S. (2001). “Interpreting the Volatility Smile: An Examination of the Information Content of Option Price.” Discussion Paper No. 706, International Finance Discussion Papers.

Young, S., G. Jabbour, M. Kramin, and T. Kramin. (2001). “Option Pricing and Higher Order Moments.” Working paper, Merrill Lynch Investor Strategies Group.

APPENDICES

Appendix 1: Obtaining the Forward Price from Options Prices

Throughout the RND estimation process it is necessary to convert back and forth between option prices and their equivalent Black-Scholes IVs. The traditional Black-Scholes model requires as input the underlying asset's price and dividend yield. Optionmetrics provides estimates of these but they are problematic. One solution is to use equation (11), which takes as input the forward price in place of the underlying price and dividend yield. Fortunately, given a full set of option prices, the forward price can easily be found. It is the strike price where the put price equals the call price.

However, strikes prices come in discrete intervals and the exact strike which equates the put and call price is likely to be between two observable strikes. In order to pinpoint the forward price, option prices need to be interpolated between the strikes.

The process is threefold: first fit a spline to the IV smile derived only from put prices; then fit a spline to the IV smile derived only from call options; and last, convert both splines back into option prices and find where the put price function intersects with the call price function. One adjustment must be made to spline fitting process: when converting between option prices and IVs via equation (11), the forward price cannot be used as an input because it has not been estimated yet. To overcome this, the underlying asset price, as quoted by Optionmetrics is used in place of the forward price. Since the IV smile is only used as a computational tool, it does not matter what forward value is used to calculate IVs, so long as the same value is used to convert back to option prices. The quoted underlying price is close enough to the forward price that it can be used as a proxy.

Figures 15 and 16 show this process for the set of options prices used in the illustrative example from the thesis. Figure 15 gives the fitted splines for puts and calls separately. Note that the IV smiles in Figure 15 do not match the IV smile in Figure 1. This is because an incorrect forward price was used in Figure 15, which would be a problem if one spline was being fit to both puts and calls (as in Figure 2), but is of no consequence when fitting only puts or calls separately. Figure 16 shows the splines

converted back into option price space and the intersection of the two splines, which is the estimated forward price.

Appendix 2: Extracting Moments from the Estimated Distribution

The implied moments (*IM*) centered about the forward price, as displayed in Figure 8 are determined by the following equations

(17) $$IM_2 = \sqrt{\sum_K [f(K) * 0.01 * (K - F)^2]}$$

(18) $$IM_3 = \sum_K [f(K) * 0.01 * (K - F)^3] \Big/ IM_2{}^3$$

(19) $$IM_4 = \sum_K [f(K) * 0.01 * (K - F)^4] \Big/ IM_2{}^4$$

where IM_2 is implied volatility un-scaled by the forward price, IM_3 is implied skewness, IM_4 is implied kurtosis, F is the option-implied forward price, and $f(K)$ is the RND at strike K. Note that implied volatility $= IM_2/F$. As explained in section 2.4.4, the RND is calculated at one cent intervals, and therefore (17)-(19) should ideally be summed at one cent intervals over the entire range of strike prices. But what constitutes the entire range of strikes? The moments could be summed from a zero strike price up to some assuredly large strike price such as two times the forward price. But this is computationally infeasible.

If the range of strikes used to calculate (17)-(19) is not wide enough, the *IM*s may not converge to a stable estimate, and the *IM*s will fluctuate widely from one day to the next. To illustrate this see Figure 17. It depicts each one cent strike's contribution to the summation in equation (19). As the strike price goes to zero or infinity, the graph tends toward zero, because the RND tends to zero. However, if too narrow a range of strikes is used to calculate equation (19), the graph will not tend to zero, and the implied kurtosis estimate will not converge. In the case of the day chosen for Figure 17, the low end of the range approaches zero at a strike price around 350, which means the left tail of the distribution must be extrapolated to this strike.

In this thesis I widen the range of strikes until $f(K) * 0.01 * (K - F)^4$ is less than 5% of its maximum value. This insures that implied kurtosis converges to a stable estimate. In the case of Figure 17, its highest value is 9,113.42, and therefore the range of strikes is widened until the graph drops below 455.67 both on the high and low side. Note that since the limit of $(K - F)^4$ tends toward infinity faster than $(K - F)^3$ tends toward +/-infinity, implied skewness always converges within a tighter range than implied kurtosis. By the same reasoning implied volatility converges within a tighter range than implied skewness. Therefore, by

insuring that implied kurtosis converges, this also insures that implied skewness and implied volatility converge to a stable estimate.

Table 1. Applications of Option-Implied *PDF*s

Panel A: Central Bank Research involving Option-Implied *PDF*s

Author (*Year*)	*Central Bank*
Neuhaus (1995)	Deutsche Bundesbank
Leahy and Thomas (1996)	Federal Reserve Board of Governors
Melick and Thomas (1997)	Federal Reserve Board of Governors
Bahra (1997)	Bank of England
Abken and Ramamurtie (1996)	Federal Reserve of Atlanta
Malz (1997)	Federal Reserve of New York
Nakamura and Shiratsuka (1999)	Bank of Japan
McManus and Watt (1999)	Bank of Canada
Coutant et al. (2001)	Banque de France
Andersen and Wagener (2002)	European Central Bank

Panel B: Event Studies using Option-Implied *PDF*s

Author (*Year*)	*Event*
Bates (1991, 2000)	Black Monday, Oct. 1987
Melick and Thomas (1997)	Crude oil prices, first Gulf War
Gemmill and Seflekos (1999)	FTSE 100, crashes and elections
Söderlind (2000)	UK Exchange Rate Mechanism crisis
Jondeau and Rockinger (2000)	French snap election, 1997
Shiratsuka (2001)	Japanese boom and bust, 80's/90's
Andersen and Wagener (2002)	Euribor rates, Sept. 11th, 2001
Campa, Chang and Refalo (2002)	Brazilian Real Plan, 1994-1999

Panel C: Implied Risk Aversion

Author (*Year*)	*Risk-Neutral / Historical Model*
Ait-Sahalia and Lo (1998, 2000)	Kernel methods for both
Coutant (2000)	Hermite polynomial /ARCH
Haynes and Shin (2002)	Cubic spline / GARCH
Perignon and Villa (2002)	Kernel methods for both
Rosenberg and Engle (2002)	Curve fitting / GARCH

Table 2. Generalized Distributions Used to Model Option-Implied *PDF*s

Author (*Year*)	*Distribution*
Sherrick et al. (1992, 1995, 1996):	Burr Type-XII
Aparicio and Hodges (1998):	Beta of the Second Kind
Posner and Milevsky (1998):	Tukey's Lambda
DeJong and Huisman (2000):	Skewed-Student-t
Corrado (2001):	Johnson's S_U

Table 3. Expansion Methods Used to Model Option-Implied *PDF*s

Author (*Year*)	*Expansion Method*
Jarrow and Rudd (1982):	Edgeworth
Longstaff (1995):	Edgeworth
Corrado and Su (1996, 1997):	Gram–Charlier
Abken et al. (1996a, 1996b):	Hermite Polynomial
Abadir and Rockinger (1997):	Confluent Hypergeometric
Brenner and Eom (1997):	Laguerre Polynomials
Potters et al. (1998):	Edgeworth
Jondeau and Rockinger (2001):	Hermite Poly./Gram–Charlier
Young et al. (2001):	Edgeworth
Giamouridis and Tamvakis (2001, 2002):	Edgeworth

Table 4. Non-Parametric Methods used to Estimate Option-Implied *PDF*s

Panel A: Maximum Entropy Methods

Author (*Year*)	*Prior Distribution*
Rubinstein (1994)	Lognormal Prior
Buchen and Kelly (1996)	Uniform and Lognormal Prior
Stutzer (1996)	Historical Distribution Prior
Branger (2002)	Multiple Distribution Priors
Rockinger and Jondeau (2002)	Normal & *t*-Distribution Priors

Panel B: Kernel Methods

Author (*Year*)	*Description*
Rookley (1997)	In strike ad time-to-expiration
Aït-Sahalia and Lo (1998)	In stock price, strike, maturity, interest rate, and dividends
Bondarenko (2000)	Convolution of kernel & std densities
Härdle and Yatchew (2002)	Non-parametric least squares

Panel C: Curve-Fitting Methods

Author (*Year*)	*Description*
Shimko (1993)	Quadratic polynomial
Mayhew (1995)	Cubic Spline
Aparicio and Hodges (1998)	Cubic B-Splines
Campa, Chang, and Reider (1998)	Cubic Splines
Rosenberg (1998, 2003)	Biv.-polynomials fitted to Log-IV
Brown and Toft (1999)	Seventh-Order Splines
Andersen and Wagener (2002)	High-Order Splines
Hayes and Shin (2002)	Cubic Splines
Rosenberg and Engle (2002)	Polynomials Fitted to Log-IV

Table 5: Summary Statistics

Option data from Januray 1, 2003 to December 29, 2017. Forward price derived from option price (see Appendix 1). CDF and spline fitting statistics based of a fourth degree spline fit to the implied volatility smile.

	Mean	Std. Dev.	Min.	Q1	Median	Q3	Max.
S&P 500 Index	1,185.93	178.73	676.53	1,085.36	1,193.86	1,305.11	1,565.15
Forward Price	1,186.76	181.29	673.57	1,083.21	1,193.91	1,307.32	1,576.67
Risk-Free Rate	2.18%	1.95%	0.16%	0.27%	1.34%	4.11%	6.09%
Days to Expiration	39.1	6.4	15.0	34.0	39.3	44.0	52.0
Number of Options Used							
Puts Used	66.1	44.7	8.0	27.0	48.0	113.0	165.0
Calls Used	54.0	28.3	10.0	27.0	53.0	76.0	164.0
Total Strikes Used	68.4	42.0	11.0	29.0	58.0	110.0	164.0
Minimum of CDF	0.0090	0.0189	0.0000	0.0018	0.0041	0.0085	0.3345
Maximum of CDF	0.9864	0.0333	0.4928	0.9850	0.9940	0.9998	1.0000
Spline Fitting							
Standard Error	0.0227	0.0171	0.0013	0.0107	0.0175	0.0295	0.1353
R-squared	0.9990	0.0013	0.9742	0.9988	0.9994	0.9997	1.0000

Table 6: Raw Option Price Data

Obtained from S&P Index options on Jan. 31, 2012 with an expiration date of Mar. 17, 2012 (time to expiration = 46 days). Underlying price = 1,312.41. Forward = 1,308.86. Implied volatilities (IV) are obtained from the Black-Scholes option pricing model using forward prices.

	Calls				Puts			
Strike	Ask	Bid	Mid	IV	Ask	Bid	Mid	IV
750					0.10	0.05	0.075	0.532
775					0.15	0.05	0.100	0.516
780					0.15	0.05	0.100	0.510
800					0.15	0.05	0.100	0.487
825					0.20	0.10	0.150	0.476
840					0.40	0.05	0.225	0.478
850					0.40	0.10	0.250	0.472
860					0.40	0.05	0.225	0.455
870					0.45	0.05	0.250	0.449
875					0.40	0.05	0.225	0.438
880					0.45	0.10	0.275	0.442
895					0.50	0.10	0.300	0.429
900					0.45	0.25	0.350	0.430
905					0.50	0.15	0.325	0.421
910					0.55	0.15	0.350	0.419
915					0.55	0.15	0.350	0.413
920					0.55	0.20	0.375	0.411
925					0.55	0.20	0.375	0.405
930					0.60	0.25	0.425	0.405
940					0.65	0.25	0.450	0.397
950					0.70	0.30	0.500	0.390
960					0.80	0.30	0.550	0.384
975					0.90	0.55	0.725	0.380
980					0.95	0.45	0.700	0.373
985					1.00	0.40	0.700	0.367
990					1.10	0.55	0.825	0.369
995					1.15	0.55	0.850	0.365
1000					1.00	0.65	0.825	0.358
1005					1.30	0.55	0.925	0.358
1010					1.35	0.60	0.975	0.355
1015					1.40	0.60	1.000	0.350
1020					1.50	0.85	1.175	0.353
1025					1.40	0.75	1.075	0.342
1030					1.60	1.00	1.300	0.346
1035					1.75	0.75	1.250	0.338
1040					1.70	1.00	1.350	0.336
1045					1.95	1.00	1.475	0.335
1050					1.55	1.10	1.325	0.323
1055					2.10	1.50	1.800	0.335
1060					2.20	1.40	1.800	0.328

1065					2.30	1.15	1.725	0.320
1070					2.40	1.25	1.825	0.317
1075					2.50	1.65	2.075	0.318
1080					2.60	1.75	2.175	0.315
1085					2.80	1.50	2.150	0.308
1090					2.85	1.60	2.225	0.303
1095					3.00	1.95	2.475	0.303
1100					2.70	1.85	2.275	0.292
1105					3.30	1.90	2.600	0.293
1110					3.40	2.05	2.725	0.290
1115					3.60	2.20	2.900	0.287
1120					3.70	2.65	3.175	0.286
1125					3.90	2.85	3.375	0.283
1130					4.00	3.00	3.500	0.279
1135					4.30	3.20	3.750	0.277
1140					4.50	3.10	3.800	0.271
1145					4.70	3.60	4.150	0.270
1150					4.50	3.80	4.150	0.263
1155					5.30	3.80	4.550	0.262
1160					5.40	4.00	4.700	0.257
1165					5.80	4.30	5.050	0.255
1170					6.10	4.60	5.350	0.252
1175					6.50	5.00	5.750	0.250
1180					6.40	5.60	6.000	0.246
1185					7.20	5.90	6.550	0.245
1190	127.20	124.10	125.650	0.236	7.60	6.10	6.850	0.240
1195	122.70	119.50	121.100	0.233	8.00	6.60	7.300	0.238
1200	118.10	115.00	116.550	0.230	8.60	7.00	7.800	0.235
1205	113.60	110.50	112.050	0.228	9.00	7.60	8.300	0.232
1210	109.10	106.00	107.550	0.225	9.50	8.10	8.800	0.229
1215	104.70	101.60	103.150	0.222	10.20	8.60	9.400	0.226
1220	99.90	97.20	98.550	0.218	10.80	9.20	10.000	0.223
1225	96.50	92.70	94.600	0.219	11.40	10.00	10.700	0.221
1230	91.60	88.60	90.100	0.214	12.20	10.70	11.450	0.218
1235	87.40	84.40	85.900	0.212	12.90	11.30	12.100	0.214
1240	83.20	80.20	81.700	0.209	13.70	12.10	12.900	0.211
1245	79.50	75.90	77.700	0.207	14.60	12.90	13.750	0.208
1250	74.90	72.10	73.500	0.204	15.60	13.90	14.750	0.206
1255	70.90	68.10	69.500	0.201	16.60	14.90	15.750	0.203
1260	66.90	64.20	65.550	0.199	17.70	15.90	16.800	0.200
1265	63.50	60.10	61.800	0.197	18.70	17.10	17.900	0.197
1270	59.70	56.50	58.100	0.194	20.00	18.40	19.200	0.195
1275	55.50	52.90	54.200	0.191	21.20	19.60	20.400	0.192
1280	52.30	49.10	50.700	0.189	22.60	21.00	21.800	0.189
1285	48.70	45.80	47.250	0.186	24.20	22.50	23.350	0.187
1290	45.20	42.20	43.700	0.183	25.60	24.00	24.800	0.183
1295	41.80	39.10	40.450	0.180	27.30	25.70	26.500	0.181
1300	38.60	35.90	37.250	0.178	29.20	27.00	28.100	0.177
1305	35.40	32.80	34.100	0.175	31.20	28.60	29.900	0.173
1310	31.60	30.00	30.800	0.170	33.30	30.60	31.950	0.171
1315	29.60	27.00	28.300	0.170	35.30	33.50	34.400	0.170

1320	26.30	24.40	25.350	0.166	37.60	35.70	36.650	0.167
1325	23.50	21.90	22.700	0.163	40.20	38.10	39.150	0.164
1330	21.10	19.30	20.200	0.160	42.80	40.70	41.750	0.162
1335	18.90	17.10	18.000	0.157	45.70	43.50	44.600	0.160
1340	16.80	15.10	15.950	0.155	48.40	46.20	47.300	0.157
1345	14.90	13.20	14.050	0.153	51.50	49.20	50.350	0.154
1350	12.90	11.40	12.150	0.150	54.70	52.40	53.550	0.152
1355	11.20	9.80	10.500	0.148				
1360	9.90	8.40	9.150	0.147				
1365	8.50	7.00	7.750	0.144				
1370	7.20	5.90	6.550	0.142				
1375	6.00	4.90	5.450	0.140				
1380	5.20	4.30	4.750	0.140				
1385	4.50	3.40	3.950	0.138				
1390	3.80	2.80	3.300	0.137				
1395	3.10	2.35	2.725	0.136				
1400	2.65	2.10	2.375	0.137				
1405	2.30	1.60	1.950	0.136				
1410	1.85	1.30	1.575	0.134				
1415	1.60	1.10	1.350	0.135				
1420	1.35	0.90	1.125	0.135				
1425	1.15	0.75	0.950	0.135				
1430	1.00	0.60	0.800	0.135				
1435	0.85	0.45	0.650	0.135				
1440	0.80	0.40	0.600	0.138				
1445	0.70	0.30	0.500	0.138				
1450	0.65	0.25	0.450	0.140				
1460	0.55	0.15	0.350	0.142				
1475	0.50	0.15	0.325	0.152				
1500	0.30	0.10	0.200	0.160				

Table 7: Time Series Stability Test, Implied Volatility

Table gives parameter and standard error estimates for an GARCH(1,1) regression of the following form:

$$\widehat{IV}_0 = \alpha_0 + \alpha_1 IV_{-1} + \varepsilon$$

$$\hat{\varepsilon}^2_{IV_0} = \beta_0 + \beta_1 \varepsilon^2_{-1} + \beta_2 \hat{\varepsilon}^2_{IV_{-1}} + z$$

where *IV* is the implied volatility as estimated using the four tail-fitting procedures. The risk-neutral density (*RND*) within the observable range of strike prices (between the tails) is estimated in all four cases using the spline-fitting method of Figlewski (2009). Numbers in parenthesis are standard errors. ***, **, and * represent two-sided tests that the parameters are different than zero at the 99%, 95%, and 90% confidence levels, respectively. *MSE* is mean squared error. *PPCC* is the probability plot correlation coefficient and *K-S* is the Kolmogorov-Smirnov test statistic, with stars representing 99%, 95%, and 90% confidence levels that the distribution of the standardized residuals is different than the normal.

