# The implied probability distribution

## – A study of Swedish equity index options

Johan Ring*

MASTER'S THESIS IN FINANCE

STOCKHOLM SCHOOL OF ECONOMICS

## Abstract

This paper examines the option-implied risk-neutral probability distribution for index options traded on the Stockholm market over two decades. The focus is on the method of derivation with the goal of being valuable as a reference for future research as well as being valuable for its results. The key results are: (i) the Swedish implied distribution is similar in shape to the US implied distribution, (ii) the Black–Scholes model doesn't sufficiently explain option prices, (iii) there is significant time-variation in the implied volatility smile and (iv) put-call parity is violated by the volatility smile.

Keywords: Option-implied risk-neutral probability distribution, volatility smile, stock index options

Tutor: Assistant professor Roméo Tédongap

---

* 21418@student.hhs.se

# Contents

# 1. Introduction

Through most of history options were traded between buyers and sellers as over-the-counter instruments rather than over an exchange and as a result were traded in relatively low volume with few observable market prices. Without the benefit of a mathematical model or a liquid options market, prices were based on educated guesses and issuers tended to be able to require a substantial premium over modern models and modern market prices as the value of the insurance was difficult to assess. The first theoretical option pricing model was by Bachelier (1900) that assumed normally distributed stock prices, but his work was well ahead of its time and did not receive much attention until decades later. And the normal distribution he used is technically unappealing as it allows negative asset prices. In the 1960s the lognormal model became popular for asset prices as it's a simple modification of the normal model that prevents negative asset prices. At around the same time CAPM was developed so that discount rates could be estimated in a theoretically appealing way, and now the first mathematical option pricing models could be created. Sprenkle (1964) accordingly propose the option pricing model

$$c = e^{rT} SN\left[\frac{\ln(S/K) + (r + \sigma^2/2)T}{\sigma\sqrt{T}}\right] - (1 - Z)KN\left[\frac{\ln(S/K) + (r - \sigma^2/2)T}{\sigma\sqrt{T}}\right] \tag{1}$$

where $c$ is the value of the call option, $r$ the average rate of growth of the share price, $T$ the time to maturity of the option, $S$ the share price, $N$ the cumulative standard normal distribution function, $K$ the exercise price of the option, $\sigma$ the volatility of the stock returns and $Z$ the degree of risk aversion. The risk aversion constant is crucial in the formula to discount the expected option price but there was no satisfying model developed to estimate this constant or estimate the option discount rate until Black and Scholes realize that it isn't necessary to estimate the option discount rate, directly, to value the option. Black and Scholes (1973) realize that the option price can be calculated in terms of the price of the underlying stock without the need of a discount rate for the option, where both the stock and the option is discounted at the risk-free rate. This result has important practical applications as there is now a satisfying model to price options and in the same year, 1973, the Chicago Board Options Exchange (CBOE) introduces exchange-traded options. Trading volume grew rapidly, and other exchanges were soon also offering option contracts. Optionsmäklarna O.M. introduces exchange-traded options on equities traded on the Stockholm Stock Exchange in 1984 and they introduce equity index options in 1987.

The Black–Scholes model assumes a certain risk-neutral distribution of securities prices, the lognormal distribution, and calculates a theoretical option price from that. It was soon realized that market prices in many cases did not adhere to the stylized model and new research areas developed to explain the empirical results. One direction of research is the development of more advanced models for the underlying asset process, models with additional factors such as stochastic volatility, stochastic interest rates, or stochastic jumps. Another direction of research, that this paper belongs to, derive valuable information from option prices, such as the implied probability distribution of the underlying asset, without explicit assumptions about the stochastic process of the underlying asset that produce the distribution. Knowledge of the risk-neutral probability distribution has many valuable applications. It allows us to price any derivative of the asset with the same time to expiration that the probability distribution is estimated for. For example exotic illiquid options can be priced better from the implied risk-neutral distribution than the assumption of a lognormal distribution. Another application is to obtain implied tail probabilities that are used in for example value at risk estimations that are concerned primarily with the probability of extreme market movements. We can gain valuable forecasting data by learning about the actual probabilities that the risk-neutral probabilities implies. For example an analyst might want to know if the probability of a market decline of more than 5% a month into the future has become higher or lower from one week to another. The risk-neutral distribution can give a better answer to that question than simply the change in the implied volatility, as the shape of the probability distribution might have changed for example. The link between the risk-neutral and the actual risk-averse probability distribution help us understand the economy wide marginal utility function for different levels of wealth which is important in economics.

Deriving the risk-neutral distribution is often considered a complex task and many different methods have been proposed. The purpose of this thesis, besides its research results, is therefore also to be a reference for future research that utilize the implied risk-neutral distribution and make the research area more accessible on a master's thesis level. As far as the author is aware there is no other master's thesis from this department that derives the Swedish risk-neutral probability distribution with a flexible non-parametric method from aggregated data, rather than data from a specific day or a specific point in time. Both the method and the research results are therefore valuable. An additional factor that makes this research interesting is the fact that it studies Swedish equity index options which is a less common research subject than larger options markets such as the US market.

# 2. Theoretical framework

We'll begin with looking at the empirical findings of option pricing. The Black-Scholes model quickly became successful but there was from a beginning an interest in how well it held up for market prices. A minimal prediction of the Black–Scholes model is that all options on the same underlying asset with the same time-to-expiration but with different striking prices should have the same implied volatility as there's only one volatility parameter that governs the underlying stochastic process. As a result, since implied volatilities were first calculated they were also plotted across strike prices for options with the same time to expiration and on the same underlying stock. Most such plots have shown a U-shaped implied volatility pattern; where deep in the money or deep out of the money options tend to have higher implied volatilities, this pattern became known as the volatility smile.

An early study of the differences between market prices and model prices is Rubinstein (1985) whose most conclusive observation is that prices for short-maturity out-of-the-money calls are overpriced relative to other calls. In terms of the implied risk-neutral distribution this would mean that the market places a higher probability on extreme negative prices movements for these options than the lognormal model implies. For options with longer maturities, 71 days and above, there was no marked deviation relative to other calls. Later research would find more pronounced volatiles smiles, also for options with longer maturities. Early explanations of the volatility smile were market imperfections such as transaction costs and bid-ask spreads. As deep out of the money options have very little value transaction costs are probably a higher proportion of their price than more valuable options, which increase their price and relative implied volatility. And since they trade less frequently than near the money options their bid-ask spread tend to be higher so that it may not be possible to profit from their mispricing. However as trading volume increased and index options became popular other explanations beyond market friction has been explored.

The leverage effect produces a volatility smile. As the value of the firm's assets decreases debt levels are often unchanged as companies don’t actively relever with the result that leverage increase and equity volatility increases. This makes in-the-money call options and out-of-the-money put options relatively more valuable. On the other hand, as the asset value increases, the equity volatility decreases. This makes out-of-the-money call and in-the-money put options relatively less valuable. This effect is larger for a firm that is highly levered to begin with. The leverage effect is a rather minor effect however and it can only explain about half of the already relatively flat smiles of individual U.S. stock options as Toft and Prucyk (1997) find. And it would therefore explain even less of the index option smiles that is steeper than the smile for the underlying equities.

Grossman (1988), Gennotte and Leland (1990), Guidolin and Timmermann (2000), and Romer (1993) suggested that information aggregation models can explain a volatility smile. In these models investors learn about the true value of the underlying asset through trading, and prices adjust rapidly once learning takes place which result in a volatility smile. However, these models treat decreases in asset prices as being equally likely as increases in asset which mean that they cannot explain the non-symmetric volatility smile that is often observed that suggests that decreases in asset prices are more likely than increases and can there.

Kelly (1994) suggests that the correlations between stocks increase in down markets where asset prices fall rapidly, thereby reducing the diversification effect of a portfolio. As a result indexes experience higher volatility in down markets which can help explain why index equity options tend to have steeper volatility smiles than the underlying equity options.

Another explanation for the volatility smile could be that investors are more risk averse in down markets than in calm markets. Franke, Stapleton, and Subrahmanyam (1999), Benninga and Mayshar (1997), Grossman and Zhou (1996) and Bates (2001) have modeled such a behavior by introducing undiversifiable background risk or heterogeneous investors with one group exogenously demanding portfolio insurance. However, these models can still not explain the steep volatility skews in the indexes since they only produce rather moderately sloped volatility smiles.

Pindyck (1984), French, Schwert, and Stambaugh (1987), and Campbelland Hentschel (1992) developed models with volatility feedback where the direct effect of prices changes, negative or positive news, are multiplied through the change in the risk premium. As prices fall for example, future cash flows seem more uncertain and the price is depressed further. Since companies are slow to relever the change the equity risk premium is greater than the change in the asset risk premium. However this model is also symmetric for up and down markets and the feedback effect is equally great in both a negative and positive direction.

Another group of studies that are offered by for example by Bates (2000), Bakshi, Cao, and Chen (1997) and Pan (2002) introduce more advanced stochastic processes with additional factors such as stochastic volatility, stochastic interest rates and stochastic jumps. But they have not produced a satisfactory model for the volatility smile.

There is no consensus in the literature about which of these explanations matter the most for the volatility smile. And while some of these models can generate a downward sloping volatility smile they normally introduce the same pattern into the probability distribution of the actual stochastic process, i.e. both the risk-neutral distribution implied from option prices and the actual distribution of returns will be leptokurtic (more peaked) and left skewed (a larger left tail), which introduce a puzzle as the actual distribution from observed returns is not leptokurtic or left skewed and the actual distribution cannot be very leptokurtic or skewed under normal risk aversion assumptions. Many agree that leverage has an effect but there is no agreement for most of the other proposed explanations about whether they are a significant explanation for a volatility smile. Especially the pronounced volatility smile in equity index options is still an open question in terms of its implications for the actual distribution of returns and risk aversion.

## 2.1. Relationship between risk-neutral and actual probabilities

One of the most important contributions of the risk-neutral distribution is what it tells us about the actual probability distribution and the risk aversion function. It’s therefore important that we understand the relationship between risk-neutral probabilities, actual probabilities and risk aversion. It’s important to understand that the risk-neutral distribution and the actual distribution is not the same, and to understand how they may be different from each other under different risk aversion assumptions.

We will illustrate the relationship with a numerical example. Consider a simple economy where a stock can evolve into one of two states, one year into the future. To our disposal we have information about two existing securities, a stock and a bond, and two state prices, $\pi_u$ and $\pi_d$, we therefore have a so called complete market, a market with the same number of states as securities and with linearly independent payoffs. The stock is worth 1 today, it’s worth

$$S_u = 1.2214$$

in the up state and it’s worth

$$S_d = 0.8187$$

in the down state.

The bond is priced at 1 today and is worth 1.1 in either of the two future states. The up state price, $\pi_u$, is what an investor is willing to pay today for the certain payment of 1.00 in the up state and the down state price, $\pi_u$, is what an investor is willing to pay today for the certain payment of 1.00 in the down state. We receive the state prices from the equations system

$$1 = \pi_u 1.2213 + \pi_d 0.8187$$

for the value of the stock and

$$1 = \pi_u 1.1000 + \pi_d 1.1000$$

for the value of the bond. The stock is worth its up state payoff times the state price investors pay for receiving 1.00 in that state, plus its down state payoff valued at the state price investors are paying for receiving 1.00 in that state. By solving the equation system the state prices are

$$\pi_u = 0.6350$$

$$\pi_d = 0.2741$$

The sum of the state prices has to be equal to the price of a bond that pays 1.00 in each state, the value of such a bond is $1/1.1 = 0.9091$. We can now obtain the risk-neutral probabilities that the states occur as

$$P_i = e^{rT}\pi_i$$

$$P_u = 0.6350 \times 1.1 = 0.6985$$

$$P_d = 0.2741 \times 1.1 = 0.3015$$

We can now obtain the call price, $c$, as

$$c = e^{-rT}\sum_i P_i X_i = \sum_i \pi_i X_i$$

$$= \frac{(0.6985 \times 0.2214) + (0.3015 \times 0)}{1.1} = 0.1406$$

In this case we know the actual probabilities to be

$$Q_u = 0.9$$

$$Q_d = 0.1$$

and we can therefore obtain the risk aversion factor, $m = \pi / Q$, from

$$m_u = \frac{\pi_u}{Q_u} = \frac{0.6350}{0.9} = 0.7056$$

$$m_d = \frac{\pi_d}{Q_d} = \frac{0.2741}{0.1} = 2.7410$$

The risk aversion factor, $m$, is the marginal utility of the representative investor in a future state of the economy, it’s also known as the pricing kernel and as the stochastic discount factor. For a risk averse investor the factor is decreasing in wealth as in the example: the poorer the investor the higher the ratio, the wealthier the investor the lower the ratio. When the investor is poor, in the down state, he treats 1.00 received as if it happened with 2.7410 times the likelihood that it really did happen. When the investor is rich, in the up state, he treats 1.00 received as if it happened with 0.7056 times the likelihood that it really did happen.

In practice we don't know the actual probabilities, $Q$, or the pricing kernel, $m$, and they are difficult to estimate since they cannot be estimated directly from securities prices as the risk-neutral probability can be. But estimating the risk-neutral probabilities, $P$, help us estimate the other two factors since the more we know about any of the three factors help us in the estimation of the other factors. If we know two of the factors the third factor can be estimated from the other two.

There are many valuable applications in obtaining the actual risk-averse probability distribution. One of the applications is to obtain tail probabilities that are used in value at risk estimations that are concerned with the probability of extreme market movements. Another application is to obtain forecasts of market movements. There are two difficulties that must be overcome though for these applications to be practical. The first difficulty concern the problem of obtaining option prices at very low moneyness levels that are required to estimate tail probabilities, as these options tend to be infrequently traded. We can overcome this difficulty by aggregating prices over several days or even longer periods to obtain enough samples of low moneyness options to be able to estimate tail probabilities.

The second difficulty is the assumption about risk aversion that concerns most applications involving the derivation of the actual risk-averse probability distribution from the risk-neutral probability distribution. The exceptions when this assumption is not important are when the underlying security is not highly correlated with aggregate wealth such as for many commodities and to a lesser extent for interest rates. In that case risk aversion will be small and the risk-neutral and actual distribution similar regardless of the assumptions about the risk aversion function.

What do we know about the relationship between the risk-neutral and the actual probability distribution? The good news is that there are good reasons to assume that the risk-neutral and actual risk-averse distribution is similar. Rubinstein (1994) and Ziegler (2003) show that if we are willing to assume that the economywide utility function is of standard type, for example a power utility, log utility or negative exponential utility that is common in research, the shapes of the risk-neutral distribution and the actual distribution are similar; the actual distribution simply shifts a little to the right to account for the risk premium, the difference between the mean of the actual distribution and the mean of the risk-neutral distribution (i.e., the risk-free rate). And for small time horizons the shift is relatively small, for example only around 0.5% per month for the equity risk premium.