	Truncated Tails	*Lognormal Tails*	*GEV Tails*	*Spline-Extrapolated Tails*
α_0	0.1897***	0.2568***	0.2007***	0.2475***
	(0.0254)	(0.0565)	(0.006548)	(0.0531)
α_1	-0.9957***	-0.9964***	-0.9854***	-0.9969***
	(0.002044)	(0.002023)	(0.002836)	(0.001847)
β_0	$1.6374e^{-6}$***	$3.0728e^{-6}$***	0.0000100***	$2.2524e^{-6}$***
	($1.5395e^{-7}$)	($3.0583e^{-7}$)	($7.3482e^{-7}$)	($2.773e^{-7}$)
β_1	0.1229***	0.1308***	0.3068***	0.1307***
	(0.007384)	(0.00973)	(0.0132)	(0.009256)
β_2	0.8592***	0.8410***	0.7033***	0.8482***
	(0.007734)	(0.0102)	(0.008205)	(0.009739)
MSE	0.0001483	0.0001943	0.0005000	0.0001876
r^2	0.9743	0.9750	0.9333	0.9743
PPCC	0.9217***	0.9240***	0.8523***	0.9232***
K-S	0.1300***	0.1278***	0.1774***	0.1305***

Table 8: Time Series Stability Test, Implied Skewness

Table gives parameter and standard error estimates for an GARCH(1,1) regression of the following form:

$$\widehat{IS}_0 = \alpha_0 + \alpha_1 IS_{-1} + \varepsilon$$

$$\hat{\varepsilon}^2_{IS_0} = \beta_0 + \beta_1 \varepsilon^2_{-1} + \beta_2 \hat{\varepsilon}^2_{IS_{-1}} + z$$

where *IS* is the implied skewness as estimated using the four tail-fitting procedures. The risk-neutral density (*RND*) within the observable range of strike prices (between the tails) is estimated in all four cases using the spline-fitting method of Figlewski (2009). Numbers in parenthesis are standard errors. ***, **, and * represent two-sided tests that the parameters are different than zero at the 99%, 95%, and 90% confidence levels, respectively. *MSE* is mean squared error. *PPCC* is the probability plot correlation coefficient and *K-S Statistic* is the Kolmogorov-Smirnov test statistic, with stars representing 99%, 95%, and 90% confidence levels that the distribution of the standardized residuals is different than the normal.

	Truncated Tails	*Lognormal Tails*	*GEV Tails*	*Spline-Extrapolated Tails*
α_0	-0.6245***	-1.1027***	-1.2198***	-1.1084***
	(0.009467)	(0.0333)	(0.0611)	(0.0500)
α_1	-0.8497***	-0.9538***	-0.9387***	0.9692***
	(0.006639)	(0.00531)	(0.008041)	(0.004005)
β_0	0.00302 ***	0.000311***	0.002793***	0.000132***
	(0.0000148)	(0.0000556)	(0.000205)	(0.0000292)
β_1	0.1597***	0.1085***	0.5149***	0.0682***
	(0.006149)	(0.0116)	(0.0131)	(0.008270)
β_2	0.8577***	0.8427***	0.6687***	0.9113***
	(0.002034)	(0.0174)	(0.003354)	(0.0101)
MSE	0.02041	0.00657	0.46235	0.00759
r^2	0.7078	0.9084	0.3426	0.9445
PPCC	0.8410***	0.9912***	0.8539***	0.9924***
K-S	0.1835***	0.0429***	0.1876***	0.0465***

Table 9: Time Series Stability Test, Implied Kurtosis

Table gives parameter and standard error estimates for an GARCH(1,1) regression of the following form:

$$\widehat{IK}_0 = \alpha_0 + \alpha_1 IK_{-1} + \varepsilon$$

$$\hat{\varepsilon}^2_{IK_0} = \beta_0 + \beta_1 \varepsilon^2_{-1} + \beta_2 \hat{\varepsilon}^2_{IK_{-1}} + z$$

where *IK* is the implied kurtosis as estimated using the four tail-fitting procedures. The risk-neutral density (*RND*) within the observable range of strike prices (between the tails) is estimated in all four cases using the spline-fitting method of Figlewski (2009). Numbers in parenthesis are standard errors. ***, **, and * represent two-sided tests that the parameters are different than zero at the 99%, 95%, and 90% confidence levels, respectively. *MSE* is mean squared error. *PPCC* is the probability plot correlation coefficient and *K-S Statistic* is the Kolmogorov-Smirnov test statistic, with stars representing 99%, 95%, and 90% confidence levels that the distribution of the standardized residuals is different than the normal.

	Truncated Tails	*Lognormal Tails*	*GEV Tails*	*Spline-Extrapolated Tails*
α_0	3.1047***	4.3736***	4.2274***	5.0639***
	(0.0226)	(0.0957)	(0.3252)	(0.2562)
α_1	-0.9047***	-0.9486***	-0.9448***	-0.9692***
	(0.006817)	(0.006758)	(0.0161)	(0.003972)
β_0	0.002234***	0.001700***	0.0113***	0.001748***
	(0.000159)	(0.00315)	(0.001974)	(0.000339)
β_1	0.2267***	0.1399***	0.6231***	0.1097***
	(0.0138)	(0.0117)	(0.0143)	(0.009275)
β_2	0.7196***	0.8267***	0.7034***	0.8990***
	(0.0113)	(0.0112)	(0.002273)	(0.007449)
MSE	0.02867	0.06425	76.47863	0.53036
r^2	0.7820	0.8648	0.1907	0.9034
PPCC	0.9466***	0.9742***	0.8162***	0.9232***
K-S	0.1024***	0.0653***	0.2317***	0.1028***

Table 10: Pricing Error Test

Table gives pricing error statistics and 95% confidence intervals.

$$ME = \frac{1}{n}\sum(\widehat{IV}_i - IV_i) \qquad MRE = \frac{1}{n}\sum\left(\frac{\widehat{IV}_i - IV_i}{IV_i}\right) \qquad RMSE = \sqrt{\frac{1}{n}\sum(\widehat{IV}_i - IV_i)^2} \qquad RMSRE = \sqrt{\frac{1}{n}\sum(\frac{\widehat{IV}_i - IV_i}{IV_i})^2}$$

	ME	*MRE*	*RMSE*	*RMSRE*
Truncated Tails	-0.3974 [-0.4056, -0.3913]	-1.0000 [-1.0000, -1,0000]	0.4179 [0.4037, 0.4359]	1.0000 [1.0000, 1.0000]
Lognormal Tails	-0.1338 [-0.1376, -0.1264]	-0.2852 [-0.2936, -0.2769]	0.1601 [0.1532, 0.1668]	0.3182 [0.3053, 0.3318]
GEV Tails	-0.0152 [-0.0175, -0.0142]	-0.0454 [-0.0500, -0.0424]	0.03258 [0.0313, 0.0337]	0.0781 [0.0770, 0.0834]
Spline-Extrapolated Tails	-0.0042 [-0.0053, -0.0036]	-0.0157 [-0.0189, -0.0143]	0.0134 [0.0127, 0.1468]	0.0442 [0.0432, 0.0459]

Table 11: Pricing Error Test: Lower Tail Only

Table gives pricing error statistics and 95% confidence intervals.

$$ME = \frac{1}{n}\sum(\widehat{IV}_i - IV_i) \qquad MRE = \frac{1}{n}\sum\left(\frac{\widehat{IV}_i - IV_i}{IV_i}\right) \qquad RMSE = \sqrt{\frac{1}{n}\sum(\widehat{IV}_i - IV_i)^2} \qquad RMSRE = \sqrt{\frac{1}{n}\sum(\frac{\widehat{IV}_i - IV_i}{IV_i})^2}$$

	ME	*MRE*	*RMSE*	*RMSRE*
Truncated Tails	-0.4363 [-0.4429, -0.4317]	-1.0000 [-1.0000, -1,0000]	0.4452 [0.4296, 0.4669]	1.0000 [1.0000, 1.0000]
Lognormal Tails	-0.1498 [-0.1556, -0.1491]	-0.3127 [-0.3349, -0.3049]	0.1721 [0.1638, 0.1811]	0.3376 [0.3257, 0.3558]
GEV Tails	-0.0156 [-0.0175, -0.0138]	-0.0327 [-0.0369, -0.0301]	0.0329 [0.0323, 0.0352]	0.0646 [0.0623, 0.0679]
Spline-Extrapolated Tails	-0.0029 [-0.0036, -0.0024]	-0.0063 [-0.0084, -0.0054]	0.0122 [0.0133, 0.0124]	0.0263 [0.0258, 0.0280]

Table 12: Pricing Error Test: Upper Tail Only

Table gives pricing error statistics and 95% confidence intervals.

$$ME = \frac{1}{n}\sum(\widehat{IV}_i - IV_i) \qquad MRE = \frac{1}{n}\sum\left(\frac{\widehat{IV}_i - IV_i}{IV_i}\right) \qquad RMSE = \sqrt{\frac{1}{n}\sum(\widehat{IV}_i - IV_i)^2} \qquad RMSRE = \sqrt{\frac{1}{n}\sum(\frac{\widehat{IV}_i - IV_i}{IV_i})^2}$$

	ME	*MRE*	*RMSE*	*RMSRE*
Truncated Tails	-0.1664 [-0.1687, -0.1637]	-1.0000 [-1.0000, -1,0000]	0.1671 [0.1514, 0.1864]	1.0000 [1.0000, 1.0000]
Lognormal Tails	-0.0194 [-0.0217, -0.0170]	-0.1103 [-0.1255, -0.0981]	0.0248 [0.0216, 0.0272]	0.1348 [0.1213, 0.1515]
GEV Tails	-0.0209 [-0.0229, -0.0189]	-0.1203 [-0.1305, -0.1101]	0.0246 [0.0226, 0.0278]	0.1382 [0.1254, 0.1545]
Spline-Extrapolated Tails	-0.0143 [-0.0162, -0.0125]	-0.0803 [-0.0901, -0.0706]	0.0185 [0.0167, 0.0207]	0.1023 [0.0925, 0.1147]

Figure 1. Binomial Option Pricing Model

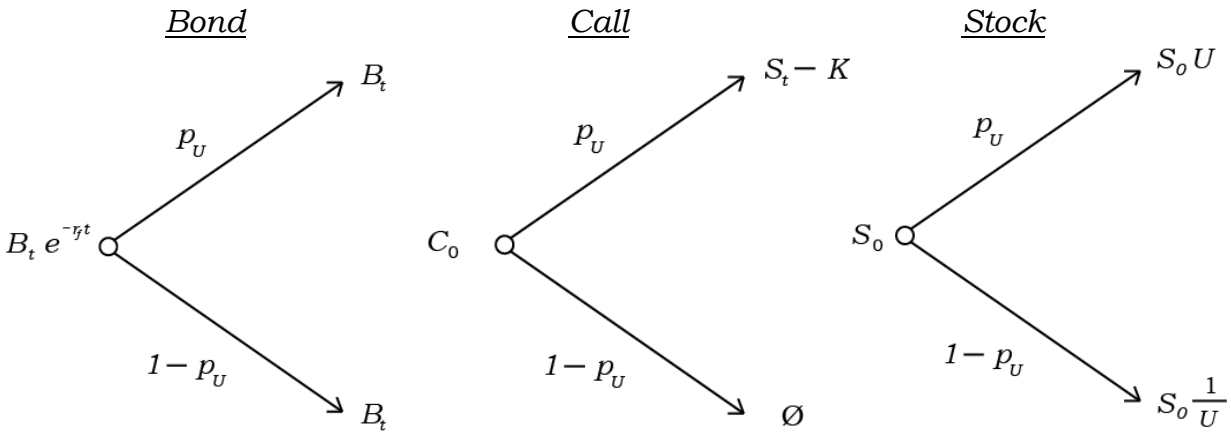

where,

B_t = price of a risk-free bond at time t
p_U = probability there is an *upward* return
S_t = underlying share price
K = strike price
C_0 = *call* price
$U = e^{rt+\sigma\sqrt{t}}$
$1/U = e^{rt-\sigma\sqrt{t}}$
r_f = risk-free rate.

Figure 2. Multi-Period Binomial versus Black-Scholes

Binomial Option Pricing Model *Black-Scholes*

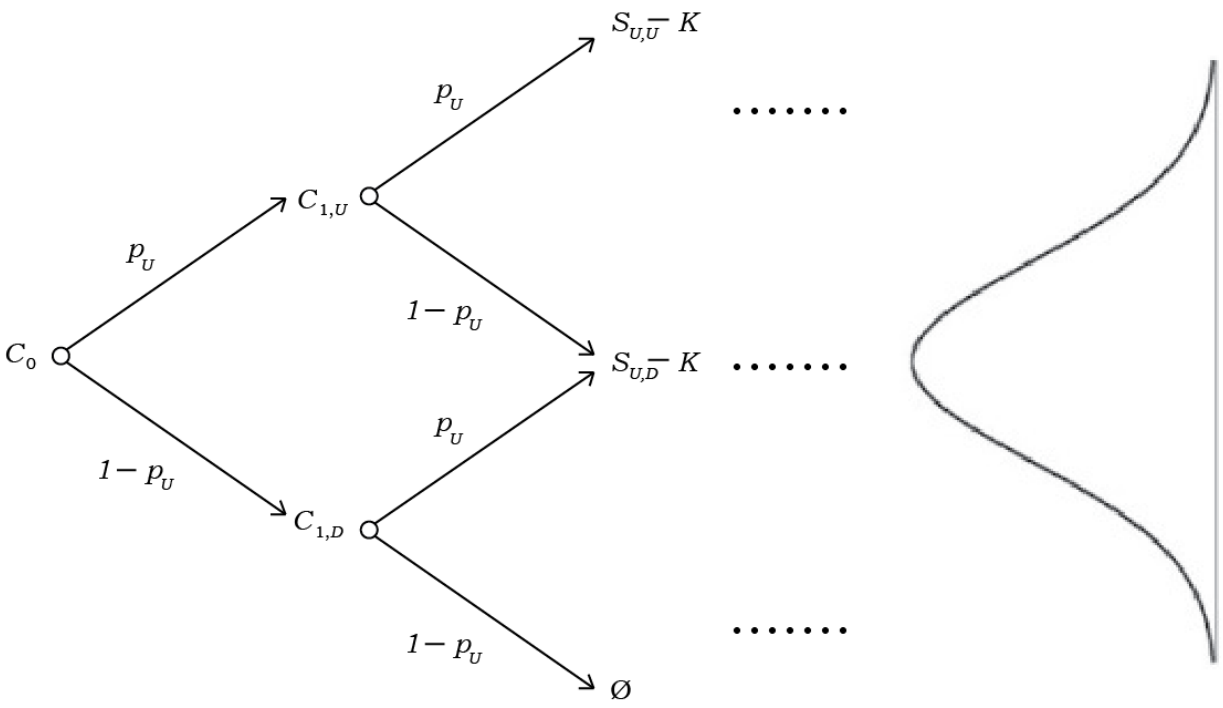

Figure 3. S&P 500 Options at the CBOE

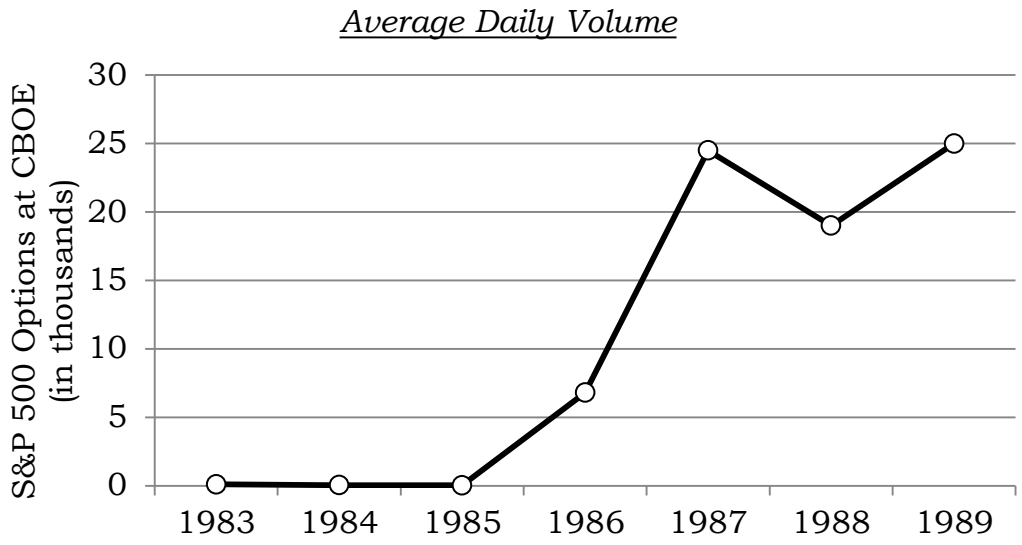

Source: *CBOE*

Figure 4. *Z*-Score of Daily S&P 100 Returns around Black Monday, 1987

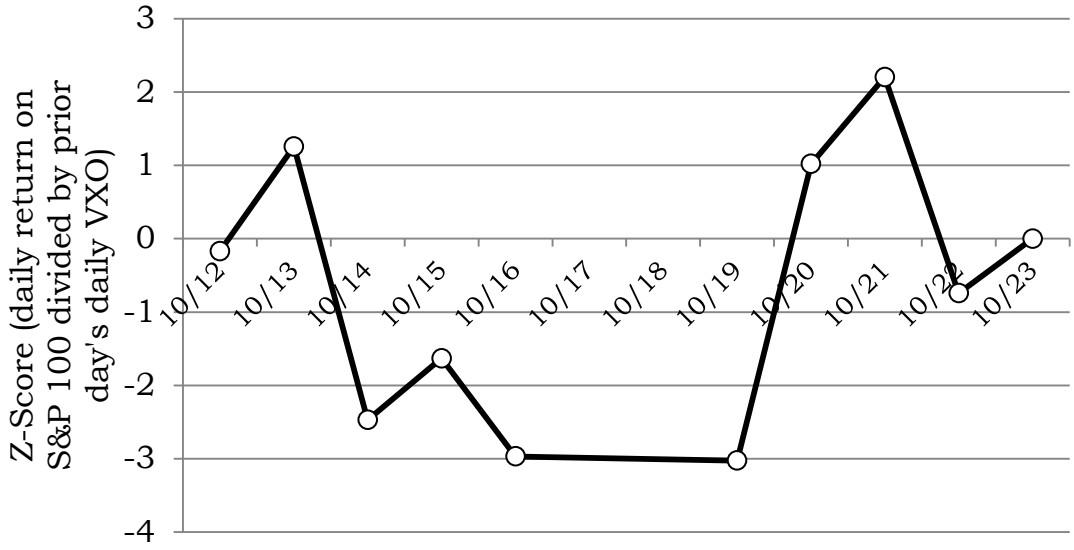

Source: *CBOE, Yahoo Finance*

Figure 5. CBOE *VIX* Index and S&P 500 Index

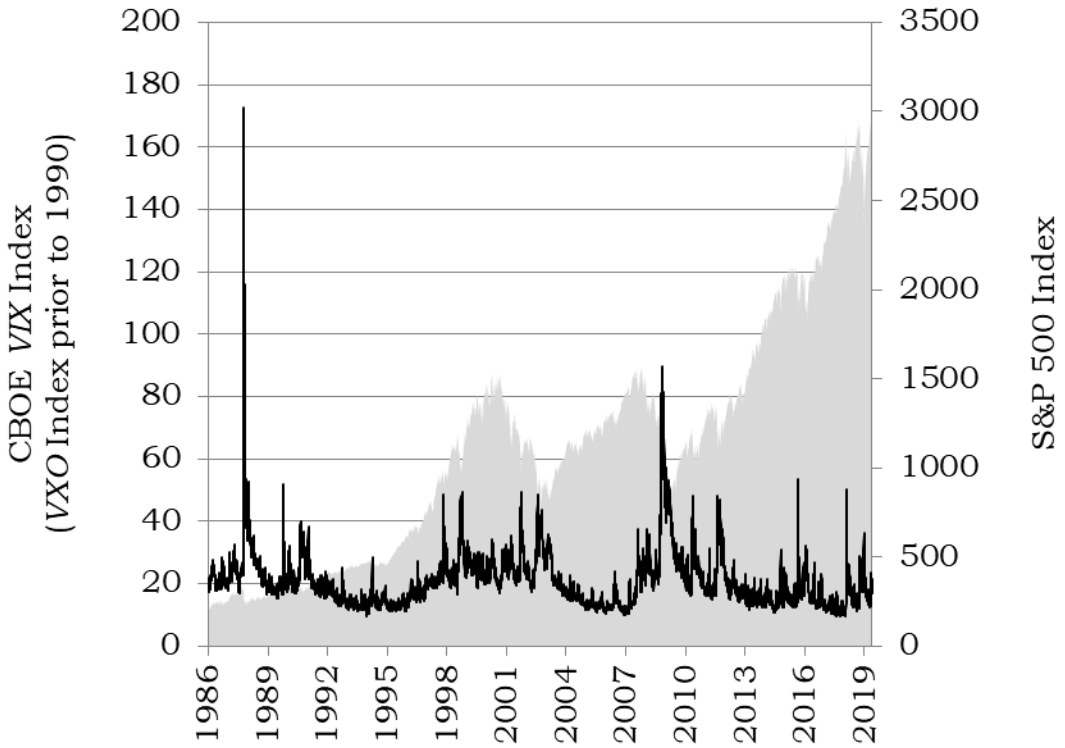

Source: *CBOE, Yahoo Finance*

Figure 6. Black-Scholes Vol. Smile for Various Assets on Dec. 29th, 2017

Dashed line is the underlying asset's price.

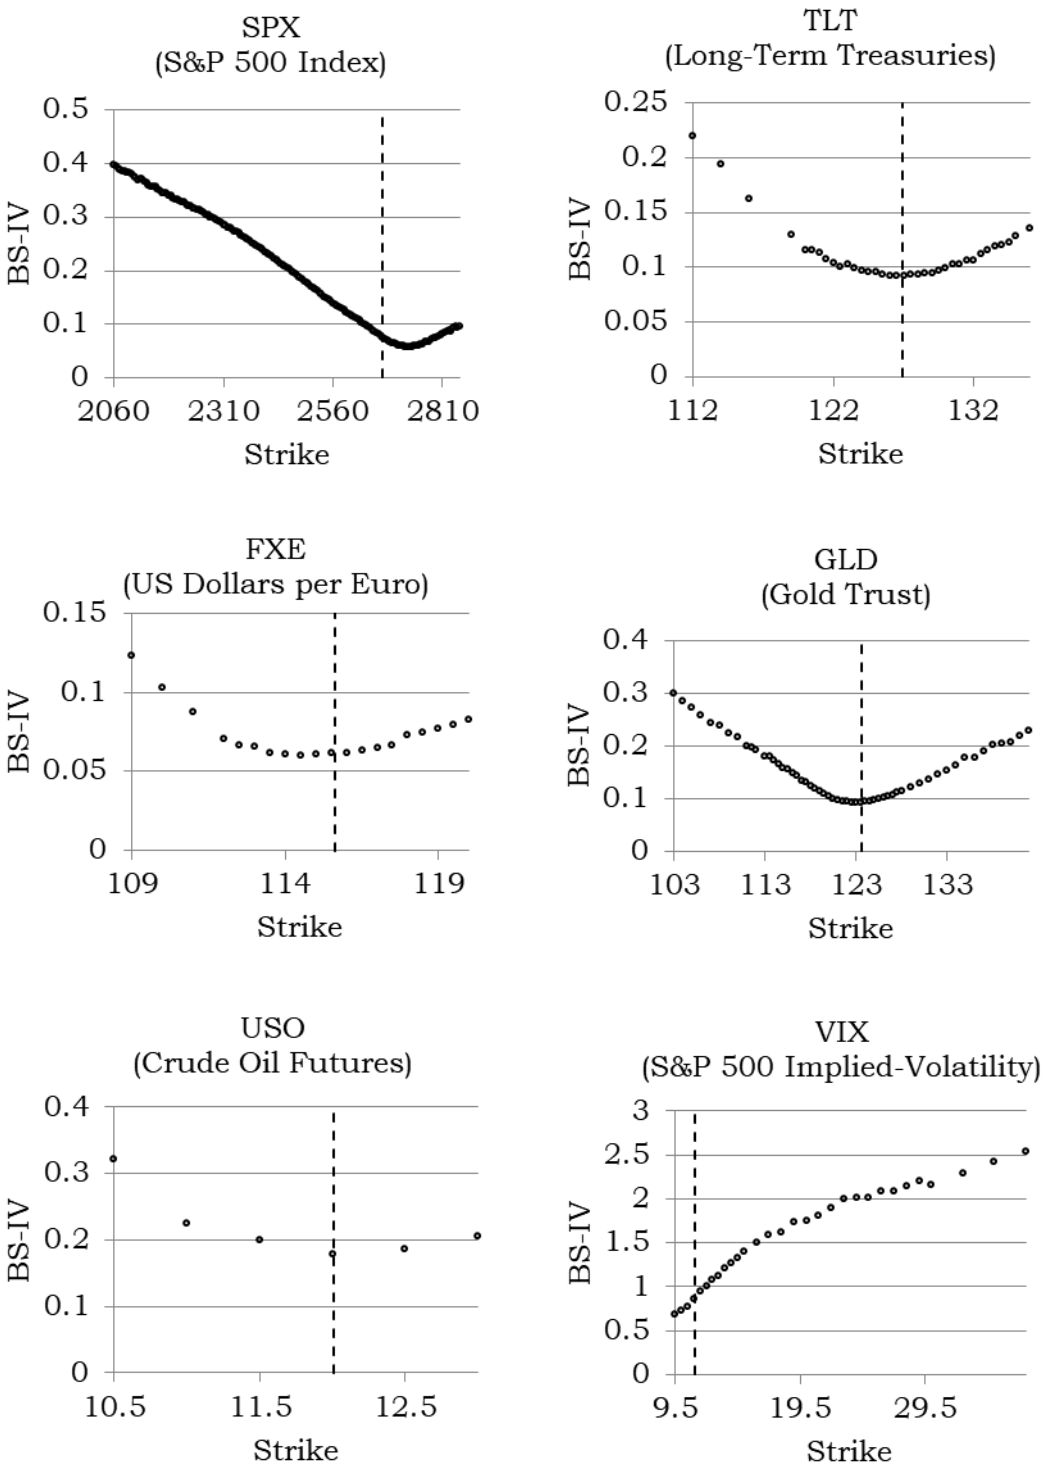

Source: *OptionMetrics*

Figure 7. Example: Estimating Imp.-Volatilities under the Binomial Model

Risk-Neutral Pricing:

Assumptions: $S_0 = \$100$, $K = \$105$, $r_f = 5\%$, $t = 1$, $C_0 = \$14.50$

Value of Call: $C_0 = p_U^Q \left(S_0 e^{r_f t + \sigma\sqrt{t}} - K\right) e^{-r_f t} = \frac{1-D}{U-D}(S_0 U - K e^{-r_f t})$

where, $U = e^{\sigma\sqrt{t}}$, $D = 1/U$

Therefore: $\$14.50 = \frac{1-D}{U-D}(\$100U - \$105e^{-5\%})$

Which implies: $U = 1.34$, $D = 0.75$, $p_U^Q = 42.8\%$, $p_D^Q = 57.2\%$, $[\boldsymbol{\sigma = 29.1\%}]$

Real-World Pricing:

Assumption: $r = 12\%$ (required return of the stock)

Therefore: $S_t = \$100e^{12\%} = p_U^P S_0 e^{12\%} U + p_D^P S_0 e^{12\%} D$

Which implies: $p_U^P = 55.1\%$, $p_D^P = 44.9\%$

$E[r_C] = 30.3\%$ (expected-return of the call)

$\sigma_C = 105.0\%$ (volatility of the call)

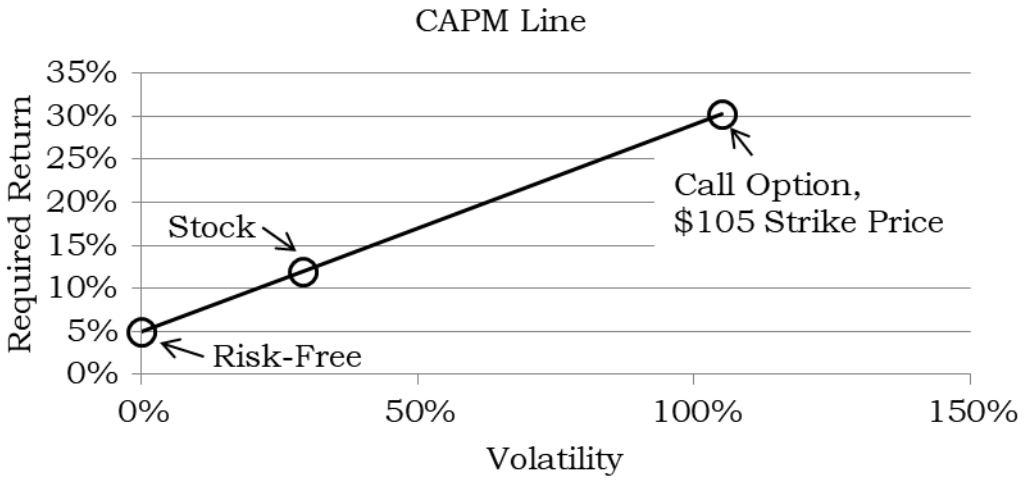

Figure 8: Example of a Volatility Smile

Obtained from S&P 500 Index options with an expiration date of Mar. 17, 2012 (time to expiration = 46 days). Underlying price = 1,312.41. Risk-free rate = 0.1995%. Forward price = 1,308.86.

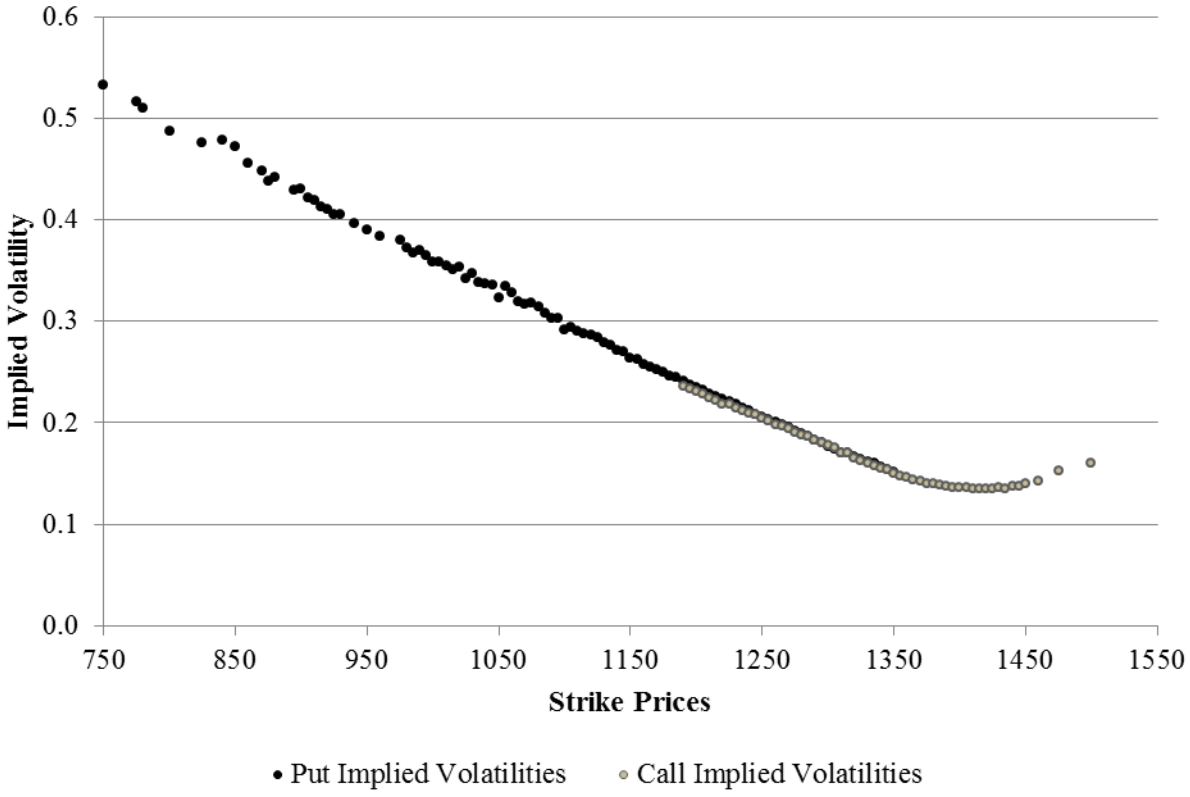

Figure 9: Fitted Spline, Fourth-Order with a Single Knot at the Money

Obtained from S&P 500 Index options on Jan. 31, 2012 with an expiration date of Mar. 17, 2012 (time to expiration = 46 days). Underlying price = 1,312.41. Risk-free rate = 0.1995%. Forward price = 1,308.86.

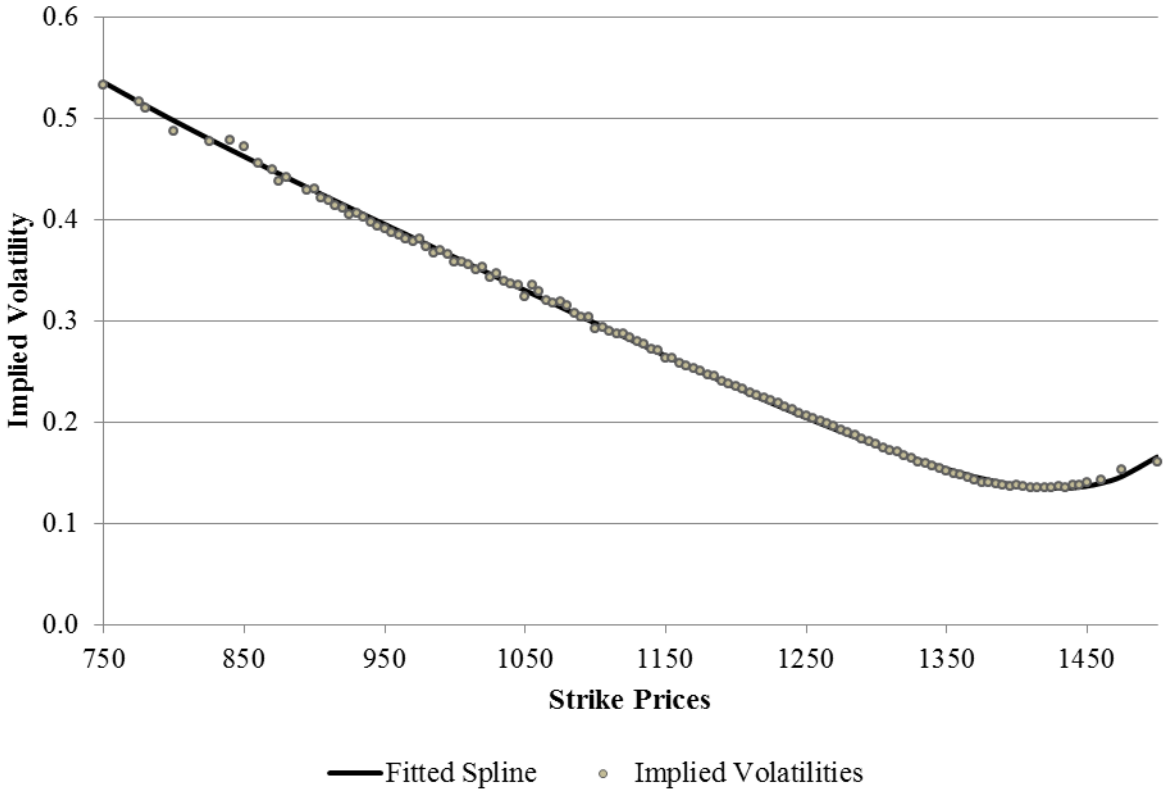

Figure 10: Implied Risk-Neutral Probability Distribution without Tails

Implied from S&P 500 Index options on Jan. 31, 2012 with an expiration date of Mar. 17, 2012 (time to expiration = 46 days). Underlying price = 1,312.41. Risk-free rate = 0.1995%. Forward price = 1,308.86. Density function and cumulative function obtained numerically via equations (5) and (6). Distribution displayed between the 2nd and 98th percentiles.