The bad news is that it's hard to reconcile the leptokurtic and skewed option-implied risk-neutral probability distribution of equity indices with the actual probability distribution implied from historical returns with the same time horizon as the studied options, for example 30 days, that are often close to lognormal rather than skewed or leptokurtic. Jackwerth (2000) find that the implied risk-neutral distribution together with a lognormal actual distribution imply a risk aversion function that differs much from commonly used risk aversion functions, most notably it increases for large portions of wealth close to the starting wealth level (where options are liquid, rather than far away from the starting wealth level where far away from the money options, that tend to be illiquid, are required to estimated probabilities) which violates the basic assumption of risk aversion that it's monotonically decreasing in wealth. Jackwerth propose that it could be a case of mispricing in the market and offer a profitable trading strategy that can be utilized with liquid options. It's however hard to argue that a mispricing continues after it has been explained and attention has been brought to it and alternative risk aversion functions have been proposed. The stochastic jump and stochastic volatility model by Pan (2002) and the assumption that the price of risk is proportional to the volatility or to the intensity of the jump can reconcile the empirical risk neutral and actual probability distribution under certain assumptions. The model of heterogeneous investor beliefs by Ziegler (2003), where the economy consist of two equally large groups, optimists and pessimists, that consistently believe in a distinctly pessimistic probability distribution and a distinctively optimistic probability distribution respectively, can reconcile the empirical risk neutral and actual probability distribution. The problem with the model lies in its fairly onerous behavior assumptions that no obvious equilibrium in the economy would support. It's more likely that investors learn from historical returns and adjust their beliefs. Brown and Jackwerth (2003) propose a risk aversion model that can be increasing in wealth near the starting wealth level by consisting of two vastly different risk aversion functions, one associated with high market volatility and one associated with low market volatility, between which the risk aversion function varies. None of these alternative risk aversion functions have gained much support as of yet and it's still an open question whether the risk-neutral and actual distribution is similar or whether a risk aversion function that's significantly different from common assumptions is represented in market prices. The result of this uncertainty is that analysts and researchers use caution when deriving the actual distribution from the risk-neutral distribution and applications of the risk-neutral distribution in estimation of the actual distribution, such as for forecasting and estimating tail probabilities, will be held back until there exist a clearer understanding of these questions.

## 2.2. The actual distribution

Studies of the risk-neutral distribution implied from equity index options often find a leptokurtic and left skewed distribution and under reasonable pricing kernel assumptions we can assume that the actual probability distribution also exhibits those traits and the lognormal assumption should therefore not be used for the actual distribution either. It's also common to observe leptokurtic distributions from sample returns. But that stems at least partly from a well-known statistical phenomenon that a nonstationary sample, with time-varying parameters, will introduce leptokurtosis even if the true underlying distribution is lognormal. This is illustrated in figure 1 where continuously compounded returns over the sample period express a high level of kurtosis for daily returns, and effect that subside when fewer observations are used in the weekly and monthly returns series. To further understand this effect a simulated sample with a time-varying parameter, the standard deviation of returns, is overlaid the sampled returns. This simulation draws random returns with a random standard deviation that try to simulate the variation in standard deviation that the underlying stock process experience. It does this by varying the standard deviation randomly for each observation according to the standard deviation of the option-implied volatility of the underlying index. The simulated sample returns are

$$\ln price_t - \ln price_{t-1} = \mu + \epsilon_t$$

$$\epsilon_t = \mu^* + \epsilon^*$$

where $\mu$ is the mean return of the sample, $\epsilon_t$ a lognormally distributed disturbance, $\mu^*$ the mean implied volatility of the sample, $\epsilon^*$ a lognormally distributed disturbance with the standard deviation of the implied volatility of the sample. This simulated sample, which is also drawn in figure 1, exhibit leptokurtosis too, although not to the extent that the actual sample does. The fact that the actual sample is more leptokurtic than the simulated sample may imply that the true underlying distribution for the sample returns is leptokurtic which magnifies the total leptokurtosis. But the reason may also be that the true underlying distribution is nonstationary in a more complicated manner, with a larger variance in the standard deviation or additional variance in higher moments that enforce the total leptokurtosis exhibited.

Because of this behavior of a non-stationary time series we cannot say if the actual probability distribution of the stock index is lognormal or not, based on sample returns. There's no clear way to differ between the leptokurtosis introduced by the nonstationarity and the leptokurtosis that may exist in the underlying return process. Historical returns still carry some weight however for estimating the actual distribution, particularly since sample returns longer return series, such as a weeks or months returns, are less affected by the described nonstationarity effect and are fairly lognormal. Part of the reasoning behind the pricing kernel puzzle is that it's hard to believe that the actual probability distribution is as leptokurtic as risk-neutral probabilities seems to imply when sample returns are fairly lognormal for longer time periods than days. It's also important to conclude that the described nonstationarity effect in sample returns limits its use in estimating the actual probability distributions and that much can be gained by combining the analysis of historical sample returns with the analysis of option-implied risk-neutral distributions and the pricing kernel.

**Figure 1. Actual probability distribution.**

*Note:* The actual probability sampled from OMXS30 returns over the same sample period as the options. The histogram is overlaid with a kernel distribution. A lognormal with the same mean and standard deviation is overlaid as a reference.

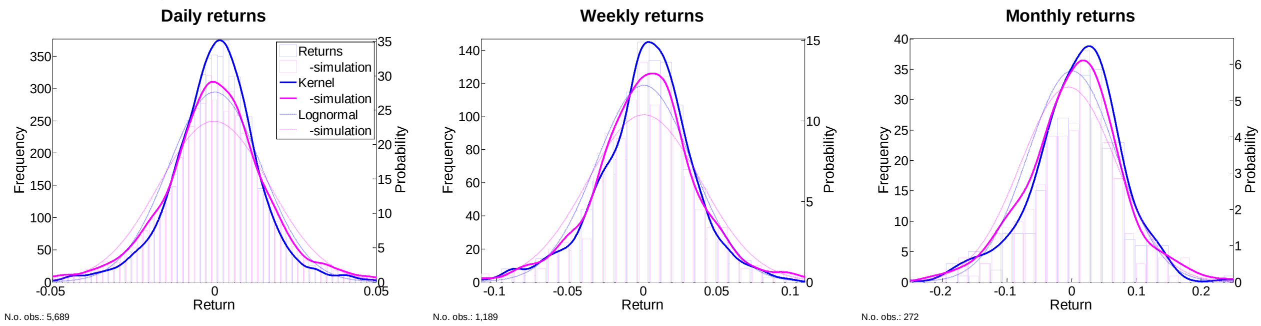

# 3. Methodology

Because of the scarcity of options prices in the 1970s and 1980s not much empirical research was conducted on options prices at the time so the history of recovering implied risk-neutral probabilities is short. One of the first papers on the subject was Banz and Miller (1978) that backed out risk-neutral probabilities out of a set of option prices for the purpose of capital budgeting, but they simply used hypothetical Black–Scholes-based option prices and as a result only backed out the associated lognormal probabilities. Longstaff (1990) backed out implied probabilities by distributing them uniformly between any two adjacent strike prices, but strike prices were too far apart to give numerically stable results. Later methods interpolate between strike prices and extrapolate outside observed strike prices to provide satisfactory tail probabilities and Rubinstein (1994) produce the first widely cited implied risk-neutral distribution for US equity index options with such a method.

The methods of recovering implied risk-neutral probabilities can be divided into two groups, parametric approaches and non-parametric approaches. Both groups use an optimization method to find the variables that best fit observed option prices. The parametric approach fit the data to a parametric probability distribution. The simplest possible parametric approach is to fit a simple parametric distribution to the data, such as a lognormal distribution, but the reason that we observe a volatility smile is that the lognormal distribution is not flexible enough, and we must therefore use a more flexible parametric approach. The parametric approaches can be divided into expansion methods, generalized distribution methods and mixture methods. For example, Jarrow and Rudd (1982) utilize an Edgeworth expansion with two expansion terms and a normal distribution as the approximating distribution, Sherrick, Irwin, and Forster (1992) utilize the Burr distribution, Ritchey (1990) utilize a mixture of normal distributions. Expansion methods start with a simple analytical distribution such as a normal or lognormal and apply correction terms to it. Generalized distribution methods use distribution functions with more parameters than just mean and standard deviation, they often add skewness and kurtosis parameters. Skewness allow the left tail to be fatter than the right tail to allow a higher likelihood of extreme negative returns than extreme positive returns. Kurtosis allow the distribution to be more peaked in the center and have fatter tails, i.e. give a higher probability to small and large price changes at the expense of a smaller probability for intermediate size price changes. Mixture methods use a weighted average of several simple distributions, for example a weighted average of two or three lognormal distributions with individual parameters. Incorporating more models in the mixture give added flexibility but with the downside of increasing the number of parameters that has to be estimated. For example a weighted average of three lognormal distributions require eight parameters to be optimized, the mean and standard deviation of the three distributions and two weight parameters, the third weight parameter comes from the constraint that the resulting model must integrate to 1. A common practical problem when optimizing mixture models is that it’s relatively easy to produce spikes in form of one of the distribution in the mix receiving a very small standard deviation and producing a spike in the resulting distribution that is certainly not a property of the true distribution.

The nonparametric methods fit the data without utilizing a parametric probability distribution with the benefit that they can be fully flexible and independent from a parametric model. The probability distribution is rarely fitted directly because of numerical problems of constraining probabilities to be positive, to sum to 1 and exhibit some smoothness. It's more common to either fit a function of call prices over strike prices and obtain the probabilities from those call prices or to fit an implied probability function over strike prices and obtain call prices and probabilities from that. The drawback of fitting call options prices is that they vary much across strike prices, deep in the money calls are valued as high as the underlying asset itself while deep out of the money calls are valued close to zero. We are therefore implicitly putting more weight on the in the money calls than on the out of the money calls. And in addition to that constraints must be applied so that the arbitrage bounds are not violated. A better method is therefore to fit a function of implied volatiles across trike prices and from that proceed to calculate option prices and probabilities. The advantage of this method is that volatilities are much more similar in magnitude across strike prices than option prices are and as long as the fitted volatility smile is reasonably smooth these methods produce positive and arbitrage free risk-neutral probabilities.

The nonparametric methods can be divided into three groups, maximum entropy, kernel, and curve-fitting methods. Maximum entropy methods find risk-neutral probability distributions that fit the options data and that presume the least information relative to a prior probability distribution. As a prior distribution, a lognormal distribution is often chosen, and the resulting posterior distribution is subject to the constraints that the risk-neutral probabilities are positive, sum to 1 and fit the option prices (with the underlying asset treated as a zero-strike option). The main problem with this group of models is that they use a logarithm of the posterior distribution in their optimization method. Probabilities can become tiny and the logarithm of such small values goes rapidly to large negative values that dominate the maximization routine and drive the results. This method also requires the use of nonlinear optimization routines which are more complicated to use than linear optimization routines.

Kernel methods are related to nonlinear regressions in that they are localized for each data point and don't specify the linear form of a standard regression. At a given data point the observation of either the call price or implied volatility is the weighted average of the implied volatilities adjacent to it. The advantage of the model is that the tradeoff between smoothness and fit can be controlled through the bandwidth parameter. The downside is that the model is data intensive and cannot smooth over gaps in the data very well.

Curve-fitting methods try to fit the implied volatility smile with some flexible function. The most typical criteria for the fit are sums of the squared differences in modeled and observed volatilities, the squared differences in modeled and observed option prices, or the squares of 1 minus the ratio of option prices. Typical functions for curve fitting are polynomials in the strike price $[\sigma(K) = \alpha_0 + \alpha_1 K + \alpha_2 K^2 + \ldots + \alpha_n K^n]$, but they can exhibit oscillatory behavior if they involve higher-order terms.

Therefore splines are a better choice here; splines piece together polynomial segments at so-called knots by matching levels and derivatives at the knots. The choice of the location of those knots is somewhat of an art, too many knots cause overfitting of the data and too few knots prevent the observed volatilities from being matched. Splines tend to be smoother and do not exhibit the oscillations that polynomials are prone to, but they need to be of an order higher than 3 for the probability distribution to turn out to be smooth. Alternatively a smooth function can be fitted explicitly by penalizing jaggedness. One way is to add into the objective function a term based on the sum or integral of squared second derivatives of the function. The more curvature the function has, the larger this term will be. A tradeoff parameter governs how much weight is given to fit versus smoothness. Its value is typically based on trial and error.

The easiest and most stable methods, as we will explain, tend to be curve-fitting methods that fit the implied volatility smile. The limiting case of a flat implied volatility function simply gives the Black–Scholes model. These methods tend to be fast and some of them require only a single calculation. As long as the implied volatility smile is reasonably smooth, the associated risk-neutral probability distribution will be strictly positive. But this condition is not guaranteed and it should be checked separately. The choice of method matter less the more option price observations we have and if the choice of method doesn't change the results much we should use a fast and stable method. For this paper the choice of method fall on the fast and stable curve fitting method proposed in Jackwerth (2004). It has several advantages, the tradeoff between fit and smoothness is easily controlled with an external parameter, as the tradeoff parameter goes to zero the volatility smile becomes flat, the calculations are simple and can be done in a spreadsheet, it doesn't use complicated iterative optimization methods, the variance minimization is done by the multiple regression formula, the resulting smooth implied volatilities are obtained as the slope coefficients in a multiple regression with suppressed constant. The method consists of creating a matrix with the strike price intervals on both axes, to which the observed implied volatilities are associated. All intervals don't need to have an observed implied volatility; a volatility will be interpolated if an observation is missing for an interval. The granularity of the matrix can be adjusted freely after how much detail we want and our data permits.

An example with three option prices will illustrate the method. Assume we have three call option prices $c_1 = 0.181$, $c_2 = 0.104$, $c_3 = 0.04$, with different strike prices $K_1 = 0.9$, $K_2 = 1.0$, $K_3 = 1.1$, and implied volatility $\sigma_1 = 0.25$, $\sigma_2 = 0.2$, $\sigma_3 = 0.15$, but that are otherwise identical with an interest rate of 5%, with no dividend yield, a time to expiration of 1 year and the underlying asset price is 1. To keep the example manageable we limit ourselves to 11 future index levels that are 0.1 apart from each other on the range 0.5 to 1.5. If the resulting probability distribution is too coarse or does not tend to 0 in the tails, we can introduce half-steps (0.55, 0.65, etc.) or add values to the ends (0.3, 0.4, 1.6, 1.7, etc.). Trial and error is required to find a proper trade-off parameter, $\lambda$. If the distribution is too jagged $\lambda$ is reduced, if the fitted volatilities do not match the observed volatilities sufficiently $\lambda$ is increased. The next step is to generate a system of equations, one for each future index level, with the entries 1, –4, 6, –4, and 1 around the diagonal and –3 and 3 as the top two values of the first column and the last two values of the last column.

The information from the observed options is then added to their corresponding strike price group 0.9, 1 and 1.1. The required multiplier is

$$\text{trade– off parameter} \times \left[\frac{\text{number of index levels}}{\text{number of options} \times (\text{index spacing }^4)}\right] =$$

$$= \lambda \left[\frac{J+1}{I(\Delta^4)}\right] = 0.01 \left[\frac{11}{3 \times (0.1^4)}\right] \approx 367$$

which is added to the matrix for the strike prices where there exist and observation in the following way:

$$\begin{array}{cccccccccccc}
3\sigma_1 & -4\sigma_2 & +1\sigma_3 & +0\sigma_4 & +0\sigma_5 & +0\sigma_6 & +0\sigma_7 & +0\sigma_8 & 0\sigma_9 & +0\sigma_{10} & +0\sigma_{11} & = 0 \\
-3\sigma_1 & +6\sigma_2 & -4\sigma_3 & 1\sigma_4 & +0\sigma_5 & +0\sigma_6 & +0\sigma_7 & +0\sigma_8 & 0\sigma_9 & +0\sigma_{10} & +0\sigma_{11} & = 0 \\
1\sigma_1 & -4\sigma_2 & +6\sigma_3 & -4\sigma_4 & +1\sigma_5 & +0\sigma_6 & +0\sigma_7 & +0\sigma_8 & 0\sigma_9 & +0\sigma_{10} & +0\sigma_{11} & = 0 \\
0\sigma_1 & +1\sigma_2 & -4\sigma_3 & +6\sigma_4 & -4\sigma_5 & +1\sigma_6 & +0\sigma_7 & +0\sigma_8 & 0\sigma_9 & +0\sigma_{10} & +0\sigma_{11} & = 0 \\
0\sigma_1 & +0\sigma_2 & +1\sigma_3 & -4\sigma_4 & +(6+367)\sigma_5 & -4\sigma_6 & +1\sigma_7 & +0\sigma_8 & 0\sigma_9 & +0\sigma_{10} & +0\sigma_{11} & = 0.25 \times 367 \\
0\sigma_1 & +0\sigma_2 & +0\sigma_3 & +1\sigma_4 & -4\sigma_5 & +(6+367)\sigma_6 & -4\sigma_7 & +1\sigma_8 & 0\sigma_9 & +0\sigma_{10} & +0\sigma_{11} & = 0.20 \times 367 \\
0\sigma_1 & +0\sigma_2 & +0\sigma_3 & +0\sigma_4 & +1\sigma_5 & -4\sigma_6 & +(6+367)\sigma_7 & -4\sigma_8 & +1\sigma_9 & +0\sigma_{10} & +0\sigma_{11} & = 0.15 \times 367 \\
0\sigma_1 & +0\sigma_2 & +0\sigma_3 & +0\sigma_4 & +0\sigma_5 & +1\sigma_6 & -4\sigma_7 & +6\sigma_8 & -4\sigma_9 & +1\sigma_{10} & +0\sigma_{11} & = 0 \\
0\sigma_1 & +0\sigma_2 & +0\sigma_3 & +0\sigma_4 & +0\sigma_5 & +0\sigma_6 & +1\sigma_7 & -4\sigma_8 & +6\sigma_9 & -4\sigma_{10} & +1\sigma_{11} & = 0 \\
0\sigma_1 & +0\sigma_2 & +0\sigma_3 & +0\sigma_4 & +0\sigma_5 & +0\sigma_6 & +0\sigma_7 & +1\sigma_8 & -4\sigma_9 & +6\sigma_{10} & -3\sigma_{11} & = 0 \\
0\sigma_1 & +0\sigma_2 & +0\sigma_3 & +0\sigma_4 & +0\sigma_5 & +0\sigma_6 & +0\sigma_7 & +0\sigma_8 & +1\sigma_9 & -4\sigma_{10} & +3\sigma_{11} & = 0
\end{array}$$

The system of equations can be solved with the multiple regression formula where $\sigma_1$ to $\sigma_{11}$ are the slope coefficients and the intercept is suppressed. Once we have obtained the implied volatilities for all strike prices the implied volatility calculation is used in reverse to obtain a call price for each implied volatility, i.e. the Black–Scholes model is used to obtain a call price for each volatility. We then use an important result by Breeden and Litzenberger (1978) that the second derivative of call prices over strike prices are the state prices and if we use forward call prices it's the risk-neutral implied probability

$$\frac{\partial^2 c}{\partial K^2} = \pi_{S=K} \Leftrightarrow \frac{\partial^2 c e^{rT}}{\partial K^2} = P_{S=K} \tag{2}$$

which can be expanded to

$$P_{S=K} = e^{rT} \frac{c_1(K_1 - \Delta K) - 2c_2(K_2) + 2c_3(K_3 - \Delta K)}{\Delta K^2} \tag{3}$$

for three adjacent strike prices, $K_i$, and their call prices, $c_i$, where $P_{S=K}$ is the probability that the underlying asset is worth the center strike price, $K_2$, at expiration. It's therefore possible to fit a function of call prices across strike prices and use the Breeden and Litzenberger result to obtain the implied probability distribution. In this curve-fitting method the distribution doesn't automatically integrate to 1 but needs to be adjusted according to

$$P_i = P_i \times \left[\int_0^\infty P \times K \, dK\right]^{-1}$$

to integrate to 1. The resulting implied distribution is shown in figure 2 for three different values of the tradeoff parameter that result in three different levels of fit to the observed volatilities. A lognormal with the same mean and standard deviation is overlaid the implied distribution as a reference. In the limit, where the fit is loosest, the fitted smile turns into a flat line with a value of the average of the observed volatilities.

**Figure 2. Implied distribution example; smoothing.**

*Note:* Illustrate the effect on different settings for the tradeoff parameter, $\lambda$. As $\lambda$ is decreased the implied volatility tends towards a flat line and the distribution towards the lognormal.

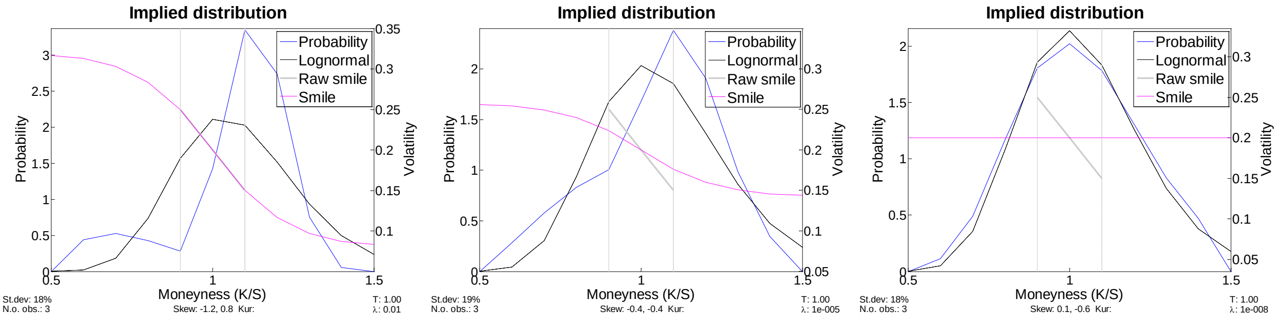


## 3.1. Data aggregation

In the previous example the option prices span a fairly narrow moneyness range and we do not obtain any information about the implied probability outside the available strike prices. That part of the distribution is just an extrapolation from the data. The same thing is often true about actual trading data too; as option price data from even the most liquid option markets at a certain point in time, such as the close of a trading day, are relative sparse is terms of their moneyness span since most trading occur for near the money options. And the prices can oscillate much from day to day especially for the illiquid strike prices. It's therefore possible to improve the implied distribution estimate by aggregating option prices over time. By doing this we accumulate option prices for deep in the money and out of the money options that are traded less often, and even out the day to day noise in the option prices. The downside is of course that we don't obtain any information about how the implied distribution has varied during the sample period and we may risk evening out important anomalies, for example. But overall data aggregation is an important benefit when analyzing option markets, especially less liquid markets.

To improve the accuracy and comparability of the results the implied volatility and moneyness are defined in relative terms. The volatility smile shift up and down as the overall market uncertainty varies up and down. These variations are fairly large as can be seen from the implied volatility over the sample period, from figure 7. If we averaged the unadjusted volatility the result would be that days with high overall volatility receive a larger weight which is undesirable as only relative volatiles, the differences between the volatility for different strike prices, are of interest. A relative implied volatility is therefore calculated as

$$\sigma_i = \frac{\sigma_i}{\sigma} \tag{4}$$

where $\sigma_i$ is the volatility for the individual option and $\sigma$ is the best estimate of the implied volatility for a certain day, taken from the most liquid options. It's calculated as the average of the implied volatility of the two most liquid nearest month or next to nearest month options for every trading day at the close of trading. As the options move towards expiration the volatility measure ranges from the 45 day volatility, the uncertainty of the underlying asset 45 days into the future, to the 15 day volatility, after which the next month's options are considered so that the volatility term is increased to 45 days again. (The exact range is alternating between a range of 42 days to 15 days and a range of 49 days to 15 days because of the schedule by which new options are issued.)

The moneyness measure is also crucial for the volatility smile and the implied distribution. For a single point in time it's common to expressed moneyness as the striking price to underlying asset price ratio, $K/S$, which has a simple interpretation and is sufficient for a single point in time. However, it doesn't allow us to aggregate data easily since this moneyness ratio is expressed in absolute terms that don't take into consideration the time to expiration. For example, close to expiration, let's say 30 days to expiration, a moneyness ratio of 0.9, i.e. a strike price 10% below the stock price, is a considerable distance between the strike price and underlying asset price in terms of the 30 day standard deviation, a 10% change in the index in a month is a roughly a two standard deviation event, but for a long term option it's not a big distance, other things equal, as a 10 percent change in a year is a $\sqrt{12}$ times smaller standard deviation event. We therefore introduce a moneyness measure that takes into account the time to expiration, or stated differently makes moneyness independent of the option's time to expiration in its interpretation, and we will use this moneyness measure for all samples that aggregate data from different times. This moneyness measure is

$$\xi = \frac{\ln(K/F)}{\sigma\sqrt{T}} \tag{5}$$

where $F$ is the corresponding futures price for the underlying asset, $K$ the strike price, $T$ the time to expiration as a fraction of a year and $\sigma$ the annualized implied volatility. Or in those cases where a futures price is missing $\xi = \frac{\ln(K/S)-(r-q)T}{\sigma\sqrt{T}}$ where $S$ is the underlying index value, $r$ the risk-free interest rates, $q$ the dividend yield. If we assume that $F = Se^{(r-q)T}$ this is equal to formula (5). At the point in time the variable is measured the implied volatility, $\sigma$, in the calculation is not the implied volatility for the option that has the strike price $K$, it's the same implied volatility for all strike prices for that point in time, the volatility consider to be the best measure of implied volatility for any given day. The reason for using only one implied volatility at a certain point in time is to scale moneyness by the same volatility for all strike prices. The underlying distribution that we are deriving has only one standard deviation at a certain point in time and we should therefore scale moenyness by only one standard deviation at a certain point in time.

If we examine the formula we can see that $\xi$ is the number of standard deviations between the strike price and expected underlying asset value at expiration with the risk-neutral expectations that the Black-Scholes model (adjusted for a dividend yield $q$) assume. The variable therefore has an intuitive interpretation as moneyness expressed as the number of standard deviations from the arithmetic mean.

To estimate the probability distribution from aggregated data implied volatility and moneyness in absolute terms is required and these are obtained by the transformation

$$\sigma_i^* = \sigma_i \times \mu_\sigma \tag{6}$$

where $\sigma_i$ is the relative volatility of an individual option and $\mu_\sigma$ is the mean implied volatility for the period considered. Moneyness in absolute terms is obtained from the transformation

$$[K/S]_i^* = \xi_i \times \left(\mu_{|\xi|} \times \mu_{\sigma_t\sqrt{T_i}}\right) + 1 \qquad (7)$$

where $\xi_i$ is the relative moneyness for an individual option, $\mu_{|\xi|}$ the mean of the moneyness in the sample in absolute terms and $\mu_{\sigma_t\sqrt{T_i}}$ the mean of the de-annualized implied volatility in the sample. This transformation gives a $[K/S]_i^*$ that is on average close to the untransformed moneyness in terms of $K/S$, that can be called $[K/S]_i$ with equivalent notation, of individual options. The implied distribution is not sensitive to small changes in these transformation terms, but if $\mu_{\sigma_t\sqrt{0.12}}$ is changed to $\mu_{\sigma_t\sqrt{1}}$ for example, so that the term becomes three times larger, more $[K/S]_i^*$ values are negative and unfeasible and the considered volatility smile is flatter as a standard deviation of moneyness is larger in absolute terms, and the transformed $[K/S]_i^*$ values are on average larger than the $[K/S]_i$ values which would introduce a measurement bias.

The choice of data for the study is index options rather than options on individual equities. This is natural as we are looking for a proxy for the economy-wide implied probability that has important implications for the economy, rather than for a small subset of the economy such as a single equity. There are also technical advantages with studying equity index options. There exist European style equity index options so that we don’t have to take into account the possibility of an early exercise which can complicate the valuation of the option. Since dividends are spread out over time early exercise is less of an issue with American style equity index options compared to American style equity options. The options can easily be hedged using the index future with the same time to maturity, the index payout can be reliably estimated or inferred from index futures, unlike for bond option the underlying asset can be assumed to follow a risk-neutral lognormal process, unlike currency exchange rates the equity index does not have an obvious non-competitive trader in its market (i.e., the government), and the equity index is less likely to experience jumps than probably any of its component equities and most other underlying assets such as commodities, currencies and bonds.

## 3.2. The lognormal distribution

It will help us to be familiar with the lognormal distribution when we discuss option pricing and deviations from the Black–Scholes model. Figure 6 illustrate one of the defining traits of the lognormal; its skew increase with standard deviation according to its skew definition

$$skew = \left(e^{\sigma^2} + 2\right)\sqrt{e^{\sigma^2} - 1}$$

where $\sigma$ is the standard deviation of its natural logarithm and as a result exhibits considerable skew even at moderate standard deviations. As its skew increase its main body and mode moves away from the arithmetic mean, as its arithmetic mean and mode is defined as

$$mean = e^{\mu+\sigma^2/2}$$

$$mode = e^{\mu-\sigma^2}$$

As a reference figure 6 compare the lognormal distribution to symmetric non-skewed distributions centered on the same arithmetic mean and with the same standard deviation. In the figure, when the risk is high, at 30%, there’s considerably less cumulative probability on the left tale beyond two standard deviations than on the right tale which will affect model pricing. The skew that the lognormal distribution exhibits may be unfitting compared to the true distribution, it’s possible that actual risk-neutral distribution is not skewed, and a distribution closer to the actual distribution may be a symmetric distribution such as the normal distribution whose density function is gradually clamped to zero to avoid negative values. As can be seen from the figure that would require a very small adjustment to the normal distribution even at a high standard deviation like 30% since only around 0.04% of cumulative probability is in the negative value region in that distribution. The downside is of course that the analytical formula for this adjusted density function would be very complex or difficult to achieve at all, so that it’s easier to use the lognormal distribution despite its unattractive skew.

The skew may be a poor representation of investor’s expectations as they may instead place equal weight on extreme negative and positive returns, or possibly more weight on extreme negative returns as some theories suggest. If so, we would expect this anomaly in the lognormal distribution to lead to far out of the money put options to be priced with a larger implied volatility premium than far out of the money call options, as the shaded area in the lognormal distribution in the left tail below –2 standard deviations is smaller than the shaded area in right tail above 2 standard deviations, especially for the most skewed lognormal. This is also what the volatility surface in figure 10 show us, at a moneyness of –2 puts are priced higher than calls at a moneyness of 2. It would take additional analysis to determine if the skew is an important reason for that, but it’s good to keep in mind that the skew may result in mispricing.

We can also mention that the skew is not always an important factor in option pricing, at small standard deviations the lognormal is hardly skewed at all. And since most traded options have a short time to expiration their un-annualized implied volatility is often below 10%, at which point the lognormal is not very skewed, the skew is for many options not an important factor in the valuation. But for long term options the skew of the lognormal can become meaningful so that their implied distribution may require a larger adjustment to the lognormal.