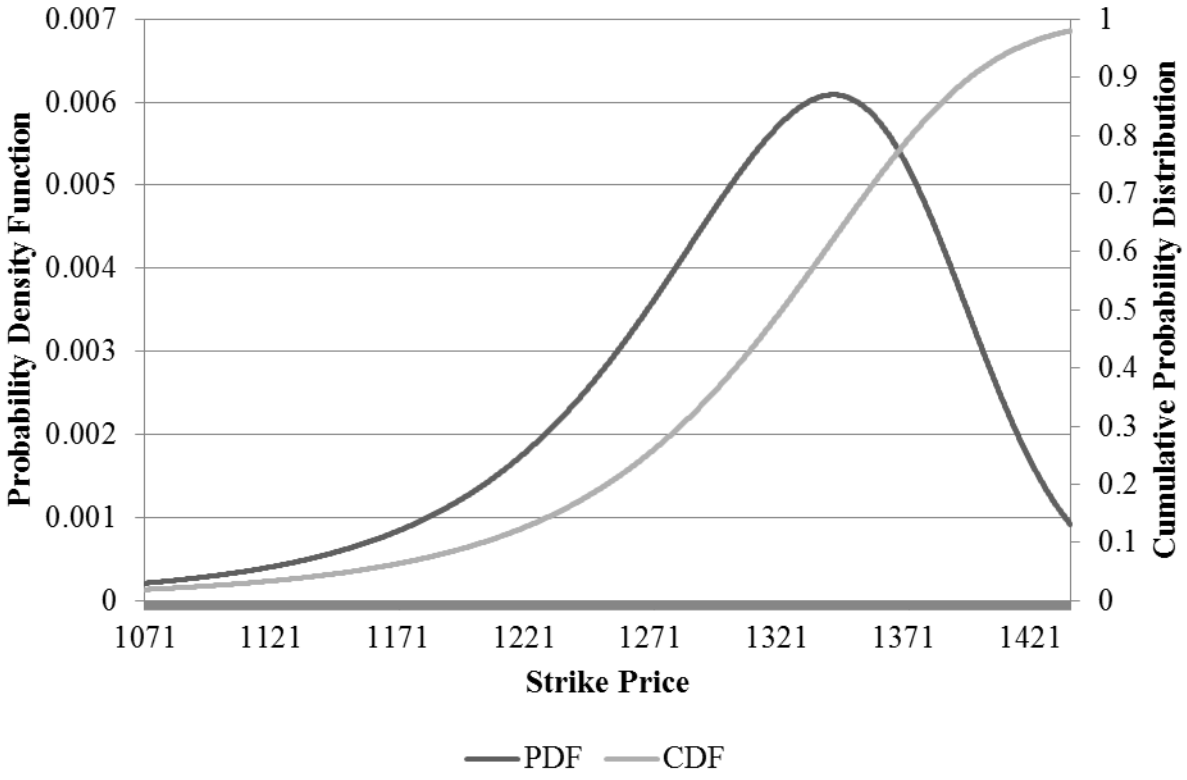

Figure 11: Truncated Distribution Compared to Non-Truncated Distribution

Implied from S&P 500 Index options on Jan. 31, 2012 with an expiration date of Mar. 17, 2012 (time to expiration = 46 days). Underlying price = 1,312.41. Risk-free rate = 0.1995%. Forward price = 1,308.86. Non-truncated distribution obtained numerically via equations (5) and (6) and displayed between the 2nd and 98th percentiles. Truncated distribution displayed over entire range of distribution.

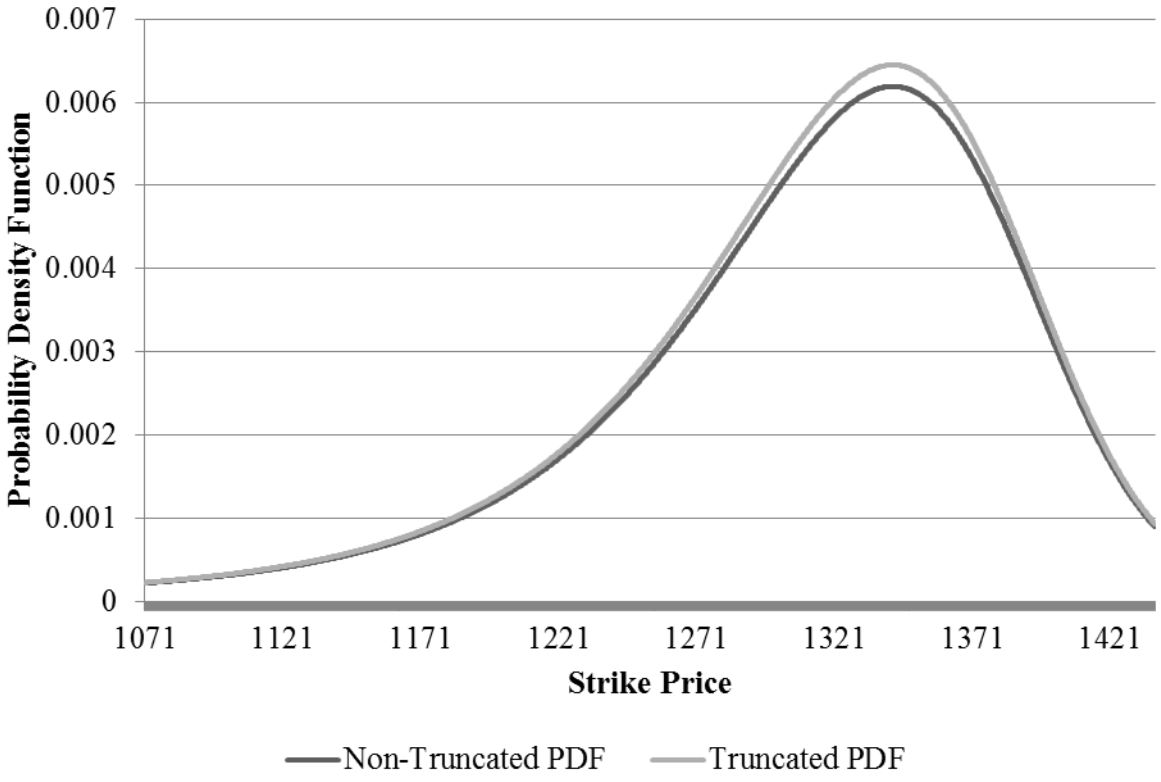

Figure 12: Lognormal Tails Compared to Tails Extrapolated from the Volatility Smile

Implied from S&P 500 Index options on Jan. 31, 2012 with an expiration date of Mar. 17, 2012 (time to expiration = 46 days). Underlying price = 1,312.41. Risk-free rate = 0.1995%. Forward price = 1,308.86. Dotted line represents the implied volatilities resulting from attaching lognormal tails at the 2nd and 98th percentiles of the probability density function. Solid line represents the fitted spline with tails extrapolated from the volatility smile. The left tail is a linear extrapolation of the volatility smile between the 2nd and 5th percentile. The right tail is a linear extrapolation of the volatility smile between the 95th and the 98th percentile. Vertical lines represent the 2nd, 5th, 95th, and 98th percentiles. 2nd percentile = 1,071.28. 5th percentile = 1,151.49. 95th percentile = 1,416.01. 98th percentile = 1,437.46.

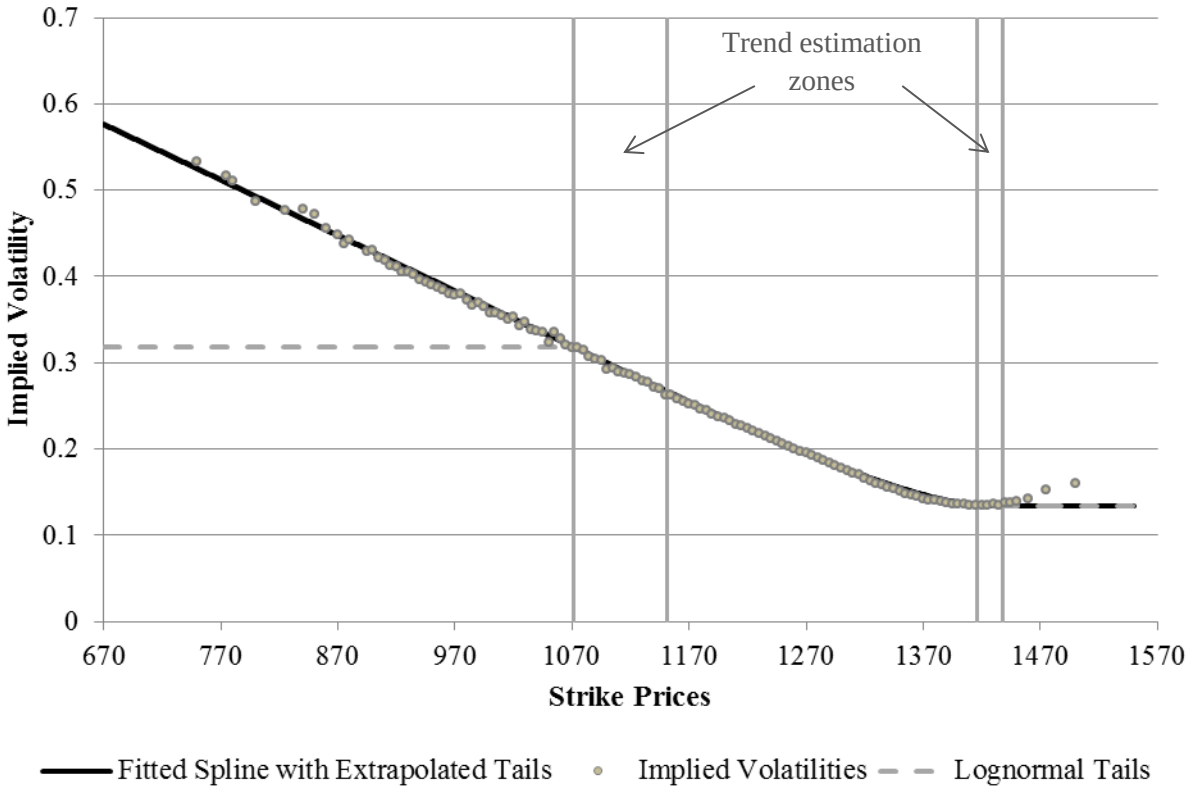

Figure 13: Generalized Extreme Value Density Function Tails Compared to Density Tails Extrapolated from the Volatility Smile

Implied from S&P 500 Index options on Jan. 31, 2012 with an expiration date of Mar. 17, 2012 (time to expiration = 46 days). Underlying price = 1,312.41. Risk-free rate = 0.1995%. Forward price = 1,308.86. Dark line represents the fitted spline with left tail extrapolated from the volatility smile. Light line represents a tail from a Generalized Extreme Value (GEV) distribution chosen such that the density and cumulative distribution match the spline-fitted distribution at the 5th percentile, and also so that their densities match at the 2nd percentile. The location parameter of the GEV is estimated to be 1280.61; the scale parameter 47.49; and the shape parameter 0.1596. Vertical lines represent the 2nd and 5th percentiles. 2nd percentile = 1,071.28. 5th percentile = 1,151.49.

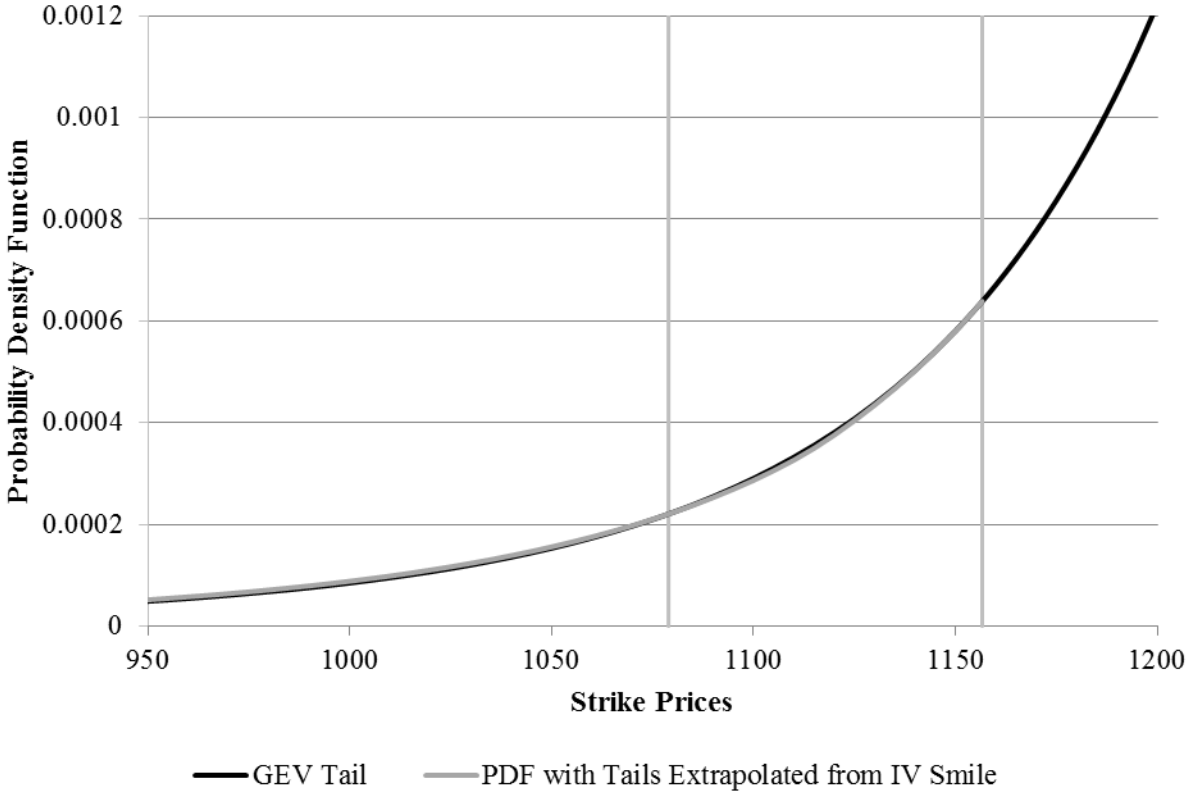

Figure 14: Generalized Extreme Value Cumulative Distribution Tails Compared to Cumulative Distribution Tails Extrapolated from the Volatility Smile

Implied from S&P 500 Index options on Jan. 31, 2012 with an expiration date of Mar. 17, 2012 (time to expiration = 46 days). Underlying price = 1,312.41. Risk-free rate = 0.1995%. Forward price = 1,308.86. Dark line represents the fitted spline with left tail extrapolated from the volatility smile. Light line represents a tail from a Generalized Extreme Value (GEV) distribution chosen such that the density and cumulative distribution match the spline-fitted distribution at the 5th percentile, and also so that their densities match at the 2nd percentile. The location parameter of the GEV is estimated to be 1280.61; the scale parameter 47.49; and the shape parameter 0.1596. Vertical lines represent the 2nd and 5th percentiles. 2nd percentile = 1,071.28. 5th percentile = 1,151.49.

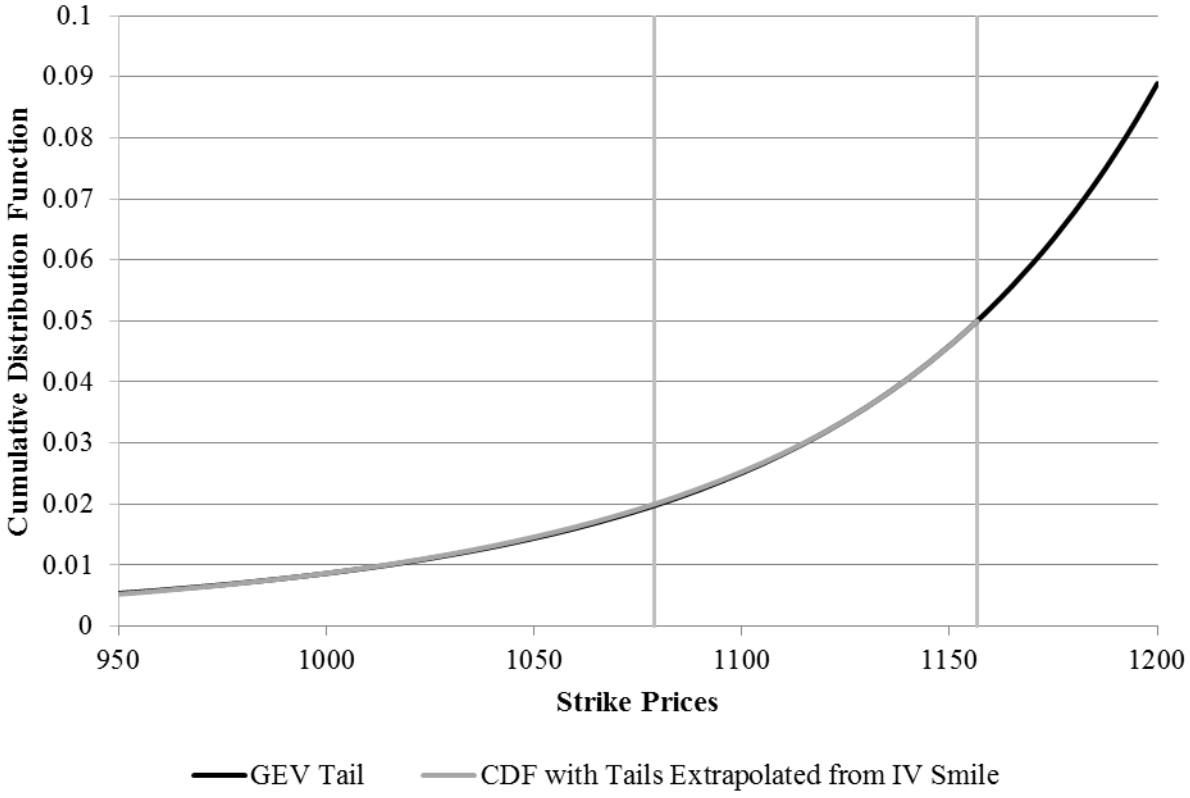

Figure 15: Time Series of Option-Implied Moments using Truncated Tails

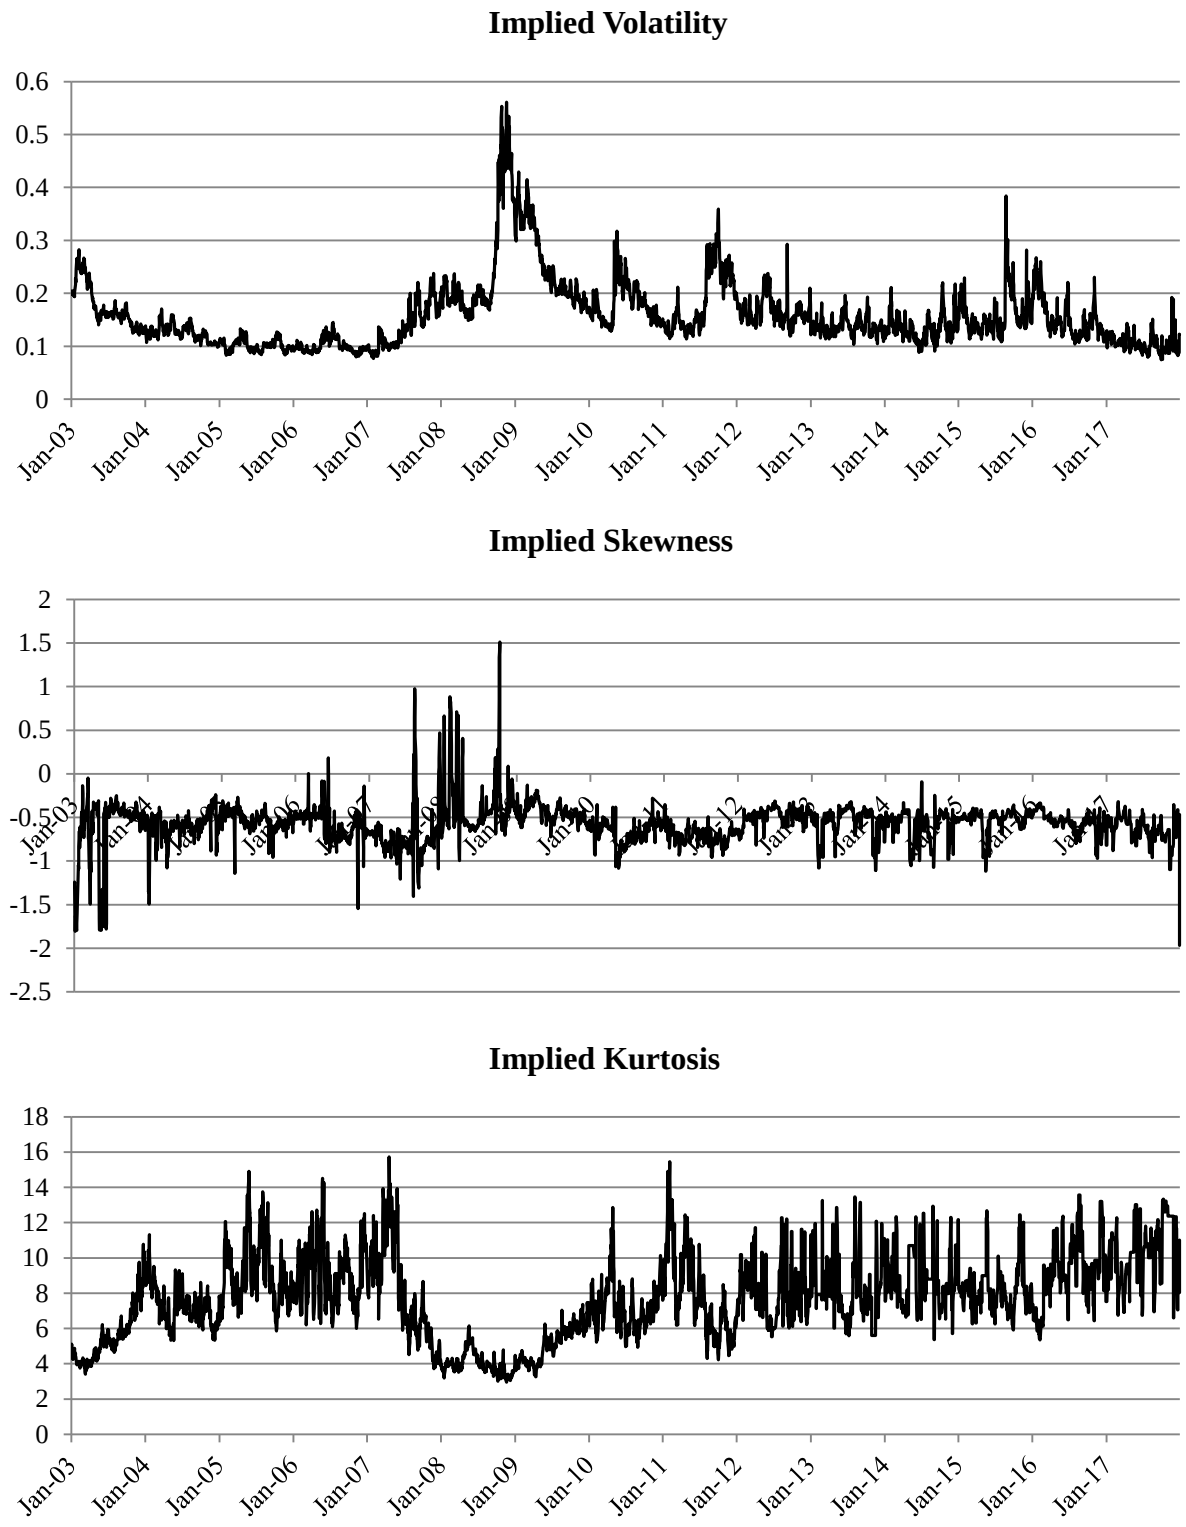

Figure 16: Time Series of Option-Implied Moments using Lognormal Tails

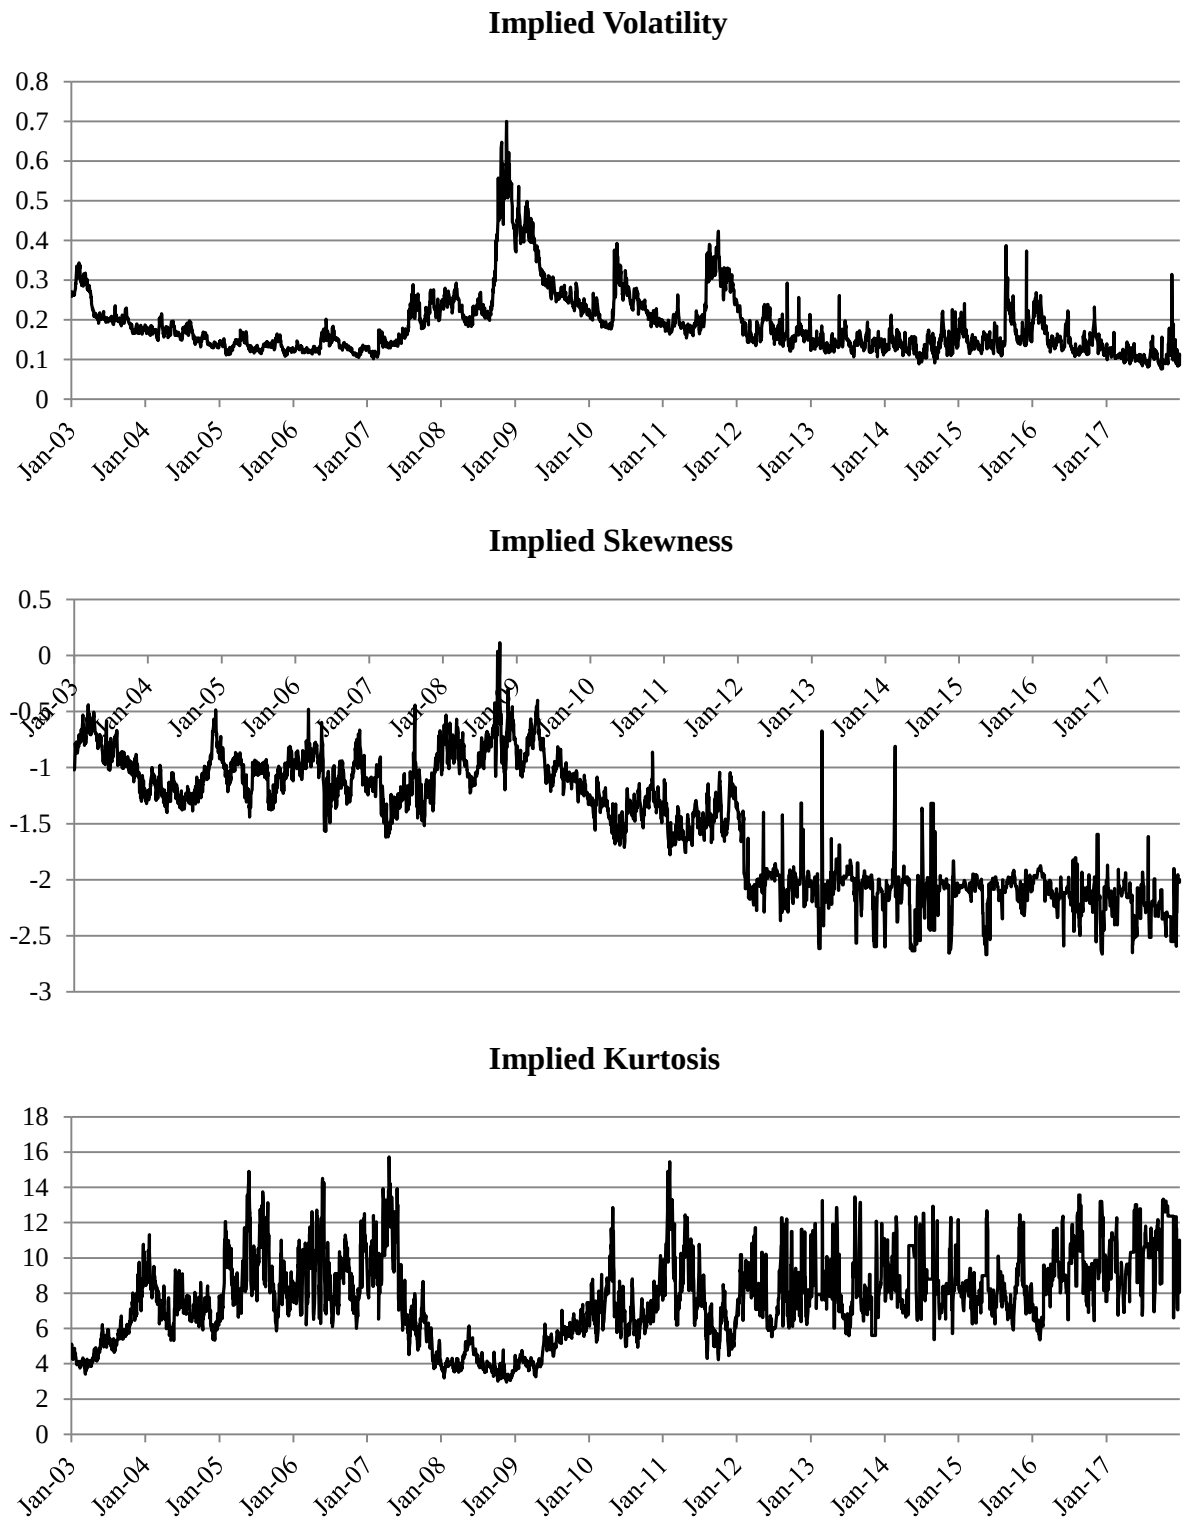

Figure 17: Time Series of Option-Implied Moments using GEV Tails

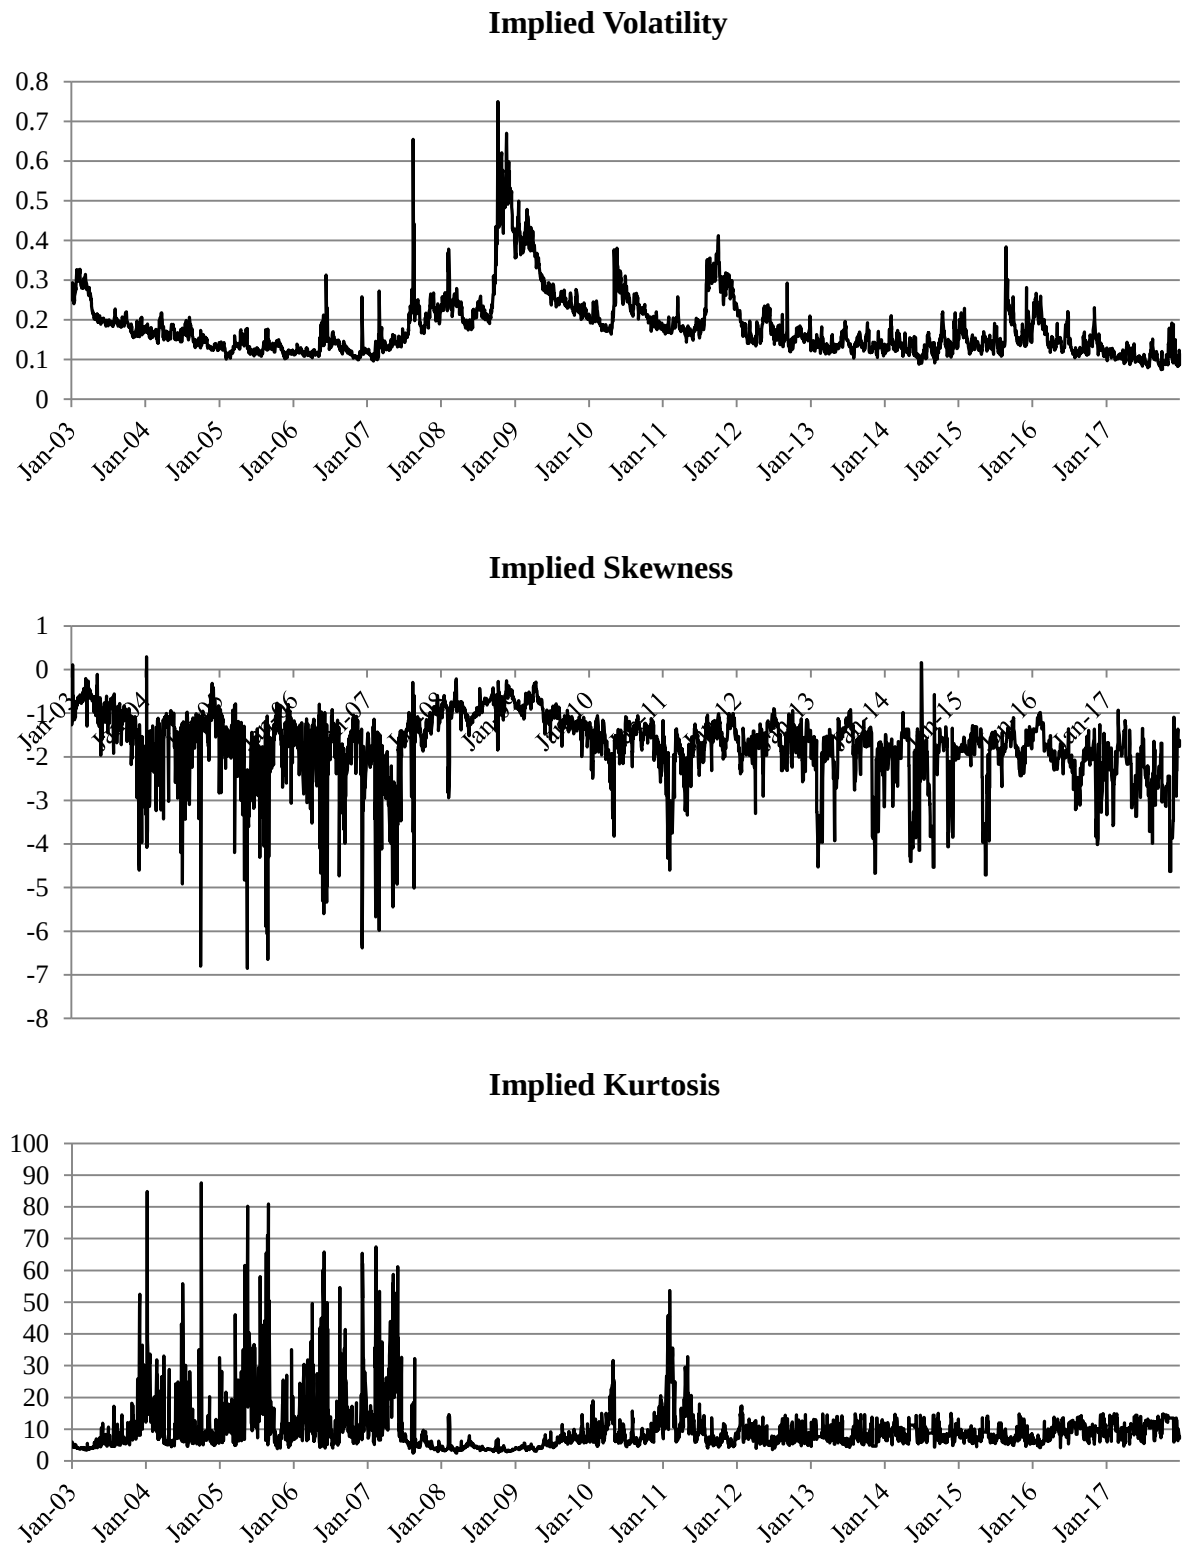

Figure 18: Time Series of Option-Implied Moments using Smile-Extrapolated Tails

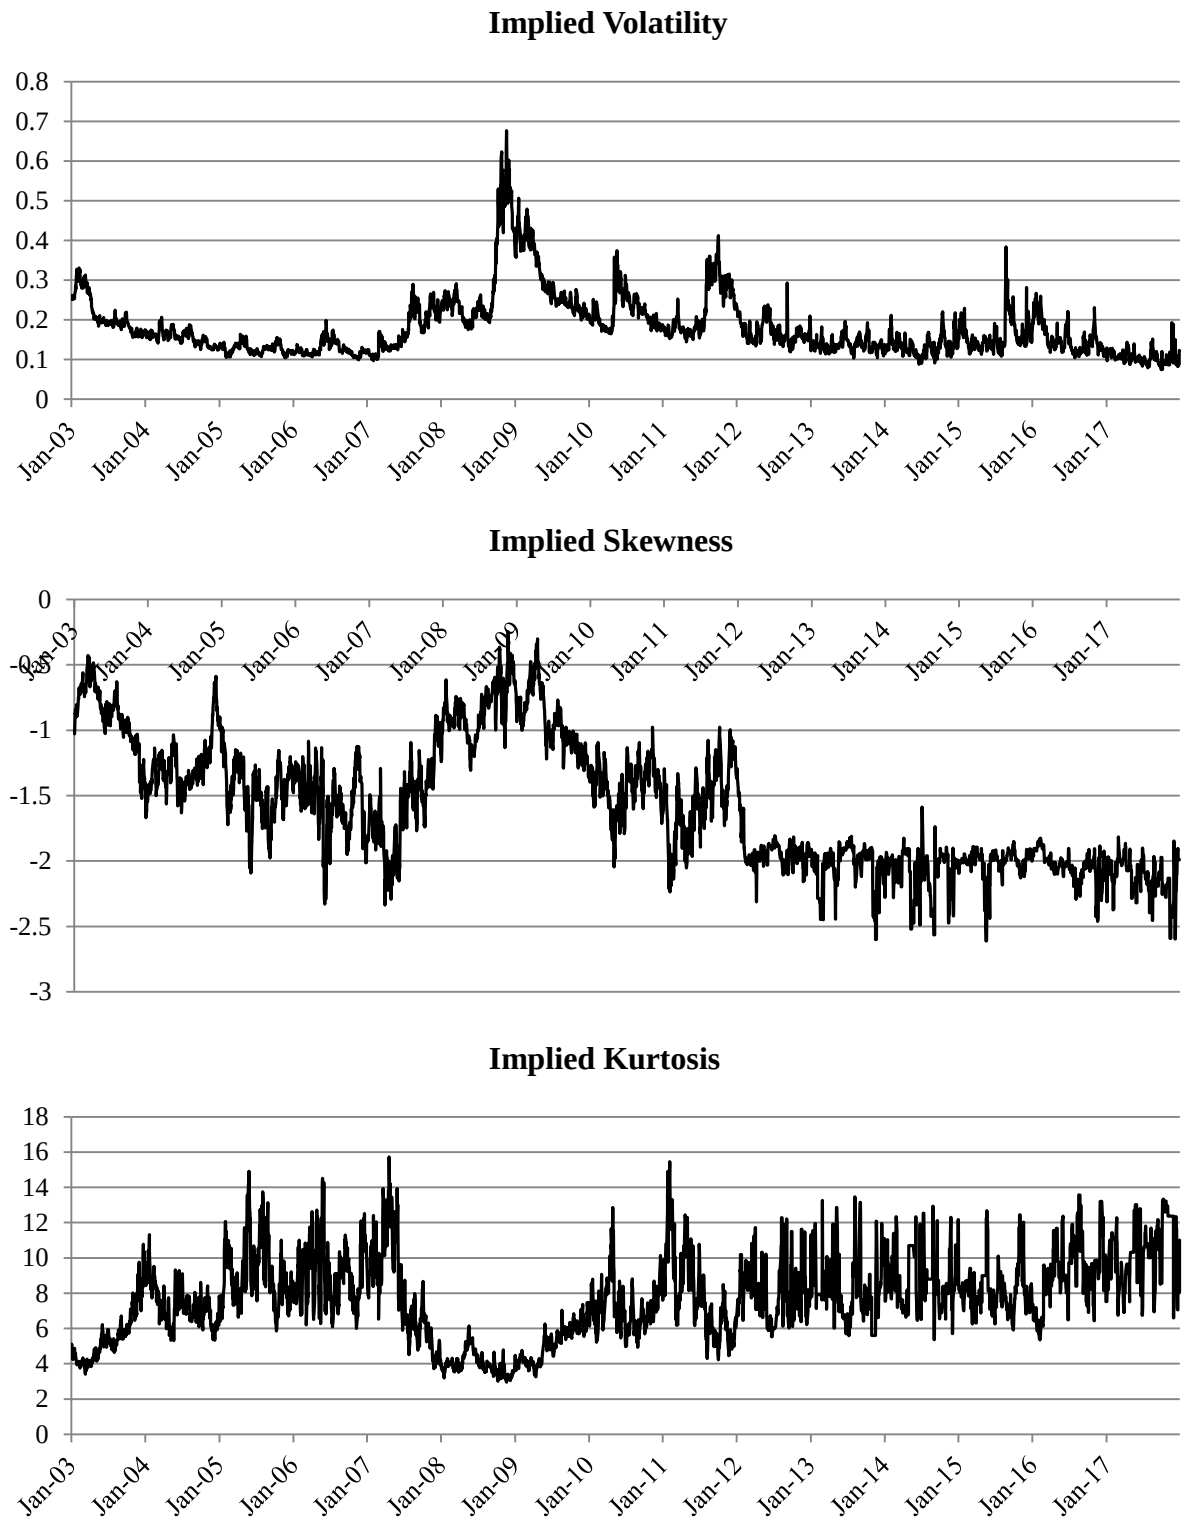

Figure 19: Normal Probability Plot, Implied Volatility

Normal probability plot comparing standardized residuals from GARCH(1,1) model against theoretical z-scores from a normal distribution. *PPCC* is the probability plot correlation coefficient. For comparison, straight line represents result if residuals are perfectly normal.