## 3.3. The option pricing formula

The original Black–Scholes (1973) formula assumes that the underlying asset doesn’t pay any dividend. Since the average annual dividend yield in our sample is 2.2% we would introduce a systematic error if we ignored dividends. Perhaps a small error, but still an error that would better be avoided. We are therefore left with the choice between the Black-Scholes formula adjusted for a dividend paying asset and the futures option model called Black 76 that use the index future as the underlying asset price.

The Black–Scholes formula for the price of a European style option on a non-dividend-paying stock is

$$c = e^{-rT}[Se^{rT}N(d_1) - KN(d_2)] \tag{8}$$

$$p = e^{-rT}[KN(-d_2) - Se^{rT}N(-d_1)] \tag{9}$$

where

$$d_1 = \frac{\ln(S/K) + (r + \sigma^2/2)T}{\sigma\sqrt{T}}$$

$$d_2 = \frac{\ln(S/K) + (r - \sigma^2/2)T}{\sigma\sqrt{T}} = d_1 - \sigma\sqrt{T}$$

$c$ is the call price, $p$ is the put price, $S$ is the price of the underlying asset, $K$ is the strike price, $\sigma$ is the standard deviation of the underlying asset, $N(x)$ is the standardized normal density function, $r$ is the continuously compounded interest rate, $SN(d_1)e^{rT}$ is the expected value of a variable equal to $S_T$ if $S_T > K$ and zero otherwise, $N(d_2)$ is the risk-neutrla probability that the option will be exercised and $KN(d_2)$ is the strike price times the probability that the strike price will be paid.

Merton (1973) provides an adjustment for an underlying asset that pays a dividend yield. By replacing $S$ by $Se^{qT}$in the Black–Scholes formula the arithmetic mean in the formula becomes $Se^{(r-q)T}$ instead of $Se^{rT}$, and we obtain the price of a European option on a stock providing a dividend yield at rate $q$ as

$$c = e^{-rT}\left[Se^{(r-q)T}N(d_1) - KN(d_2)\right] \quad (10$$

$$p = e^{-rT}\left[KN(-d_2) - Se^{(r-q)T}N(-d_1)\right] \quad (11$$

where

$$d_1 = \frac{\ln(S/K) + (r - q + \sigma^2/2)T}{\sigma\sqrt{T}}$$

$$d_2 = \frac{\ln(S/K) + (r - q - \sigma^2/2)T}{\sigma\sqrt{T}} = d_1 - \sigma\sqrt{T}$$

Black (1976) show that the same formula can be used to value futures options. This model, that's sometimes called Black 76, replaces the arithmetic mean with the $F$ futures price so that we obtain

$$c = e^{-rT}[FN(d_1) - KN(d_2)] \tag{12}$$

$$p = e^{-rT}[KN(-d_2) - FN(-d_1)] \tag{13}$$

where

$$d_1 = \frac{\ln(F/K) + (\sigma^2/2)T}{\sigma\sqrt{T}}$$

$$d_2 = \frac{\ln(F/K) - (\sigma^2/2)T}{\sigma\sqrt{T}} = d_1 - \sigma\sqrt{T}$$

Since the futures price is given by $F = Se^{(r-q)T}$ under no-arbitrage conditions this formula is equivalent to the previous formulas.

The futures market is often very liquid so that $F$ cannot stray from its no-arbitrage price for a long time or to a large extent so that this formula should give similar results as the Black-Scholes formula for a dividend paying stock. But there's a possibility that temporary mispricings in index futures may spill over to option pricing since option traders use index futures for their hedging activity. There's also an uncertainty as to which risk-free interest rate is the correct rate. For example in USD markets there's a documented LIBOR spread over the Treasury risk-free rate and the opinion that LIBOR should be used as the correct interest rate for option formulas. Hull (2003) chapter 20 shows that the historical LIBOR to Treasury bills spread has been approximately 0.43%-points, a difference that is primarily thought to come from differences in taxation and liquidity. In addition to that the correct dividend yield is also an uncertainty in the Black-Scholes formula.

For these reasons most research, including this paper, use the Black 76 futures option formula when a futures price with the same time to maturity is available, which it mostly is. When a futures price is not available the STIBOR SEK rate is assumed to be the correct interest rate to use for the implied volatility calculation. The STIBOR rates are available for different maturities up to one year. The STIBOR rate is selected with the same time to maturity as the option through interpolation over the STIBOR rate curve. When there's no futures price available we also need the value of the underlying OMXS30 index. The dividend yield is obtained from another Swedish stock index as no dividend yield is available for the OMXS30. This dividend yield is assumed to be a good proxy for the dividend yield on the OMXS30 index.

## 3.4. Hypotheses

We can now formulate our hypotheses. Much of the discussion about the implied risk-neutral and actual probability distribution is of a qualitative nature and many important hypotheses are not easily answered by quantitative statistical tests. Previous research of the volatility smile in most other markets than the Swedish market, for example the US market, show a pronounced volatility smile for equity index options. It's therefore trivial to assume that the Swedish market also exhibits a volatility smile. Less trivial and more interesting is the question about relative differences between different equity markets. As US option data is available for this paper we will compare the US and Swedish option-implied probability distribution. As the US and Swedish market is similar in many respects, except the lower liquidity in the Swedish market, the hypothesize is that the US and Swedish implied probability distribution is similar. The hypothesis will be answered with a qualitative discussion of the derived probabilities. The volatility smile in this hypothesis refers to, as always, the function of volatilities over strike prices. With similar properties (of the probability distribution) we refer to similar visually, in terms of shape and form, and similar in terms of quantitative higher moments of the distribution, for example skew and kurtosis.

> **Hypothesis 1**: OMXS30 and S&P 500 implied risk-neutral probabilities have similar properties.

To investigate the overall deviation from lognormality of the implied distribution we can apply a quantitative test of the size of the deviations from Black–Scholes options prices. As studies from other market show significant deviations from Black–Scholes options prices the hypothesis is that the Swedish data also show a significant deviation. A nonparametric test of the lower bound of the pricing errors will be used to determine the answer to hypothesis 2.

**Hypothesis 2**: The Black–Scholes model cannot sufficiently explain option prices.

It's also of interest to understand if the volatility smile varies over time. With variations over time we mean variations in the shape of the volatility smile, in terms of steepness or levelness or other differences, between different points in time. The hypothesis is that it does vary over time as investors reevaluate factors such as the actual probabilities, their understanding of options prices and as the options market efficiency varies. This hypothesis will be tested in a parametric test.

**Hypothesis 3**: The volatility smile is time-varying.

Another important relationship that can be tested in terms of the volatility smile is put-call parity. Put-call parity, defined as

$$c = p + S - Ke^{rT} \tag{14}$$

$$p = c - S + Ke^{rT} \tag{15}$$

where $c$ is the call price and $p$ is the put price tell us that the implied volatility for a put option with the same strike price, time to expiration and other parameters must be the same as for the call option. Otherwise there is an opportunity to set up a risk free portfolio to benefit from the fact that the put and the call are not valued in parity. It's therefore of interest to examine the difference between the volatility implied by put options and the volatility implied by call options. Since we expect market imperfections to be significant, primarily because of illiquidity, we hypothesize that put-call parity doesn't hold and that the volatility smile differs between puts and calls. This hypothesis will be tested in a parametric test.

**Hypothesis 4**: The call option volatility smile is different from the put option volatility smile.

# 4. Data

The sample consists of daily closing prices of OMXS30 index options between 1987 and 2005 provided by OMX. Counting all options that had a bid or ask quoted at the close of trading there are around 400 000 option prices in the sample. Only counting options that also traded during the day there are around 180 000 options prices. This gives on average 35 option prices per day of both calls and puts and on average 28 options per day that has a unique combination of strike price and time to expiration. The reason for the relatively high frequency of such unique combinations compared to the total number of observations is that there's not a perfect overlap between calls and puts which stems from the fact that in the money options are more likely to be traded than out of the money options so that there's a larger proportion of calls with moneyness, $\xi$, below zero and a larger proportion of puts with moneyness above zero. About 65% of traded calls are in the money in terms of moneyness and about 75% of traded puts are in the money. Additional sample statistics is found in figure 7.

The option price used in the paper is the average between the latest bid and latest ask, which is considered to be a better price estimate than the last traded price, also called the closing price. There is no implied volatility for expiration day options as volatility is zero at expiration. All options are European style so early exercise don't have to be taken into account. The futures options model Black 76 (which is explained more closely elsewhere in this paper) is used instead of the original Black–Scholes when a futures price is available, which it mostly is. When a futures price is not available the original Black–Scholes formula adjusted for a dividend yield is utilized. The STIBOR rates are obtained from riksbank.se, the underlying OMXS30 index quotes and the dividend yield AFFWALZ(DY) are obtained from Datastream.

The US option data consist of daily closing prices for S&P 500 index options between 1995 and 1996 provided by OptionMetrics. Implied volatiles are included with the sample and are assumed to be estimated from index futures with the futures options formula as that's the most common estimation method. The option price is assumed to be the average of the closing bid and ask as that's most common. Expiration week options are excluded from the sample as close to expiration implied volatilities are sometimes considered to be less reliable than implied volatilities further from expiration. Expiration for the US options occurs on the third Friday of the month, expiration for the Swedish options occur on the fourth Friday of the month. Trading volume and bid ask quotes are not available for the US data.

# 5. The implied distribution

We now turn to analyzing the results. The results for the whole sample are found in figure 3. Results per year are found in figure 7 and figure 7.1. Figure 7 consider call price data alone, without any put option prices. Figure 7.1 compare the call price results to put price results. The reason for making a distinction between call volatilities and put volatilities is the uncertainty around the extent to which market imperfections allow deviations from the put-call parity, discussed elsewhere in this paper.

If we turn to the overall results in figure 3 we observe the familiar volatility smile that have been found to prevail in most equity index option markets. This results in a visually leptokurtic distribution that is more pointed and has a fatter left tail than the lognormal distribution that is placed on top of the implied distribution as a reference. Although measured excess kurtosis is fairly small, only 1.0, and measured skew is close to zero. The lognormal overlay has the same mean and standard as the implied density, obtained from a discreet trapezoidal numerical integration of the continuous definition of mean and standard deviation

$$\mu = \int_{-\infty}^{\infty} P \times \xi \, d\xi$$

$$\sigma = \sqrt{\int_{-\infty}^{\infty} P \times (\xi - \mu)^2 d\xi}$$

The skew and kurtosis is also calculated from numerical integration as

$$skew = \int_{-\infty}^{\infty} P \times (\xi - \mu)^3 \, d\xi \bigg/ \left[\int_{-\infty}^{\infty} P \times (\xi - \mu)^2 \, d\xi\right]^{3/2}$$

$$excess\ kurtosis = \int_{-\infty}^{\infty} P \times (\xi - \mu)^4 \, d\xi \bigg/ \left[\int_{-\infty}^{\infty} P \times (\xi - \mu)^2 \, d\xi\right]^{2} - 3$$

where $P$ is the risk-neutral probability and $\xi$ is the moneyness. Figure 3 also show the variation in skew and kurtosis over time. The implied distribution has been consistently leptokurtic throughout the sample period but skew has varied between positive and negative. If we look at the volatility smile the implied volatility at the left tail of the distribution is on average double the volatility of the near the money options and the implied volatility at the right tail is on average around 50% higher than the implied volatility near the money. Data for individual years are shown in figure 8 and figure 8.1. Most plots use a fixed span for the volatility smile on the right-hand axis to enable an easier comparison of the differences in the range of the smile year to year. The volatility smile was the least pronounced in the calm years of 1994 to 1996 and 2004 to 2005, and more pronounced in other more turbulent years in terms of the implied volatility. 2004 and 2004 had record low implied volatilities, on level with the low around 1995.

If we turn to the plot that compare the put option smile and call options smile in figure 3, the results show that there can exist large differences between average put volatility and average call volatility for far away from the money options. On average for the whole sample in the money calls tend to be priced with a higher volatility than the equivalent out of the money put and vice versa. This gives a put option smile that is somewhat of a mirror image of the call option smile. At low moneyness (deep in the money calls and deep out of the money puts) there is an average deviation of several percentage points in implied standard deviation and the difference tends to increase as moneyess get lower. The same thing is true for high moneyness (deep out of the money calls and deep in the money puts). In terms of the implied distribution puts imply a more left skewed distribution that’s also more leptokurtic, even if only to a small degree. The put option implied distribution is slightly more left skewed with a measured skew of –0.5. The measured leptokurtosis is very similar, it's only 0.1 higher from the distribution implied by put options. The small differences in measured skew and kurtosis for call options when compared to put options comes from the fact that formula (6) and (7) in that case use the overall implied volatility for both puts and calls, as the options depend on the same underlying process and volatility.

For both puts and calls the aggregation of thousands of options makes the volatility smile smooth and little additional smoothing is needed. The smiles are visually smooth and with a pattern almost as consistent as if it was produced by an analytical model; an advantage with aggregating data is that the noise in the data is averaged away.

**Figure 3. The implied distribution.**

*Note:* Top left: Implied distribution and volatility smile for the entire sample. Top right: annual skew and kurtosis of the implied distribution. Only calls are considered. Skew and kurtosis for 1987 is deliberately outside the plot because of its large value and can be found in figure 7. Bottom left: Put option data overlaid the call option data. Bottom right: Volatility term structure.