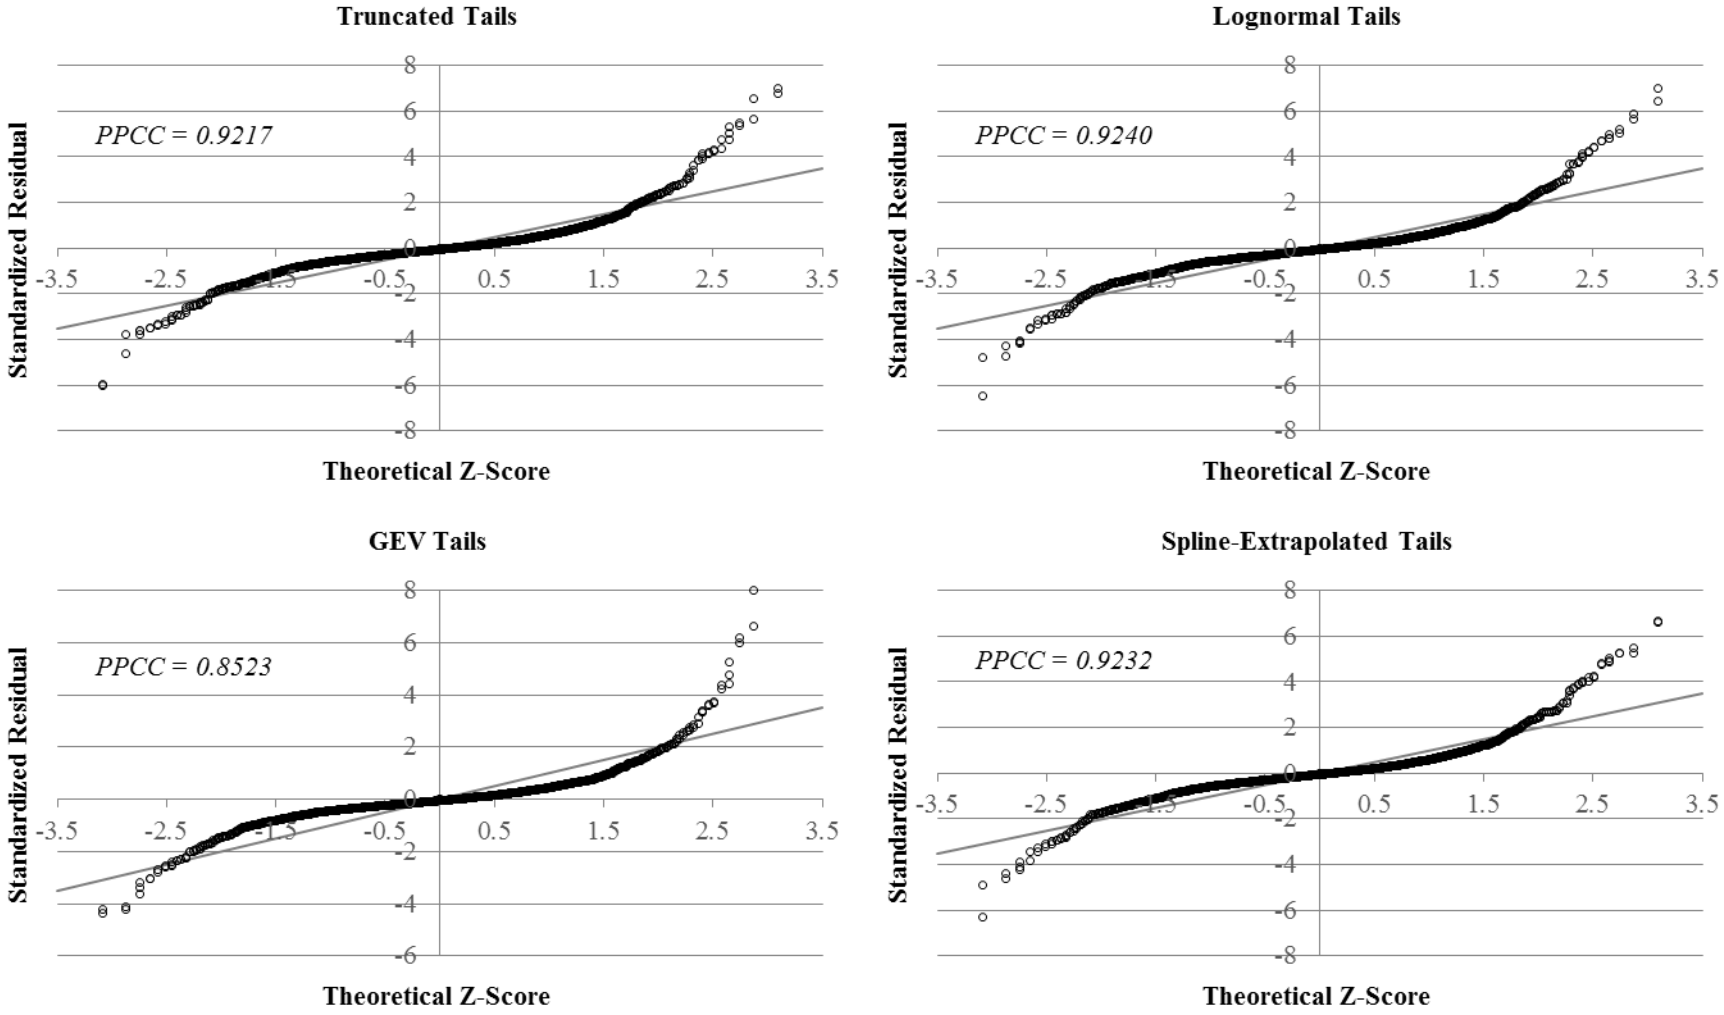

Figure 20: Normal Probability Plot, Implied Skewness

Normal probability plot comparing standardized residuals from GARCH(1,1) model against theoretical z-scores from a normal distribution. *PPCC* is the probability plot correlation coefficient. For comparison, straight line represents result if residuals are perfectly normal.

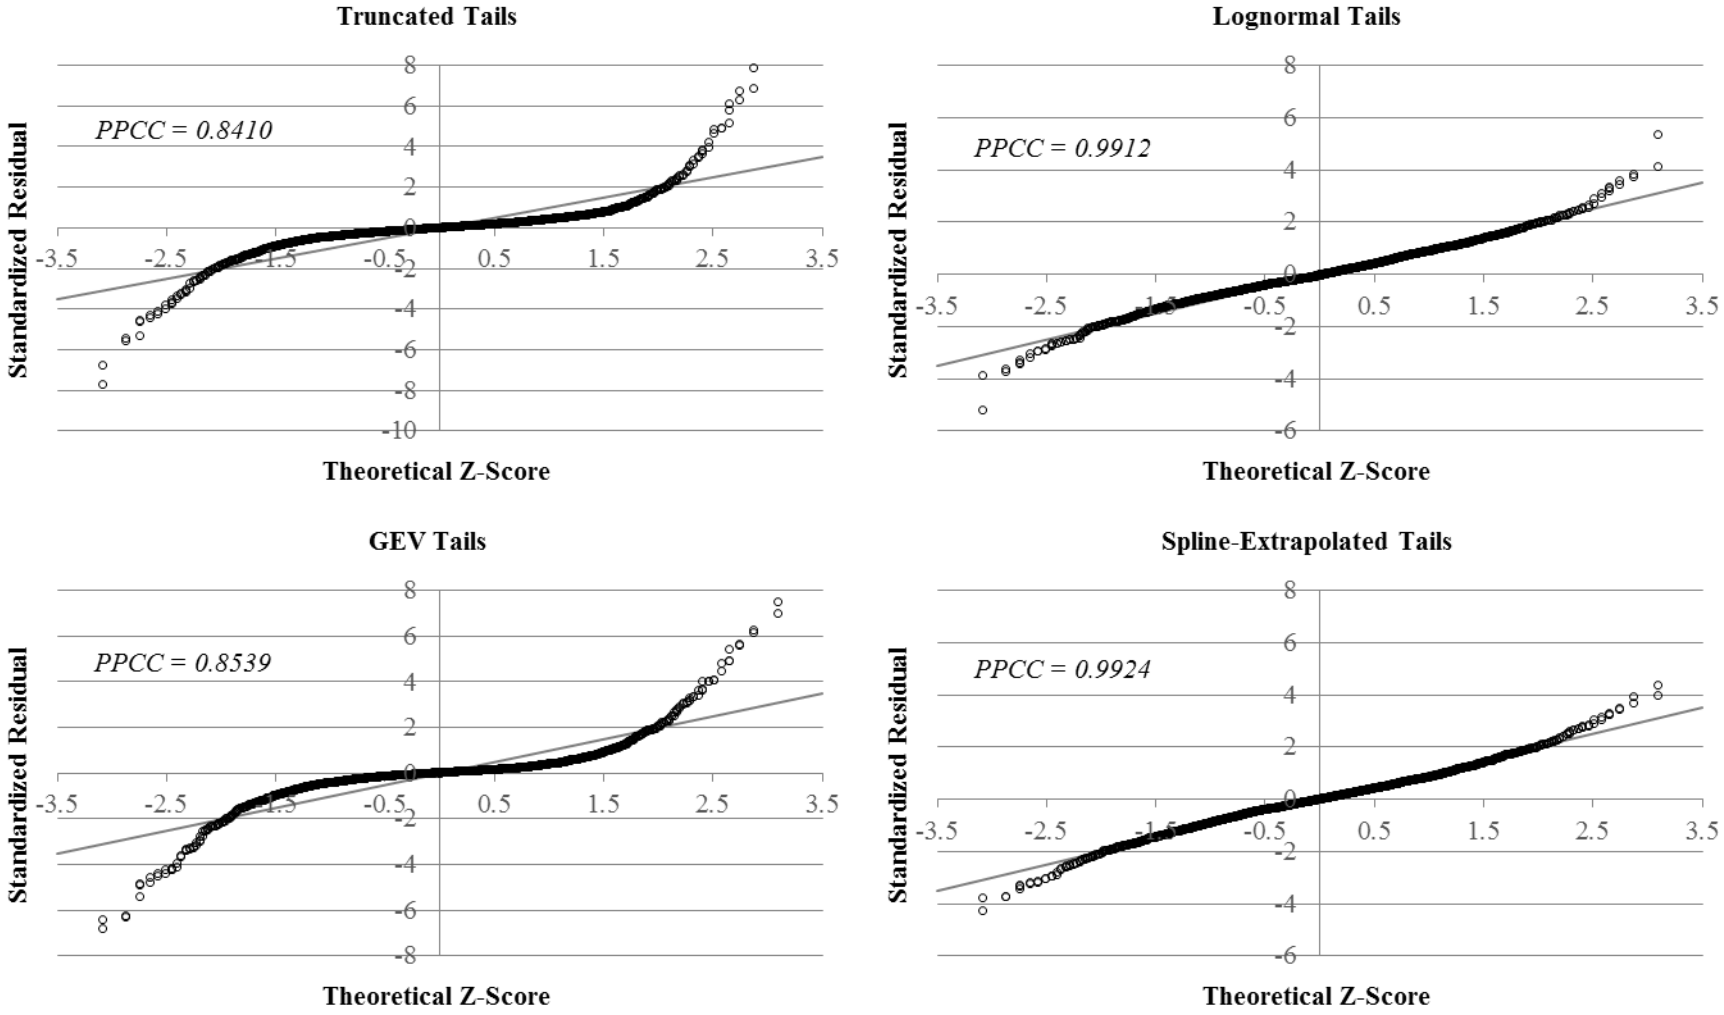

Figure 21: Normal Probability Plot, Implied Kurtosis

Normal probability plot comparing standardized residuals from GARCH(1,1) model against theoretical z-scores from a normal distribution. *PPCC* is the probability plot correlation coefficient. For comparison, straight line represents result if residuals are perfectly normal.