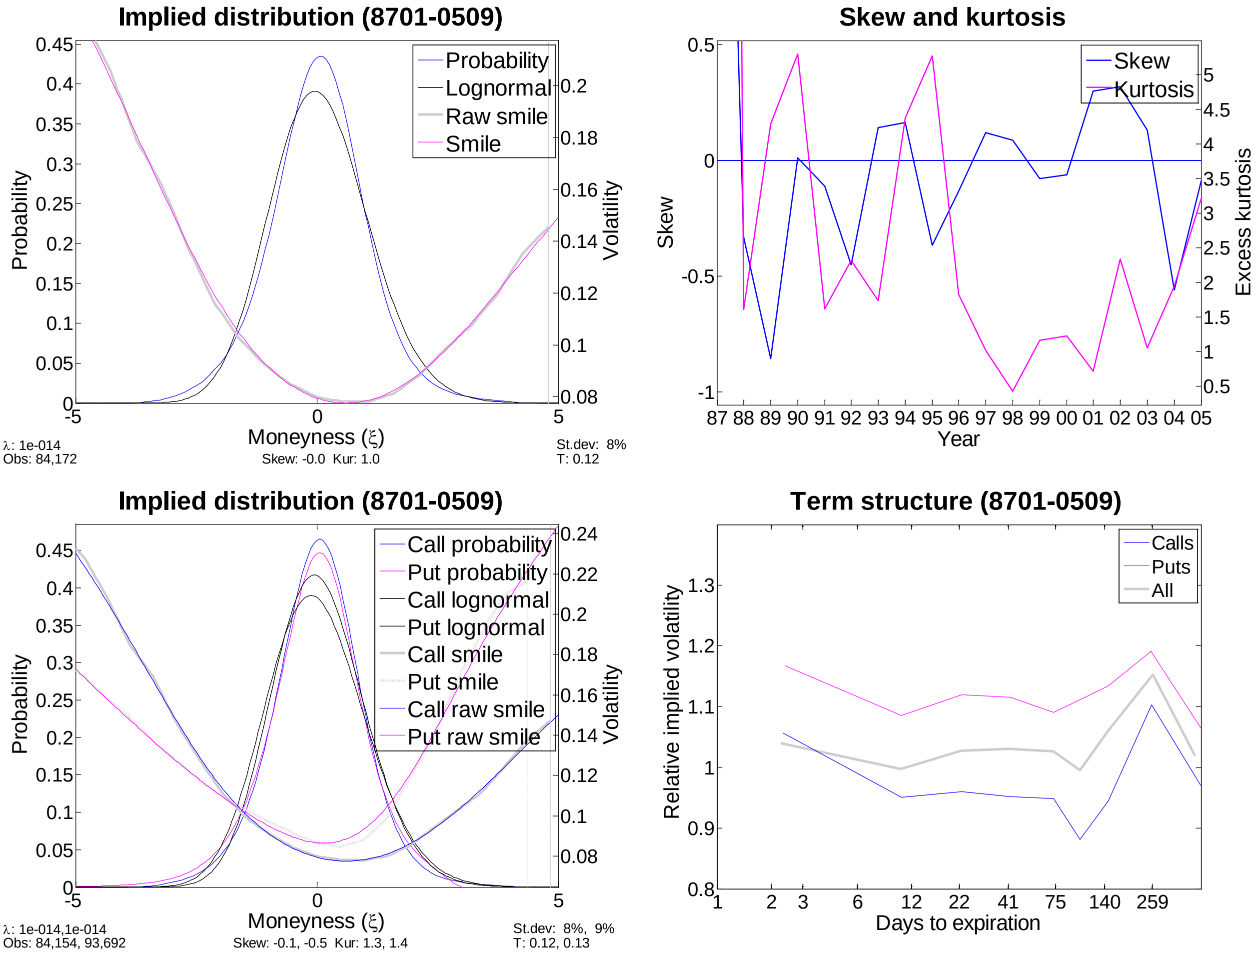


## 5.1. US comparison

We hypothesized that there wouldn't be a marked difference between the US and Swedish market and the results in figure 4 point in that direction. It's at least hard to argue that there is a marked difference between the probability distributions as the volatility smiles are close to symmetrical in slope and the probability distributions are similar. There is no clear pattern in skew or kurtosis between the years. Both distributions have a very small measured skew and excess kurtosis. In the first year the Swedish distribution is more leptokurtic than the US distribution while the second year show the reverse relationship since the US distribution is more right skewed than the Swedish distribution. Kurtosis doesn't show a consistent relationship either, in the first year the Swedish distribution is more leptokurtic while the other measures of moments are nearly indistinguishable. Visually there are no constant differences either. The US distribution is more peaked in the first year while the Swedish distribution is more peaked in the second year. It's hard to say if either distribution has a thinner tail than the other. The similarities are important for further studies of the actual distribution and the pricing kernel as a similar risk-neutral distribution makes it more probable that the actual distribution and the pricing kernel are similar too. If the pricing kernel is the same in both markets it means that the same relationships exist between the actual distributions of the two markets. We can also draw the important conclusion that a riskier market like the Swedish market doesn't need to a have a more leptokurtic distribution. In the first year the annualized standard deviation from the average of daily implied volatilities is 17% and 11% in the Swedish and the US market respectively, and in the second year it's 18% and 14% respectively. The lower liquidity in the Swedish market would be assumed to make the volatility smile steeper, with a lower liquidity leading to a higher premium and higher volatility for options far away from the money. Other things equal it's therefore possible that the Swedish market would have a less pronounced volatility smile if it was more liquid. It's hard to measure the significance of that effect though. The US option data in this thesis doesn't include bid and asks, only the average of the closing bid ask quote, so it's not possible with the data available for this paper to compare the bid asks spreads along the volatility smile. In conclusion the volatility smiles and implied distributions are almost identical symmetrically between the US and Swedish market and an affirmative answer is given to hypothesis 1, the US and Swedish distributions have similar properties.

**Figure 4. The implied distribution; US comparison.**

*Note:* Implied distribution from Swedish and US data. Only calls are considered.

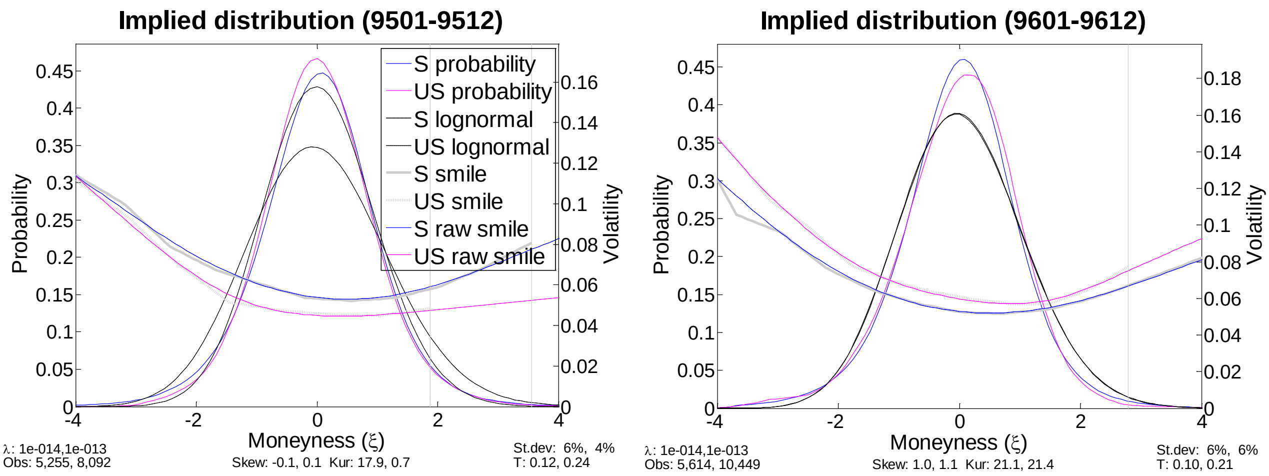

## 5.2. Volatility term structure

In addition to the volatility smile over moneyness there is a tendency for options with the same strike price but with different strike months to be priced with different implied volatilities. The primarily reason behind the term structure of volatilities is thought to be that implied volatility is expected to be mean reverting so that long term options are priced with higher volatilities if the short term volatility is unusually low, and vice versa. There may also be a tendency for the volatility smile over strike prices to be more or less pronounced for long- or short term options because of the relative usefulness of long- and short term options as insurance instruments or that the long term risk-neutral probability distribution of the underlying asset is different from the short term probability distribution.

As we aggregate data over such a long period as a year we would expect the variation in the short term volatility to even out and much of the term structure to disappear and this seems to largely be the result from the estimates presented in figure 3 and figure 9. What we see is that the variation in short term volatility evens out and there’s no clear pattern of an upward or downward sloping term structure throughout the series. There are years with both upward sloping and downward sloping and relatively flat term structures.

The method for investigating the term structure is to calculate the average implied volatility within different time to expiration groups. The primary source of uncertainty in this aggregated measure is the fact that some periods may have a larger prevalence of extreme moneyness options and therefore a larger prevalence of extreme relative implied volatilities. For example the expiration week seem to have the largest prevalence of such extreme observations judging by figure 10. If the average was over all options this would show up in the results, only options in the upper quartile of trading volume is therefore considered for every subgroup. This is considered to produce a result that’s representative of the relative volatility level and comparable between the different time to expiration groups as this excludes most options with extreme moneyness as their prices tend to come from very thin trading if they are traded at all in any particular day of the sample.

There’s not much known about the structural aggregated term structure of volatilities as studies of aggregated volatility structures does not have a very long history, the early studies mostly use volatility structures from option prices available in a given day or a given point in time. Another factor that could result in a term structure of volatilities is that the lognormal model’s skew is increasing in standard deviation which is increasing in time, and that would be able to manifest itself in a slope in the aggregated volatility term structure if the skew is increasingly deviating from the true risk-neutral distribution as time to expiration and volatility increase. Another possible reason is that many close to expiration options have almost no value so that trading costs and other market imperfections has a larger relative impact on these options and give them a relative premium. This could give close to expiration options a tendency for a higher volatility.

# 6. Nonparametric tests

So far the discussion has been largely of a qualitative nature and we now turn to a more quantitative discussion. By measuring the overall deviations from the Black–Scholes model we can measure the overall deviation from lognormality in the risk-neutral distribution. We conduct a nonparametric test to measure the economic significance of the deviations from the Black–Scholes model that is observed in market prices. The test was suggested by Rubinstein (1985) and consists of comparing matched pairs of options. It rests on the assumption in the Black–Scholes model that the implied volatility is the same for different strike prices for otherwise identical options, and the same for different time to expiration for otherwise identical options. Pairs are matched from options with the same strike prices, and from options with the same time to expiration. Only call options are considered. As in Rubinstein (1985) options with less than 21 days to expiration are not considered. The primary difference of including these short term options would be visible in the strike price paired options as many close to expiration options are far away from the money and accordingly have a very high implied volatility, and therefore increase the size of those minimum errors.

For the time to expiration comparison, to be matched, two records must belong to the same underlying stock, be observed on the same day during the same constant stock price interval, and have the same striking price. For the striking price comparison, two records are considered matched if they belong to the same underlying stock, are observed on the same day during the same constant stock price interval, and have the same time to expiration. All possible nonoverlapping matched pairs of option records is selected. That is, no option record appears in more than one matched pair.

As a comparison to the nonparametric test we can conduct a parametric test of statistical significance from the matched pairs. For a certain year there is 1029 matched pairs for options with the same time to expiration but with different striking prices. In 897 of these the option with the smaller moneyness, $K/S$, has the higher Black-Scholes implied volatility. Our null hypothesis is that the Black-Scholes formula produces unbiased values. If this were true, then for any matched pair, the probability is 0.5 that the option with the shorter maturity should have the higher implied volatility. So in a sample of 1029 matched pairs, we would expect 514.5 cases where the smaller moneyness option has the higher implied volatility. But in our sample we find 897 cases out of 1029. The probability that a result this extreme could have occurred given the truth of the null hypothesis is approximated by:

$$1 - N\left[\frac{(897 + 0.5) - (0.5 \times 1029}{0.5 \times \sqrt{1029}}\right]$$

where $N$ is the standard normal distribution function. This probability is at 24 standard deviations which is a very small probability. As a result we can safely reject the null hypothesis for these matched pairs. But the statistical significance doesn’t tell us about the economic significance of the pricing errors.

To complement the sign-test measure of statistical significance we will use the following nonparametric measure of economic significance. If the null hypothesis is true, then if the matched pairs are ranked by the difference in implied volatilities of the options in each matched pair, the median difference will be expected to be zero. Unfortunately, the median difference in implied volatilities is not readily translated into a dollar or percentage pricing bias from Black–Scholes values. Therefore, a similar but more meaningful measure of economic significance is used. For each matched pair we calculated the implied volatility which equalizes the percentage difference of the price of each call from its value. This is the $\sigma$ which satisfies

$$\frac{c_1(\sigma) - c_1^*}{c_1^*} = \frac{c_2^* - c_2(\sigma)}{c_2^*} \equiv 0.01\alpha$$

where $c_i$ is the Black–Scholes values of a call at volatility $\sigma$, and $c_i^*$ is the market price of the corresponding call for $i = 1, 2$. Because $\partial C/\partial\sigma > 0$, $\alpha$ can be interpreted as the minimum percentage deviation of the market prices of the options in the matched pair from their corresponding Black–Scholes values, over all estimates of volatility. In other words, we can say that, irrespective of the level of stock volatility, one of the options in the matched pair must be mispriced relative to Black-Scholes values by at least this percent. This lower bound measure of economic bias from the Black-Scholes formula does not require an estimate of stock volatility.

Let's look at a matched pair to understand how the calculation takes place. One of the pairs in the sample is

$$c_1^* = 16.5, K_1 = 740, \sigma_1 = 0.1313, T_1 = 25$$
$$c_2^* = 33, K_2 = 740, \sigma_2 = 0.1339, T_2 = 200$$

where $c_i$ is the call price, $K_i$ the strike price, $\sigma_i$ the implied volatility and $T_i$ the time to expiration in days. The Black–Scholes model assume these options to have the same volatility as they have the same strike price and time to expiration. We calculate Black-Scholes prices for the two options for a range of different volatilities, and for each volatility we save the percentage error, $(c_i(\sigma) - c_i^*)/c_i^*$ , that is the largest from the two errors. We find that the volatility 0.1302 produce the smallest percentage error. This volatility is always found between the volatilities of the two options which simplify the calculation. At this volatility option two has the largest error of the two options, $(c_2(\sigma) - c_2^*)/c_2^* = 0.0226$. We conclude that one of the two options must be mispriced, relative to Black-Scholes values, by at least this much.

The results are reported in figure 5 as the median pricing errors from the annual subsets of matched pairs. The percentage pricing errors show an overall declining trend, meaning that the Black–Scholes violations have declined overall. The smallest median percentage error for any year in the sample is around 1% as the figure show. This lowest percentage error occurs when pairs are constructed from the same striking price. When pairs are constructed from the same time to expiration percentage errors tend to be slightly higher. The test therefore also tells us that model violations across strike prices are larger than model violations across time to expiration. The first year in the sample has the largest percentage errors, which can be expected as a large part of the year was very turbulent with high levels of uncertainty. In summary the lower bound pricing errors of several percent can be considered large so from these results we can accept hypothesis 2 and conclude that the Black-Scholes model cannot sufficiently explain option prices.

It's hard to compare the size of these pricing errors with those found in Rubinstein (1985) as that article considers options on underlying equities rather than the index that often has a higher trading volume, and the aggregated samples consist of several subsamples of strike prices and time to expiration groups. What we can say is that many of the pricing errors are in the range of 1% to 5% in Rubinstein (1985) which these results also are. We can make a direct comparison to the US market from the US option data available for this paper. For the US S&P 500 options the time to expiration pairing percentage pricing errors are 2.4% and 4.4% for 1995 and 1996 respectively, compared to the Swedish pricing errors of 3.0% for both years. For strike price paired options the percentage pricing errors are 2.9% and 3.2% for 1995 and 1996 respectively, compared to 1.0% and 1.2% respectively for the Swedish data. In most cases the pricing errors are therefore smaller in the Swedish market and that indicates that the Swedish option-implied risk-neutral distribution is closer to the lognormal distribution than implied distribution for the US market.

We should also emphasize that this doesn't test if options are correctly priced; it tests the deviations from the lognormal risk-neutral distribution. If the lognormal risk-neutral distribution doesn't sufficiently explain option prices these pricing errors can prevail even in the absence of arbitrage opportunities.

**Figure 5. Nonparametric tests.**

*Note:* A lower bound measure of economic significance of Black–Scholes violation calculated from matched pairs annually. The 1987 percentage error is deliberate out of bounds as it has a large value of 9.2 with T pairing and 13.5 with K pairing.

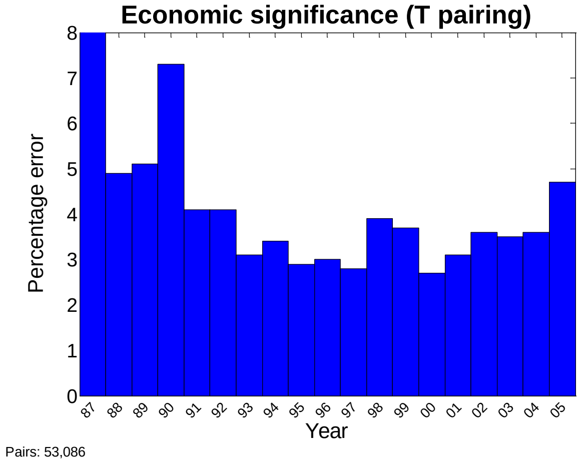


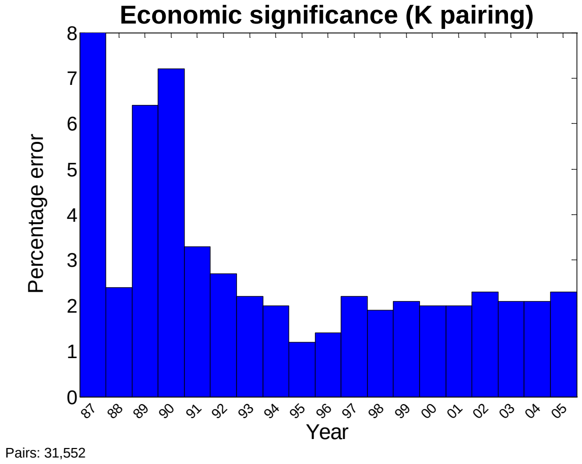


# 7. Parametric tests

To answer the hypothesis about time-variation and put-call parity in the volatility smile we turn to parametric tests. The standardized implied volatility and moneyness variables allow use to aggregate data and employ statistical inference to answer our hypothesis about implied volatility.

In terms of model selection for the test we cannot use a linear regression as the data don't suggest a liner relationship between implied volatility and moneyness. We have to use at least a quadratic equation. But that may not be flexible enough, so to allow additional flexibility we use a cubic equation of the form

$$\widehat{\sigma_\iota} = \widehat{\beta_1} + \xi_i\widehat{\beta_2} + \xi_i^2\widehat{\beta_3} + \xi_i^3\widehat{\beta_4} + \widehat{u_\iota}$$

where $\sigma_i$ is the relative implied volatility and $\xi_i$ the moneyness.

The tests are done for all years of the sample, rather than the whole sample at the same time, to be able to identify possible trends over time. The samples are still large enough as each year contains thousands of options prices. The sample use only call options to remove the impact of the potential uncertainty over how put-call parity holds for the volatility smile or if puts have their own volatility smile.

To test the hypotheses that the regression coefficients have been constant over the entire period against the hypotheses that they have not been constant (i.e. that the volatility smile has been time-varying) we can compare the size of the residuals in the two models in a goodness of fit $\chi^2$ test with an observed

$$\hat{\chi}_p^2 = \frac{\sum_{j=1}^{p} \hat{e}_{n+j}^2}{\sigma_e^2}$$

We can also compare the size of the residuals in the two models in an F test with an observed

$$\hat{F} = \frac{(\widehat{SSE_1} - \widehat{SSE_2})/m}{\widehat{SSE_2}/(n-k)}$$

To understand how the hypothesis test works for these test statistics let's look at the hypothesis test for the 1987 vs. 1988 period $\chi^2$ test. The result from this test is given in table 1.

$H_0$: $\chi_p^2 = 0$ (the same model)
$H_1$: $\chi_p^2 \neq 0$ (not the same model)
Test statistic: $\hat{\chi}_p^2 = \frac{\sum_{j=1}^{p} \hat{e}_{n+j}^2}{\sigma_e^2}$
$n$ = number of observations in the base period = 2415
$p$ = number of observations in the period we compare to the base period = 1767
Significance level: 5%
Degrees of freedom: $df = p = 1767$
Confidence levels (5% on both sides): $1670 < \chi_{obs}^2 < 1866$
Decision rule: Reject $H_0$ if $\chi_{obs}^2 > 1866$
Result: $\chi_{obs}^2 \approx 94841$
Conclusion: If both periods had the same true model as the base period we would expected our $\chi_p^2$ value to be $\chi^2$ distributed with 1767 degrees of freedom. Our critical level is 1866 (the 5% $\chi^2$ value) and our observed $\chi^2$ value of 94841 is larger than that so that we reject $H_0$. Both periods are not explained by the same model.

Let's also look at the hypothesis test for the F-test for the 1987 vs. 1988 period from the test given in table 1 to understand how this hypothesis test is constructed.

$H_0 : F = 0$ (the same model)
$H_1 : F \neq 0$ (not the same model)
Test statistic: $\hat{F} = \frac{(\widehat{SSE_1} - \widehat{SSE_2})/m}{\widehat{SSE_2}/(n-k)}$
$\widehat{SSE}_1$= the sum of the squared errors from a regression for both sub-periods ≈ 37758
$\widehat{SSE}_2$ = the sum of the squared errors from a regression of the base period ≈ 115
$n$ = number of observations in the base period = 2415
$m$ = number of observations in the period we compare the base period to = 1767

$k$ = number of coefficients in the regression model = 4
Significance level: 5%
Degrees of freedom: df $= \{m, n-k\} = \{1767, 2411\}$
Confidence levels (5% on both sides): $0.9 < F_{obs} < 1.1$
Decision rule: Reject $H_0$ if $F_{obs} > 1.1$
Result: $F_{obs} \approx 304$
Conclusion: If both periods had the same true model as the base period we would expected our $\hat{F}$ value to be F distributed with {1767, 2411} degrees of freedom. Our critical level is 1.1 (the 5% F value) and our observed F value of 304 is larger than that so that we reject $H_0$. Both periods are not explained by the same model.

We can now interpret the test results which are given in table 1. 12 of 18 time periods have a significant $\chi^2$ statistic and therefore significant variations in the volatility smile from one year to the next. The F tests also show significant variations, the F statistic is significant for all time periods. We therefore answer yes to hypothesis 3 and conclude that there is significant time-variation in the volatility smile.

We can expect the volatility smile and the implied risk-neutral distribution to vary in shape and form over time as the economy changes and investors adapt to changes and adapt their forecasts. We have not by this test specified which factors are responsible for the variations in the volatility smile. Important possible explanations are changes in the likelihood for crash scenarios and possibly changes in risk aversion over time or as overall implied volatility increase or decrease.

Now let's compare the volatility smile implied by calls and the volatility smile implied by puts. In the absence of market imperfections the put-call parity holds according to formula (14) and (15) which mean that a call's put equivalent cannot be priced according to a different implied volatility. We now allow put volatilities into the sample and compare them to call volatilities with the same $\chi^2$ test and F test that we introduced for the previous test. The test results are given in table 2. 16 of 19 $\chi^2$ statistics are significant. Only in three of the 19 years are the smiles similar enough to not be statistically significantly different. The F tests all show a significant difference. We therefore answer hypothesis 4 affirmatively that the call volatility is different from the put volatility.

What kind of market imperfections can cause this departure from put-call parity? The bid-ask spread, trading costs and lack of liquidity will allow a level of deviation from perfect arbitrage free conditions. We also emphasize that the statistical significance doesn't tell us about the economic significance. We don't know from these parametric tests how big the deviation is in economic terms.

**Table 1. Time-variation of volatility smile test.**

*Note:* The columns are sample period, number of observations in base sample, number of observations in compared sample, $\hat{\chi}_p^2$, p-value of $\hat{\chi}_p^2$, $\hat{F}$ and p-value of $\hat{F}$. A star, *, mark test statistics significant at 5%.

| Period | n1 | n2 | X2 | p_X2 | F | p_F |
|---|---|---|---|---|---|---|
| 87 vs 88 | 2,415 | 1,767 | *94,841 | 0 | *304 | 0 |
| 88 vs 89 | 1,767 | 2,336 | *7,491 | 0 | *295 | 0 |
| 89 vs 90 | 2,336 | 2,790 | *2,798 | 2.02e-06 | *78 | 0 |
| 90 vs 91 | 2,790 | 2,886 | 2,298 | 1 | *120 | 0 |
| 91 vs 92 | 2,886 | 3,277 | *113,154 | 0 | *164 | 0 |
| 92 vs 93 | 3,277 | 3,156 | *3,029,699 | 0 | *152 | 0 |
| 93 vs 94 | 3,156 | 4,236 | *14,469 | 0 | *36 | 0 |
| 94 vs 95 | 4,236 | 5,402 | *13,829 | 0 | *13 | 0 |
| 95 vs 96 | 5,402 | 5,766 | 487 | 1 | *5 | 0 |
| 96 vs 97 | 5,766 | 7,077 | *14,628,774 | 0 | *133 | 0 |

| | | | | | | |
|---|---|---|---|---|---|---|
| 97 vs 98 | 7,077 | 6,641 | 447 | 1 | *2 | 0 |
| 98 vs 99 | 6,641 | 6,699 | *83,447 | 0 | *164 | 0 |
| 99 vs 00 | 6,699 | 5,824 | 4,460 | 1 | *135 | 0 |
| 00 vs 01 | 5,824 | 6,966 | 4,847 | 1 | *163 | 0 |
| 01 vs 02 | 6,966 | 7,012 | *7,754 | 1.99e-12 | *266 | 0 |
| 02 vs 03 | 7,012 | 4,490 | *18,229 | 0 | *405 | 0 |
| 03 vs 04 | 4,490 | 4,985 | 4,907 | .3747186 | *94 | 0 |
| 04 vs 05 | 4,985 | 3,751 | *8,516 | 0 | *219 | 0 |

**Table 2. Call vs. put volatility smile test.**

*Note:* The columns are sample period, number of observations in base sample, number of observations in compared sample, $\hat{\chi}_p^2$, p-value of $\hat{\chi}_p^2$, $\hat{F}$ and p-value of $\hat{F}$. A star, *, mark test statistics significant at 5%.

| Period | n1 | n2 | X2 | p_X2 | F | p_F |
|---|---|---|---|---|---|---|
| 87 | 2,415 | 2,324 | *6,961 | 0 | *507 | 0 |
| 88 | 1,767 | 2,102 | *81,262 | 0 | *639 | 0 |
| 89 | 2,336 | 2,900 | *19,890 | 0 | *148 | 0 |
| 90 | 2,790 | 2,843 | *122,237 | 0 | *256 | 0 |
| 91 | 2,886 | 2,748 | *12,287 | 0 | *427 | 0 |
| 92 | 3,277 | 3,686 | 3,358 | .9962338 | *280 | 0 |
| 93 | 3,156 | 3,491 | *3,891 | 3.16e-07 | *94 | 0 |
| 94 | 4,236 | 4,902 | *7,692 | 0 | *30 | 0 |
| 95 | 5,402 | 5,553 | *13,690 | 0 | *11 | 0 |
| 96 | 5,766 | 6,105 | *12,379 | 0 | *329 | 0 |
| 97 | 7,077 | 7,222 | *218,582 | 0 | *4 | 0 |
| 98 | 6,641 | 7,236 | *12,533 | 0 | *330 | 0 |
| 99 | 6,699 | 7,902 | 7,228 | .9999992 | *214 | 0 |
| 00 | 5,824 | 6,107 | *7,944 | 0 | *402 | 0 |
| 01 | 6,966 | 7,467 | *32,816 | 0 | *541 | 0 |
| 02 | 7,012 | 7,497 | *50,062 | 0 | *528 | 0 |
| 03 | 4,490 | 4,949 | *130,968 | 0 | *205 | 0 |
| 04 | 4,985 | 5,828 | 4,461 | 1 | *303 | 0 |
| 05 | 3,751 | 4,783 | *20,891 | 0 | *158 | 0 |

# 8. Conclusions

We find that the Swedish equity index implied probability distribution is similar in shape to the US equity index implied probability distribution which is an important result for further studies of the actual probability distribution and the pricing kernel in the Swedish market. We find that the Black–Scholes model cannot sufficiently explain option prices for the Swedish equity index market. This result is in line with research for other options markets. Compared to other option markets the Swedish option market may be more difficult to analyze as it's less liquid and therefore allows more anomalies and larger pricing errors compared to the true pricing model. We find that it's relatively easy to detect statistically significant deviations from the Black-Scholes model, while the economic significance is still ambiguous, but we also found economically significant deviations. We find that put-call parity is violated by analyzing the volatility smile, and that the volatility smile is time-varying. The implied volatility is not static but evolve as a stochastic process, which needs to be modeled. We have gained valuable insights about the estimation and analysis of implied probability distributions and there are many natural extensions this research, such as what the risk-neutral probability distribution implies for the actual probability distribution and the pricing kernel in the Swedish equity index market, and what factors can explain the pronounced volatility term structure smile. Other research directions that are on a more general level is the development of new parametric option-pricing models that incorporating additional stochastic factors, such as stochastic jumps, stochastic volatility, and stochastic interest rates.

We find that the risk-neutral probability distribution has many valuable applications, such as pricing exotic options, investigating the impact of news on probability assessments and for forecasting. Risk-neutral probabilities in conjunction with actual probabilities also tell us about implied utility functions and therefore about the economy-wide preferences that investors exhibit. This information has important applications for many areas of economics. The relationship between the risk-neutral probability, actual probabilities and the pricing kernel is still a puzzle and it will take time to understand it. The most important future developments lie in finding an improved understanding of how to translate the risk-neutral to actual distributions, especially for securities that are correlated with aggregate wealth but not perfectly so. Forecasting and risk management will become easier when research developments allow us to use the risk-neutral distributions even more than we do now.