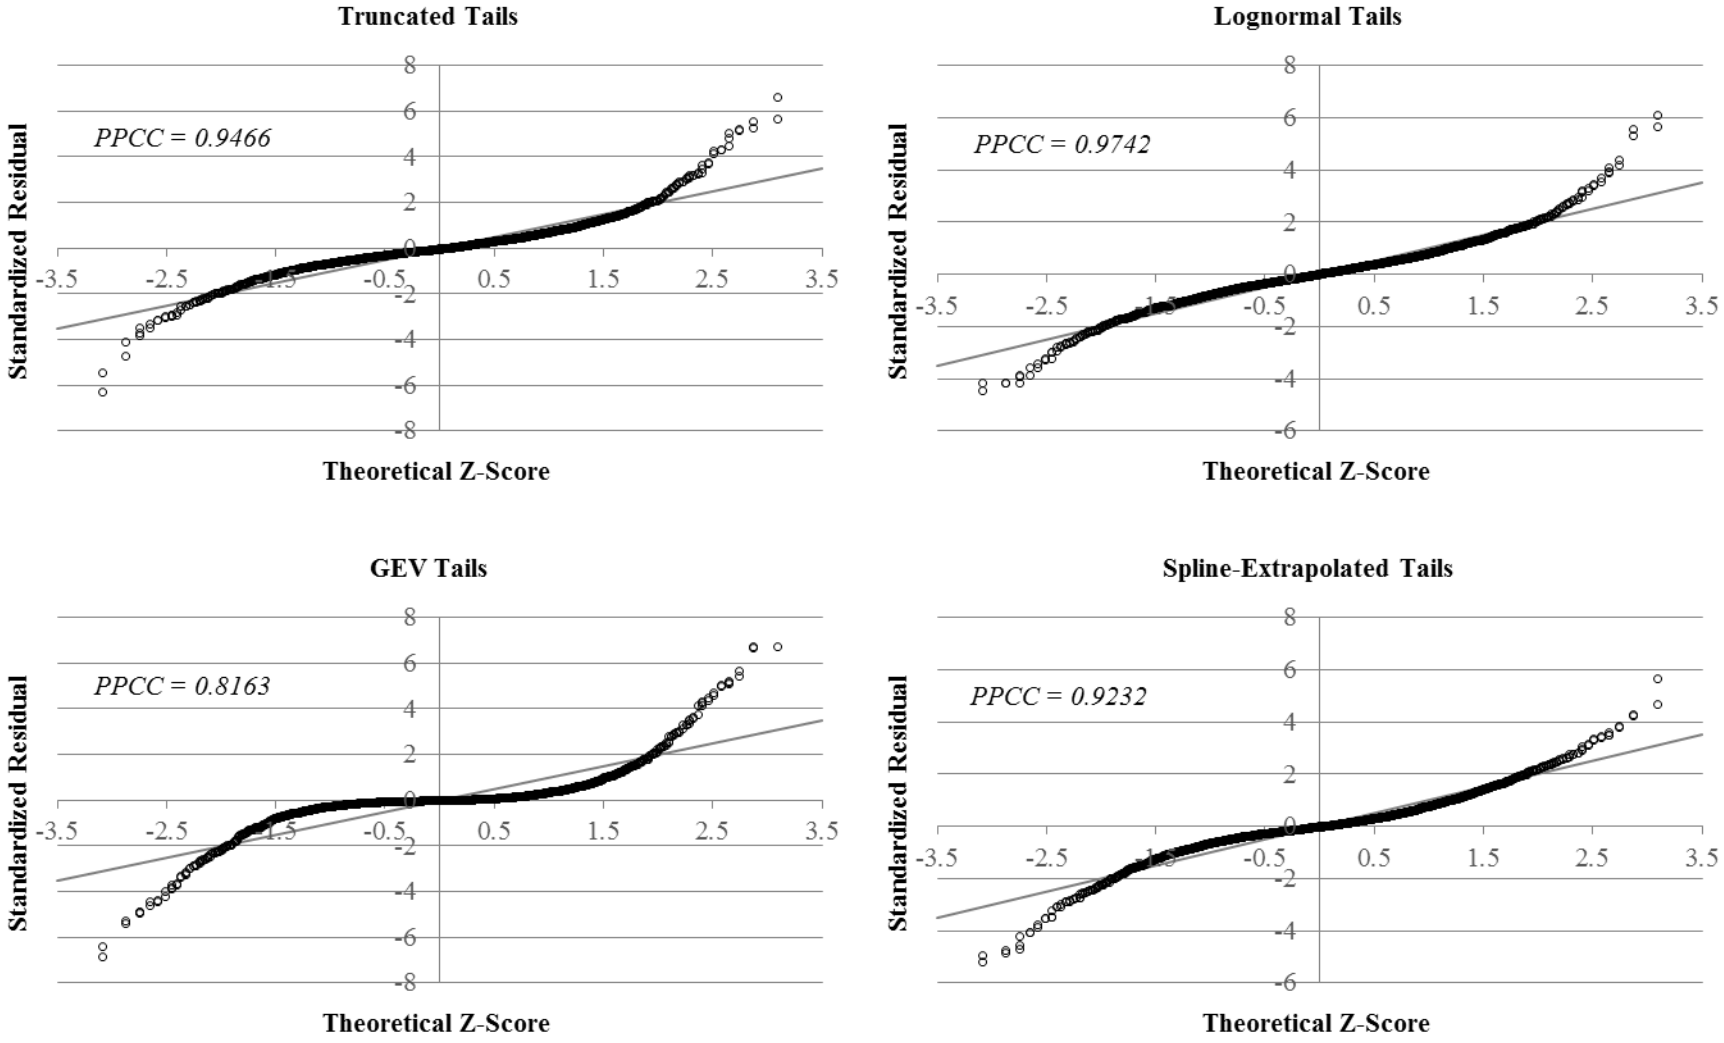

Figure 22: Cumulative Distribution Function, Implied Volatility

Dark line is the cumulative distribution function (*cdf*) of the standardized residuals from GARCH(1,1) model. For comparison, dashed, light line is the *cdf* for the normal distribution. Kolmogorov-Smirnov (*KS*) statistic is the supremum of the absolute distance between the two *cdf*s.

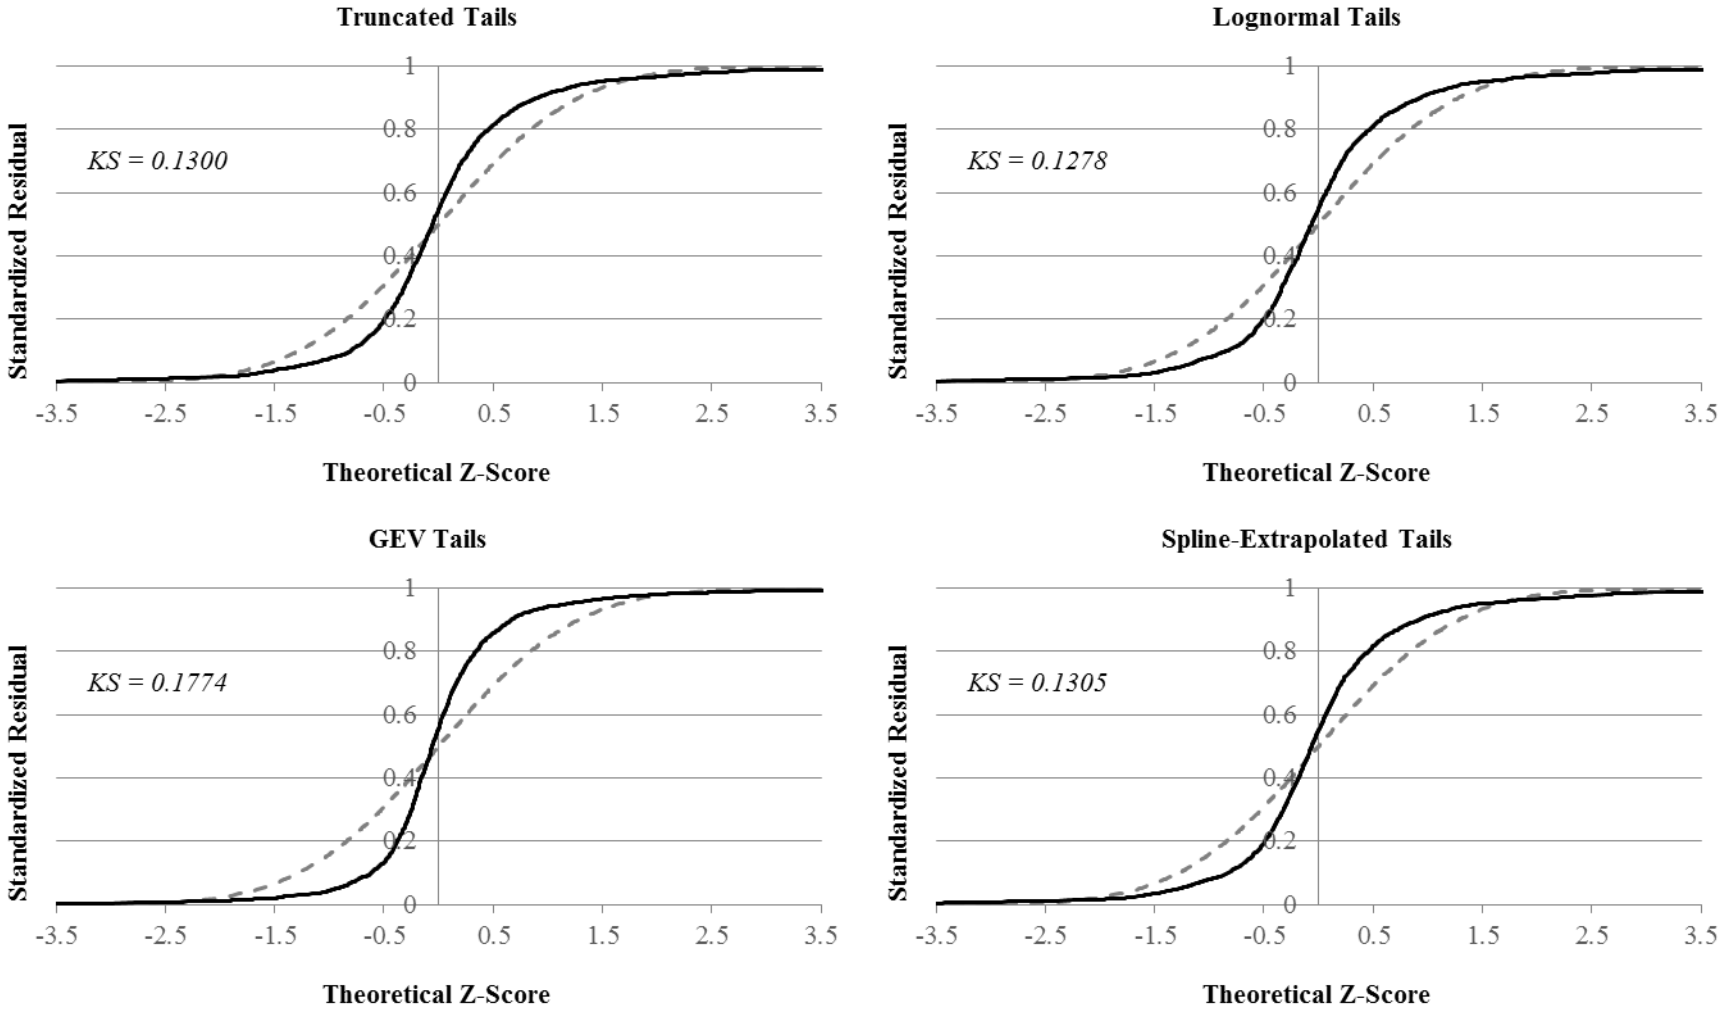

Figure 23: Cumulative Distribution Function, Implied Skewness

Dark line is the cumulative distribution function (*cdf*) of the standardized residuals from GARCH(1,1) model. For comparison, dashed, light line is the *cdf* for the normal distribution. Kolmogorov-Smirnov (*KS*) statistic is the supremum of the absolute distance between the two *cdf*s.

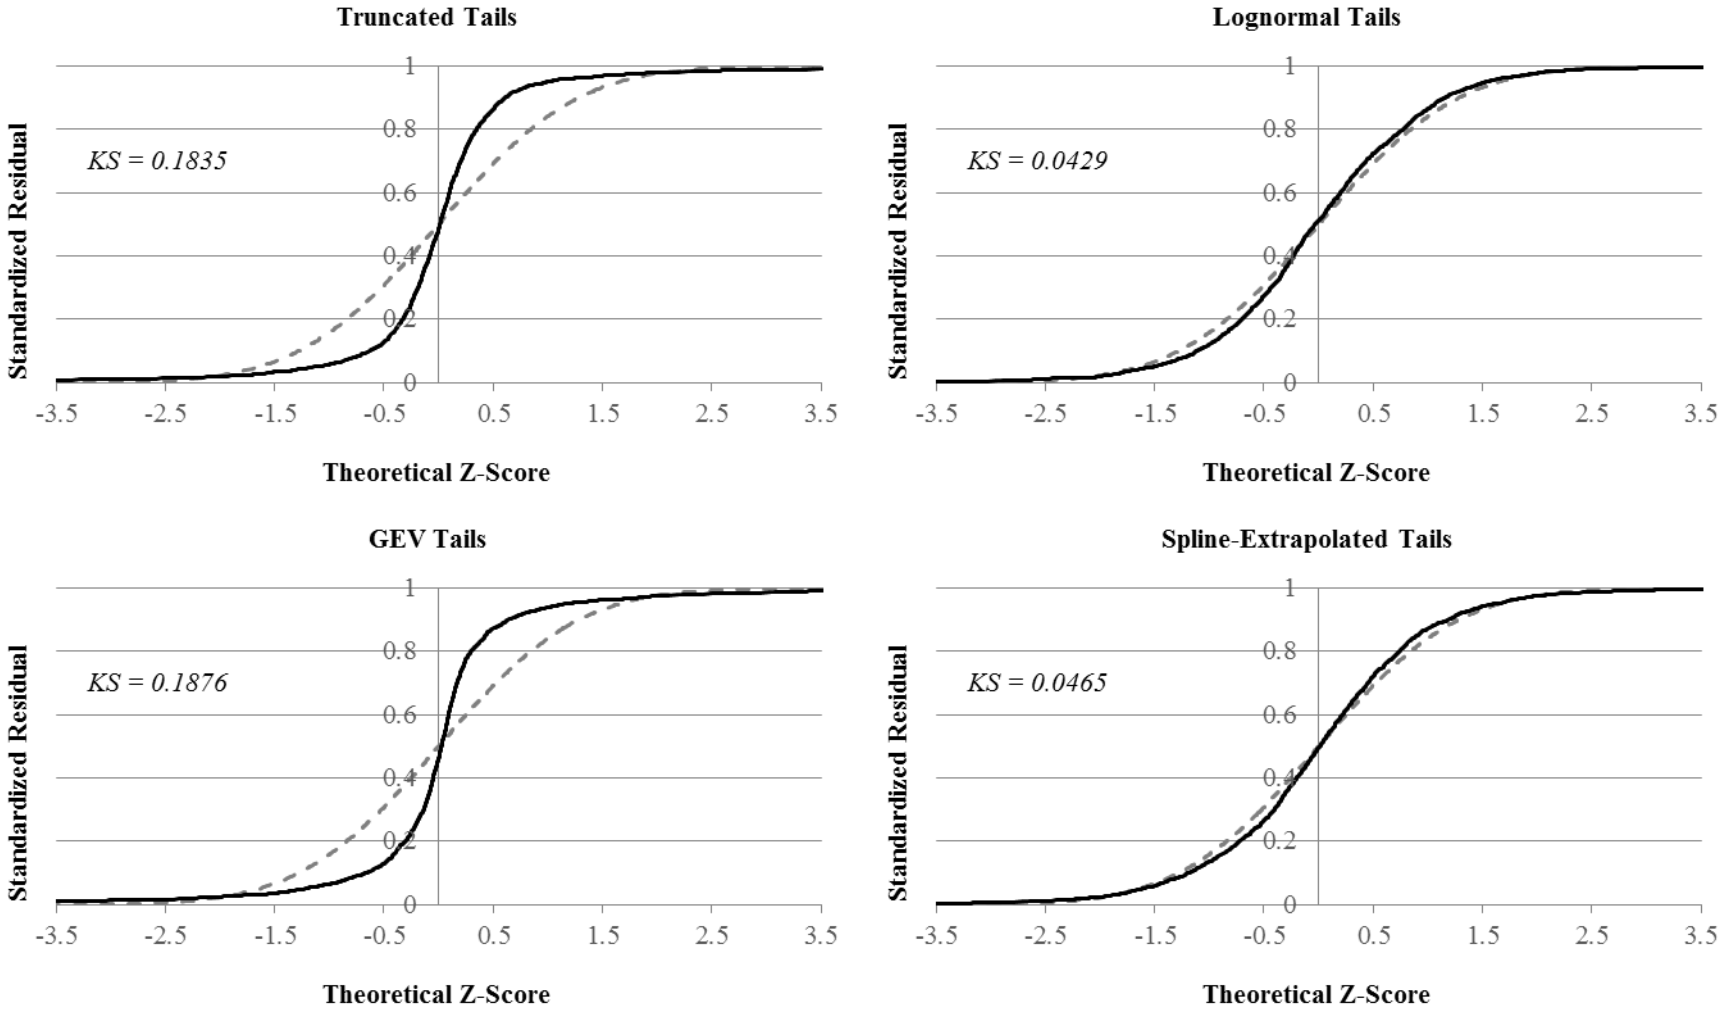

Figure 24: Cumulative Distribution Function, Implied Kurtosis

Dark line is the cumulative distribution function (*cdf*) of the standardized residuals from GARCH(1,1) model. For comparison, dashed, light line is the *cdf* for the normal distribution. Kolmogorov-Smirnov (*KS*) statistic is the supremum of the absolute distance between the two *cdf*s.