# References

Bachelier, L, 1900, Theorie de la Speculation, *Annales de l'Ecole Normale Superieure*, vol. 17

Bakshi, G., C. Cao, and Z. Chen, 1997, Empirical Performance of Alternative Option Pricing Models, *Journal of Finance* 52, 2003–49

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

Bates, D., 2000, Post-'87 Crash Fears in S&P 500 Futures Options, Journal of Econometrics 94, 181–238

———, 2001, *The Market for Crash Risk*, Working Paper No. 8557, National Bureau of Economic Research, Cambridge, MA

Benninga, S., and J. Mayshar, 1997, *Heterogeneity and Option Pricing*, Working paper, University of Pennsylvania

Black, F. and M. Scholes, 1973, The pricing of options and corporate liabilities, *Journal of Political Economy* 81, 637-654

Black, F., 1976, The Pricing of Commodity Contracts, *Journal of financial Economics* 3, 167-179

Brown, D., and J. Jackwerth. 2003, The Kernel Pricing: Reconciling Index Option Data and Economic Theory, Working paper, University of Wisconsin at Madison

Campbell, J., and L. Hentschel, 1992, No News Is Good News: An Asymmetric Model of Changing Volatility in Stock Returns, *Journal of Financial Economics* 31, 281–318

Franke, G., D. Stapleton, and M. Subrahmanyam, 1999, When Are Options Overpriced? The Black–Scholes Model and Alternative Characterisations of the Pricing Kernel, *European Finance Review* 3, 79–102

French, K., G. Schwert, and R. Stambaugh, 1987, Expected Stock Returns and Volatility, *Journal of Financial Economics* 19, 3–29

Gennotte, G., and H. Leland. 1990, Market Liquidity, Hedging, and Crashes, *American Economic Review* 80, 999–1021

Grossman, S., 1988, An Analysis of the Implications for Stock and Futures Price Volatility of Program Trading and Dynamic Hedging Strategies, *Journal of Business* 61, 275–298

Grossman, S., and Z. Zhou, 1996, Equilibrium Analysis of Portfolio Insurance, *Journal of Finance* 51, 1379–1403

Guidolin, M., and A. Timmermann., 2000, *Option Prices and Implied Volatility Dynamics under Bayesian Learning*, Working paper, University of California at San Diego

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

Jackwerth, J. C., 2000, Recovering Risk Aversion from Options Prices and Realized Returns, *Review of Financial Studies* 13, 433-451

———, 2004, *Option-Implied Risk-Neutral Distributions and Risk Aversion*, Research Foundation of AIMR

Jarrow, R., and A. Rudd, 1982, Approximate Valuation for Arbitrary Stochastic Processes, *Journal of Financial Economics* 10, 347–369

Kelly, M., 1994, *Correlation: Stock Answer*, Risk 7, 40–43

Longstaff, F., 1990, *Martingale Restriction Tests of Option Pricing Models*, Working paper, University of California at Los Angeles

Merton, R., 1973, Theory of Rational Option Pricing, *Bell Journal of Econometrics and Management Science* 4, 141-183

Pan, J., 2002, The Jump-Risk Premia Implicit in Options: Evidence from an Integrated Time-Series Study, *Journal of Financial Economics* 63, 3–50

Pindyck, R. 1984, Risk, Inflation, and the Stock Market, *American Economic Review* 74, 334–351

Ritchey, R. 1990, Call Option Valuation for Discrete Normal Mixtures, *Journal of Financial Research* 13, 285–295

Romer, D., 1993, Rational Asset-Price Movements without News, *American Economic Review* 83, 1112–30

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* 40, 455–480

———, 1994, Implied Binomial Trees, *Journal of Finance* 49, 771-818

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* 12, 275–290

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

Toft, K., and B. Prucyk., 1997, Options on Leveraged Equity: Theory and Empirical Tests, *Journal of Finance* 52, 1151–80

Ziegler, A, 2003, *Why Does Implied Risk Aversion Smile?*, Working paper, HEC

**Figure 6. The lognormal distribution.**

*Note:* A lognormal distribution with an overlay of a normal and Student's t distribution for reference, to illustrate the properties of the lognormal distribution. The distributions are centered on the same arithmetic mean. The diagonal dotted lines illustrate the direction that option payoff increase. The vertical dotted line marks the zero-value location when it's within the bounds of the plot.

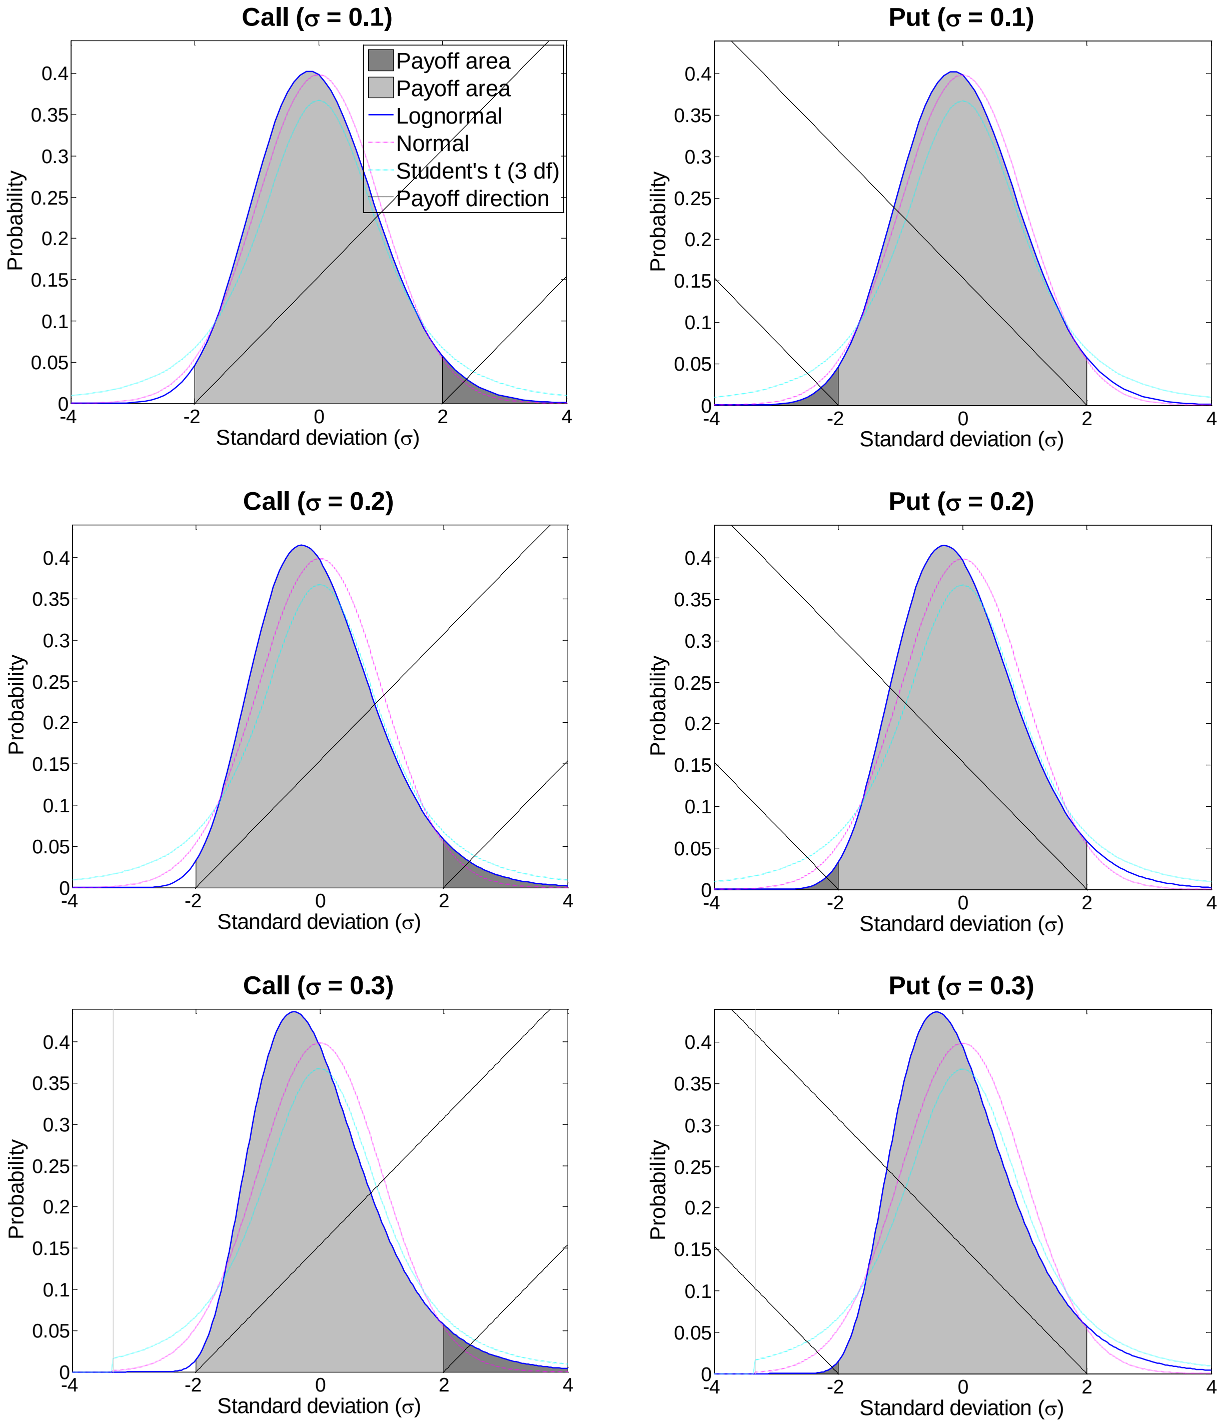


**Figure 7. Sample statistics.**

*Note:* Top-left figure: The daily implied volatility calculated as the mean of the implied volatility for the two most traded options during that day. Right hand table: selected sample statistics for options trading volume and where implied volatility is non-negative and therefore can be estimated. Bottom-left figure: annual trading volume.

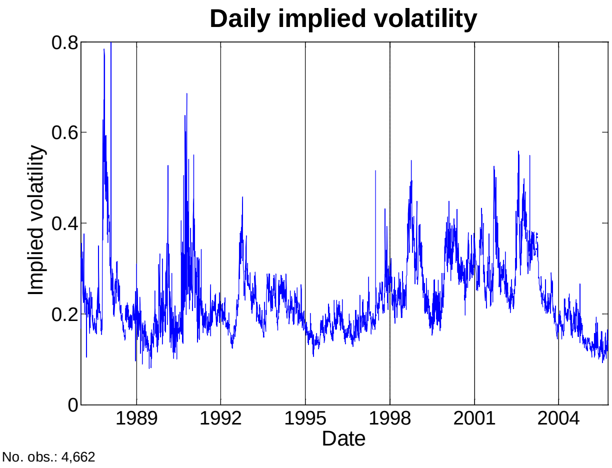


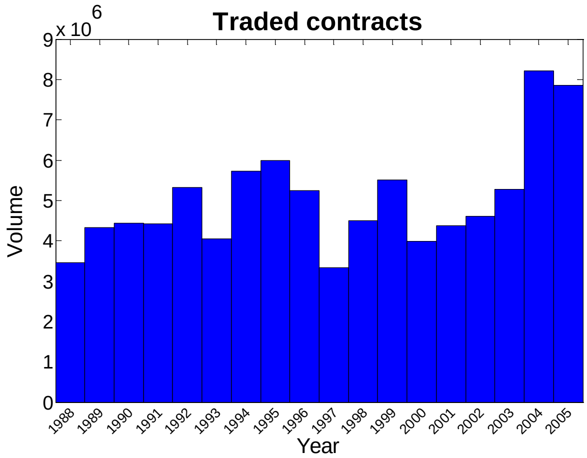


**Sample statistics**

```
Number of observations:
Total                 : 401,962 100%
Volume > 0            : 178,217  44%
Imp. volatility > 0: 177,611  44%

Moneyness (xi) distribution:
xi > -1.5 & < 1.5: 143,415 81%

Distribution of calendar days to expiration:
<= 5       :  7%
<= 30      : 47%
<= 30 * 3 : 93%
<= 30 * 12: 99%

Average calendar days to expiration: 45 days
... for daily implied volatility: 31 days
Average implied volatility: 23%
... for all options: 27%

Distribution of option types:
Calls: 47%
Puts: 53%

In the money options:
Calls as % of all calls: 35% (xi < 0)
Puts as % of all puts: 28% (xi > 0)
Of all options: 31%
```

## Figure 8. Implied distribution aggregated per year, from calls.

*Note:* The implied risk-neutral distribution over moneyness from call price data aggregated annually. The standard deviation of the implied density and sample average time to expiration is given in the upper left corner of the plot, the value of the tradeoff parameter $\lambda$ in the upper right corner of the plot and the number of aggregated observations in the bottom left corner of the plot.