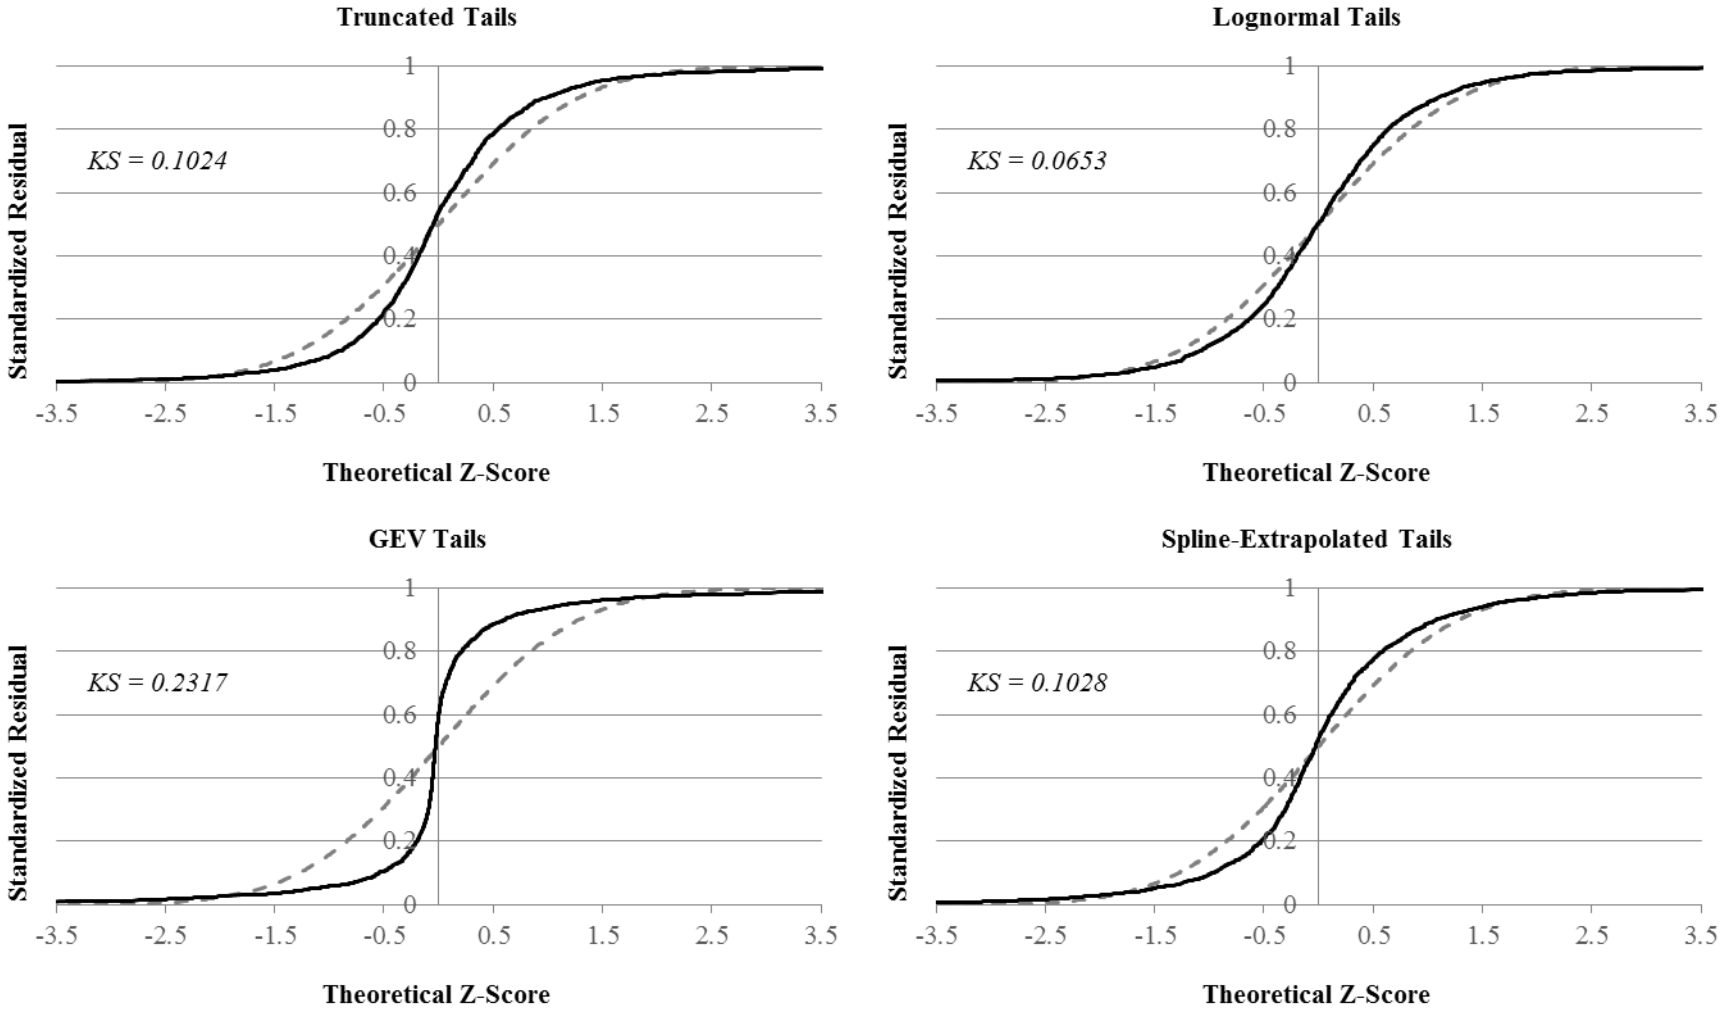

Figure 25: Separate Fitted Splines for Put and Calls

Obtained from S&P 500 Index options with an expiration date of Mar. 17, 2012 (time to expiration = 46 days). Underlying price = 1,312.41. Risk-free rate = 0.1995%. Implied volatilities calculated using the underlying price as the forward price.

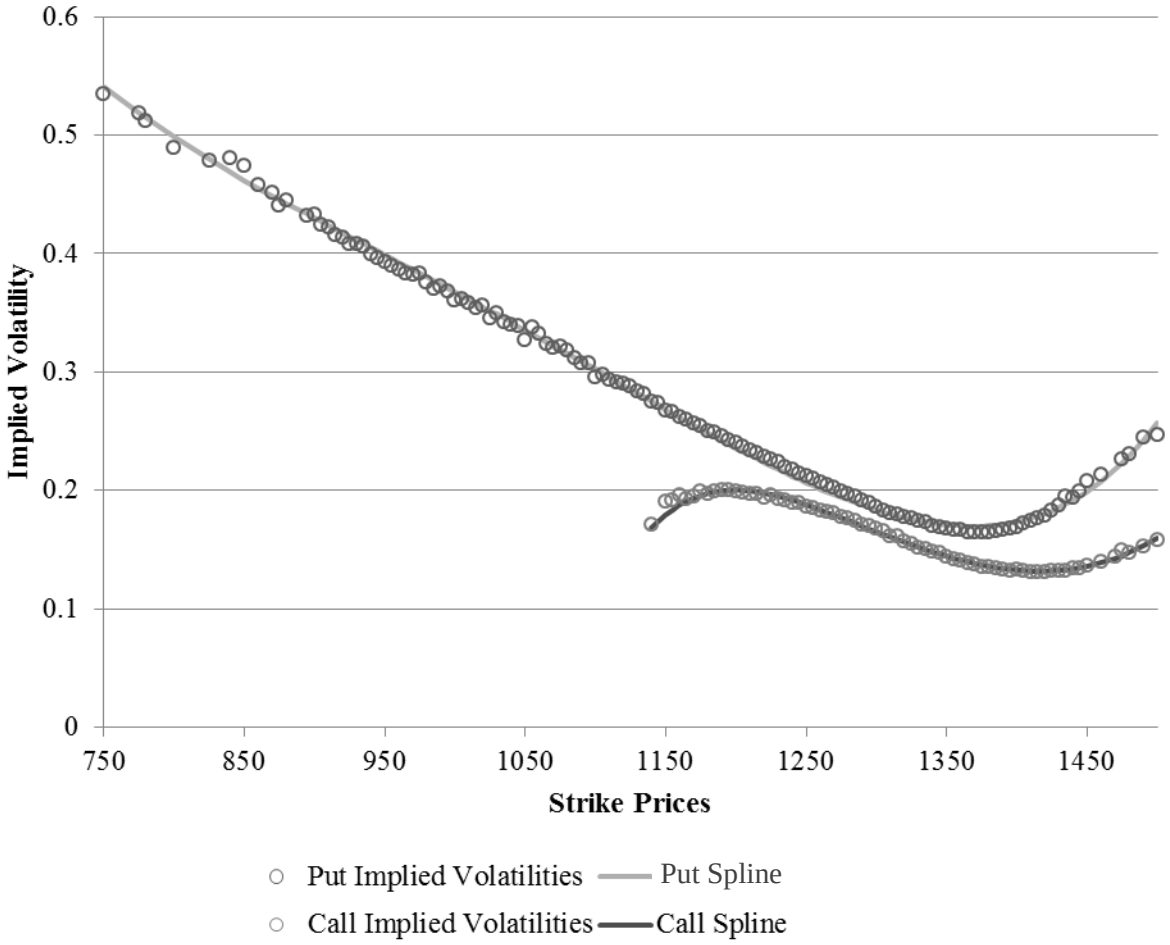

Figure 26: Finding the Forward Price from Interpolated Put and Call Prices

Obtained from S&P 500 Index options with an expiration date of Mar. 17, 2012 (time to expiration = 46 days). Underlying price = 1,312.41. Risk-free rate = 0.1995%. Implied volatilities calculated using the underlying price as the forward value. Vertical line represents the option-implied forward price (1,308.86), where the put price equals the call price.

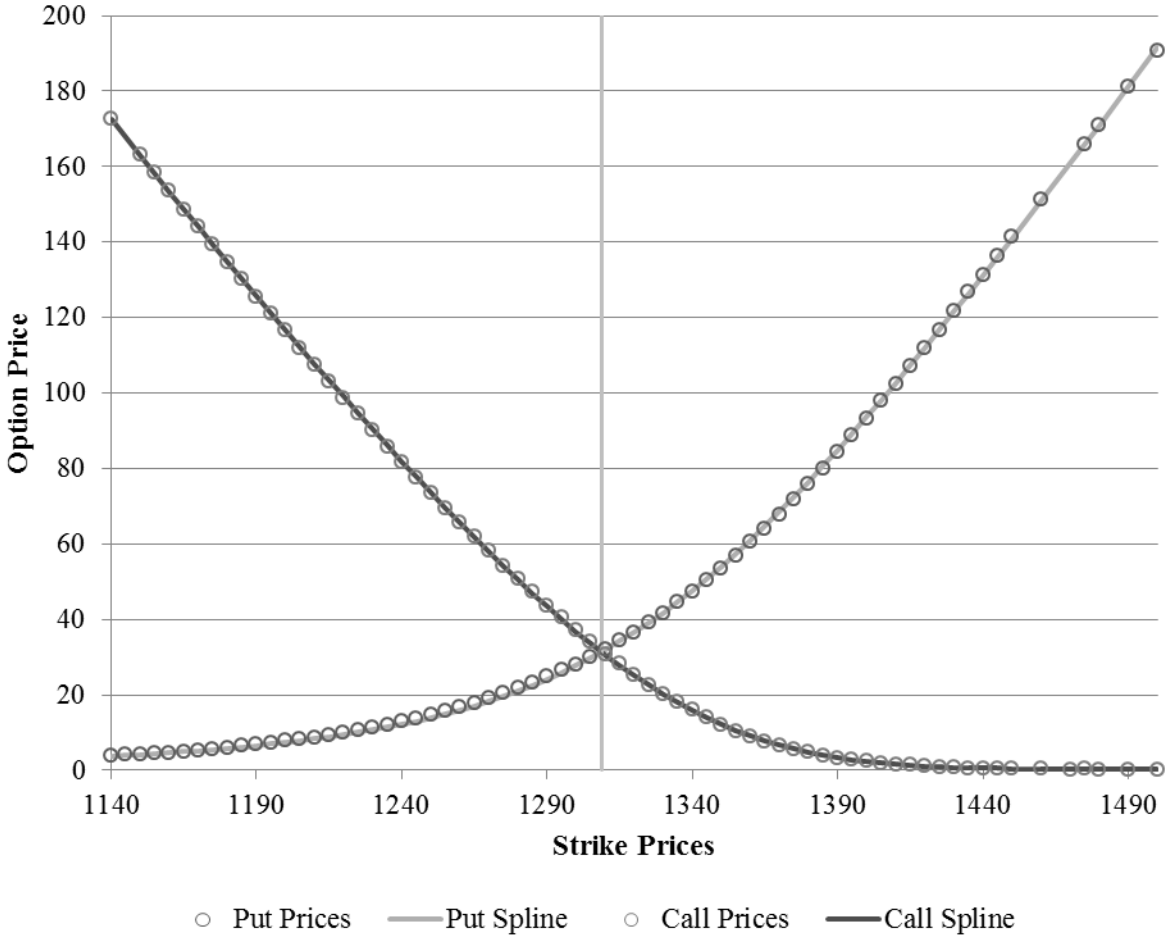

Figure 27: Convergence of Implied Kurtosis Estimate

Obtained from S&P 500 Index options with an expiration date of Mar. 17, 2012 (time to expiration = 46 days). Underlying price = 1,312.41. Risk-free rate = 0.1995%. Vertical line represents the option-implied forward price (1,308.86). Graph depicts each one cent strike's contribution to implied kurtosis. See equation (19).

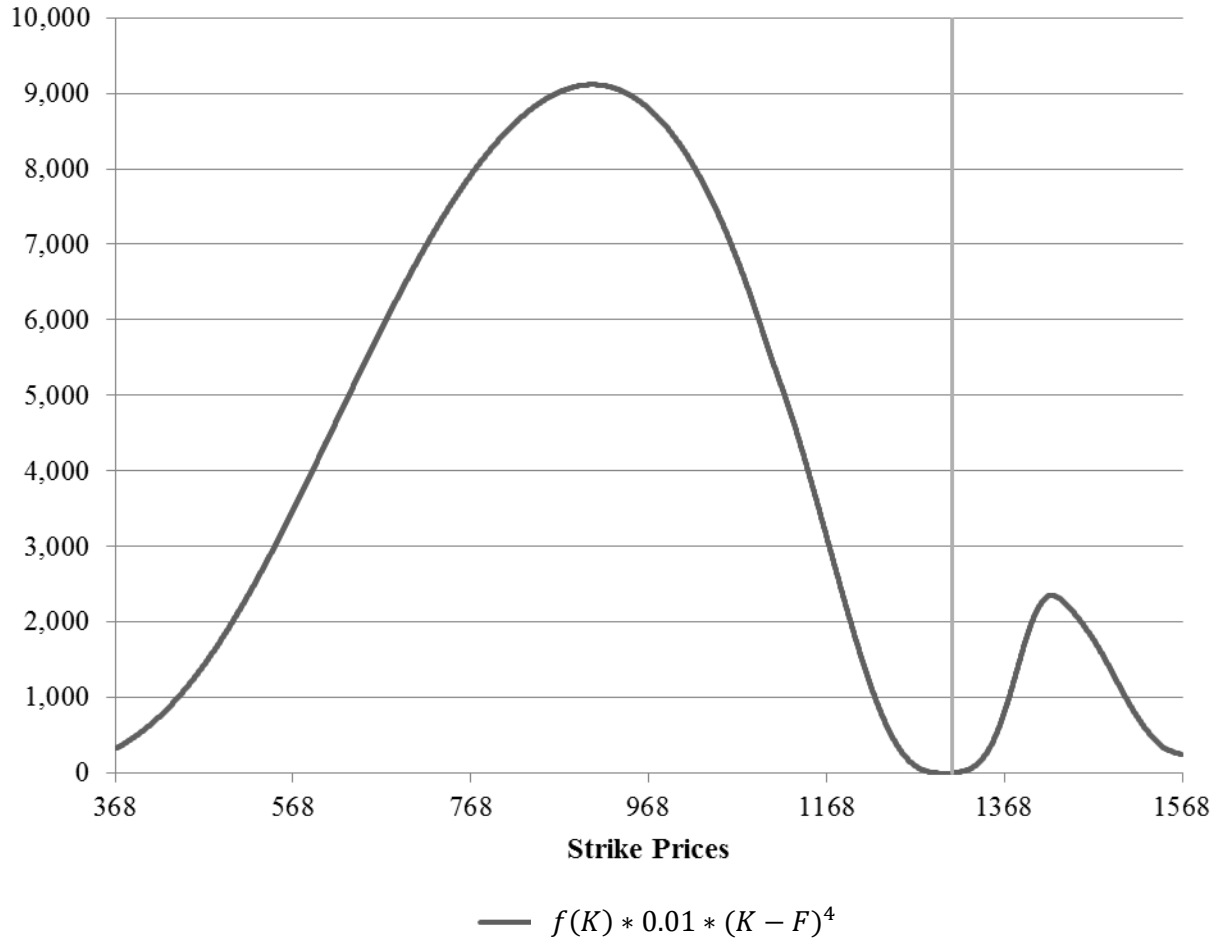

Vita

JORDAN BLAKE ZIMBELMAN

Candidate for the Degree of

DOCTOR OF PHILOSOPHY

Dissertation: THE TAILS OF OPTION-IMPLIED PROBABILITY DISTRIBUTIONS

Major Field: FINANCE / RISK-MANAGEMENT

Education:

Ph.D, Finance, Oklahoma State University, Stillwater, OK, May. 2020
M.B.A, Kansas State University, Manhattan, KS, Aug., 2006
B.S., Computer Science, Kansas State University, Aug., 2004.

Teaching Experience:

2013-2020, Oklahoma Christian University, MBA Adjunct Professor of Finance: Financial Services – Online/On-Site MBA
2018, Wichita State University, Adjunct Professor: Microeconomics
2016-2017, Tabor College – Wichita, Adjunct Professor of Finance Economics I – Online
2013-2015, The University of Oklahoma, Adjunct Professor of Finance: Commercial Banking; Corporate Finance – MBA; Derivative Securities and Markets; Financial and Energy Risk Management; Global Economics – MBA; International Financial Management; Investments; Quantitative Methods & Modeling – MBA
2009-2013, Oklahoma State University - Tulsa, Visiting Faculty: Introduction to Corporate Finance; Financial Markets and Institutions; Financial Management; Real Estate Investment and Finance
2006-2012, Oklahoma State University, Teaching Associate: Introduction to Corporate Finance; Financial Markets and Institutions; Real Estate Investment and Finance
2004-2006, Kansas State University, Teaching Associate: Statistics for Economics and Business I; Statistics for Economics and Business II

Professional Memberships:
Fin. Mngt. Ascn., Eastern Fin. Ascn., South Western Fin. Ascn.