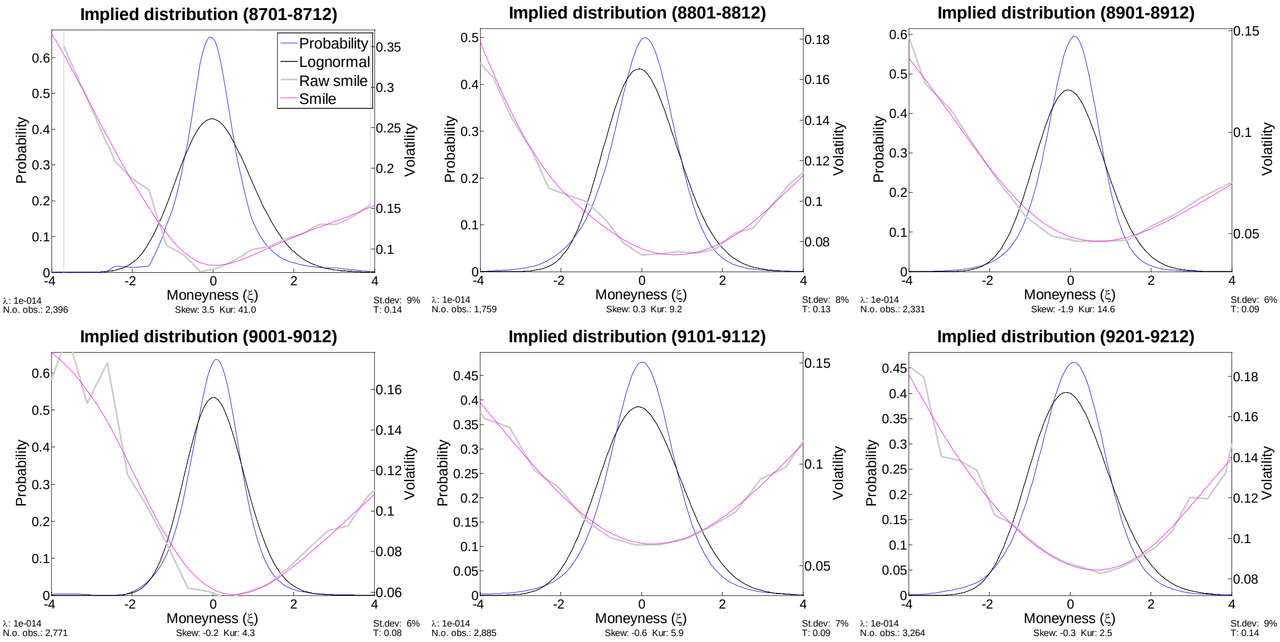

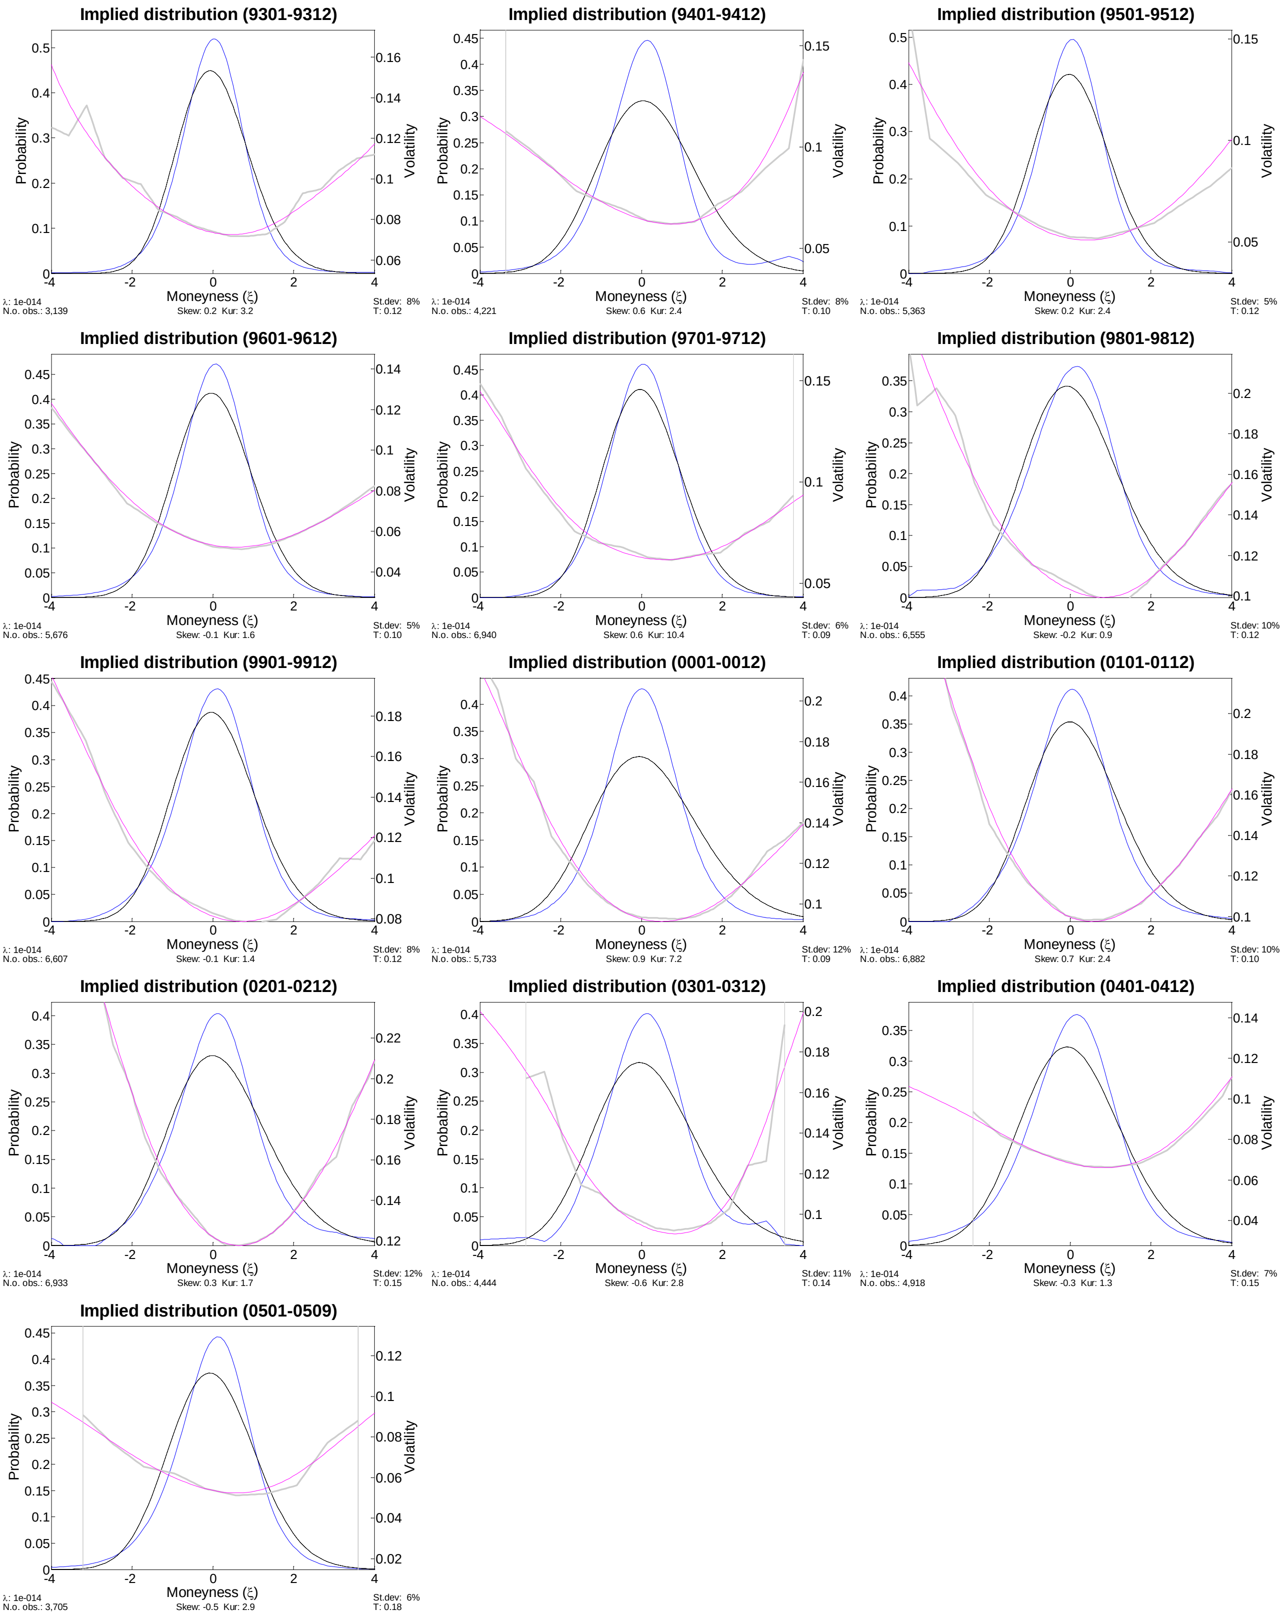


**Figure 8.1. Implied distribution aggregated per year, calls and puts.**

*Note:* The call data from figure 6 compared to put data.

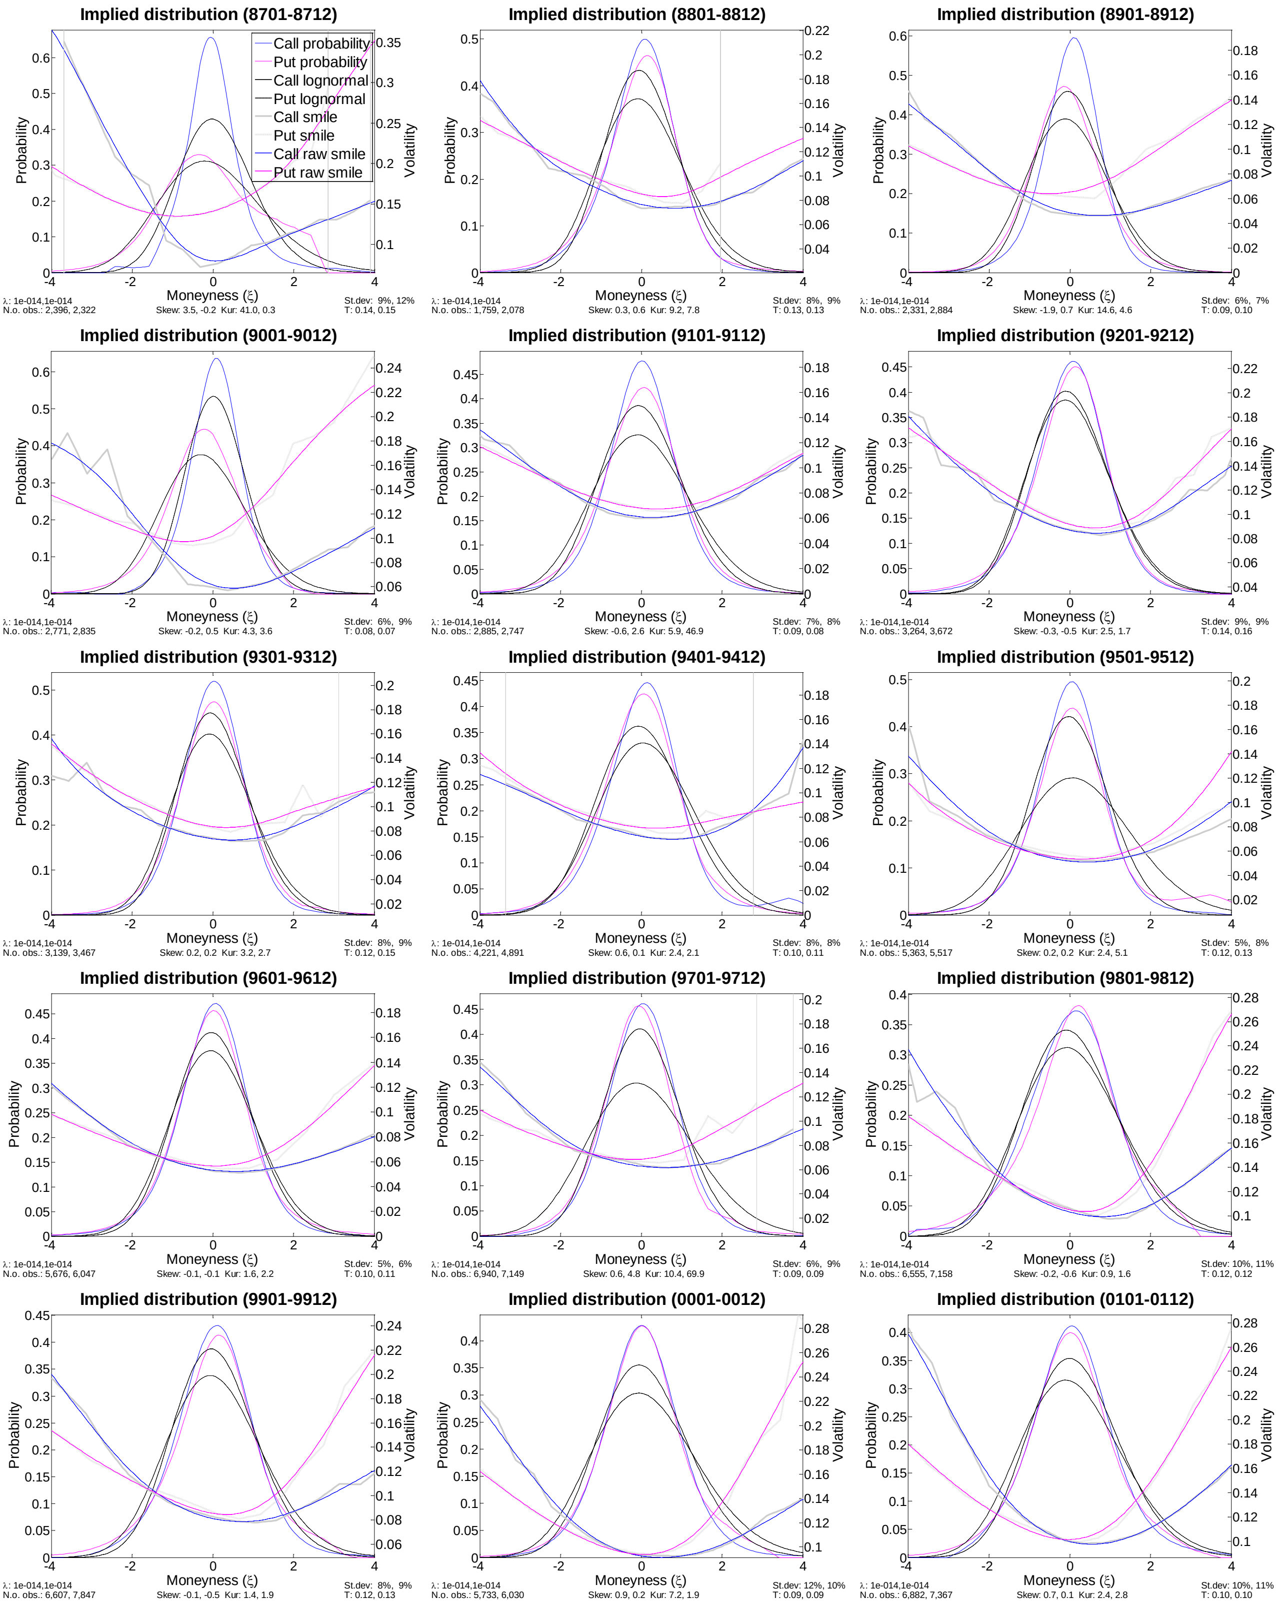
Implied distribution (8701-8712)
Call probability
Put probability
Call lognormal
Put lognormal
Call smile
Put smile
Call raw smile
Put raw smile
Probability
Volatility
Moneyness (ξ)
λ: 1e-014,1e-014
N.o. obs.: 2,396, 2,322
Skew: 3.5, -0.2 Kur: 41.0, 0.3
St.dev: 9%, 12%
T: 0.14, 0.15
Implied distribution (8801-8812)
λ: 1e-014,1e-014
N.o. obs.: 1,759, 2,078
Skew: 0.3, 0.6 Kur: 9.2, 7.8
St.dev: 8%, 9%
T: 0.13, 0.13
Implied distribution (8901-8912)
λ: 1e-014,1e-014
N.o. obs.: 2,331, 2,884
Skew: -1.9, 0.7 Kur: 14.6, 4.6
St.dev: 6%, 7%
T: 0.09, 0.10
Implied distribution (9001-9012)
λ: 1e-014,1e-014
N.o. obs.: 2,771, 2,835
Skew: -0.2, 0.5 Kur: 4.3, 3.6
St.dev: 6%, 9%
T: 0.08, 0.07
Implied distribution (9101-9112)
λ: 1e-014,1e-014
N.o. obs.: 2,885, 2,747
Skew: -0.6, 2.6 Kur: 5.9, 46.9
St.dev: 7%, 8%
T: 0.09, 0.08
Implied distribution (9201-9212)
λ: 1e-014,1e-014
N.o. obs.: 3,264, 3,672
Skew: -0.3, -0.5 Kur: 2.5, 1.7
St.dev: 9%, 9%
T: 0.14, 0.16
Implied distribution (9301-9312)
λ: 1e-014,1e-014
N.o. obs.: 3,139, 3,467
Skew: 0.2, 0.2 Kur: 3.2, 2.7
St.dev: 8%, 9%
T: 0.12, 0.15
Implied distribution (9401-9412)
λ: 1e-014,1e-014
N.o. obs.: 4,221, 4,891
Skew: 0.6, 0.1 Kur: 2.4, 2.1
St.dev: 8%, 8%
T: 0.10, 0.11
Implied distribution (9501-9512)
λ: 1e-014,1e-014
N.o. obs.: 5,363, 5,517
Skew: 0.2, 0.2 Kur: 2.4, 5.1
St.dev: 5%, 8%
T: 0.12, 0.13
Implied distribution (9601-9612)
λ: 1e-014,1e-014
N.o. obs.: 5,676, 6,047
Skew: -0.1, -0.1 Kur: 1.6, 2.2
St.dev: 5%, 6%
T: 0.10, 0.11
Implied distribution (9701-9712)
λ: 1e-014,1e-014
N.o. obs.: 6,940, 7,149
Skew: 0.6, 4.8 Kur: 10.4, 69.9
St.dev: 6%, 9%
T: 0.09, 0.09
Implied distribution (9801-9812)
λ: 1e-014,1e-014
N.o. obs.: 6,555, 7,158
Skew: -0.2, -0.6 Kur: 0.9, 1.6
St.dev: 10%, 11%
T: 0.12, 0.12
Implied distribution (9901-9912)
λ: 1e-014,1e-014
N.o. obs.: 6,607, 7,847
Skew: -0.1, -0.5 Kur: 1.4, 1.9
St.dev: 8%, 9%
T: 0.12, 0.13
Implied distribution (0001-0012)
λ: 1e-014,1e-014
N.o. obs.: 5,733, 6,030
Skew: 0.9, 0.2 Kur: 7.2, 1.9
St.dev: 12%, 10%
T: 0.09, 0.09
Implied distribution (0101-0112)
λ: 1e-014,1e-014
N.o. obs.: 6,882, 7,367
Skew: 0.7, 0.1 Kur: 2.4, 2.8
St.dev: 10%, 11%
T: 0.10, 0.10

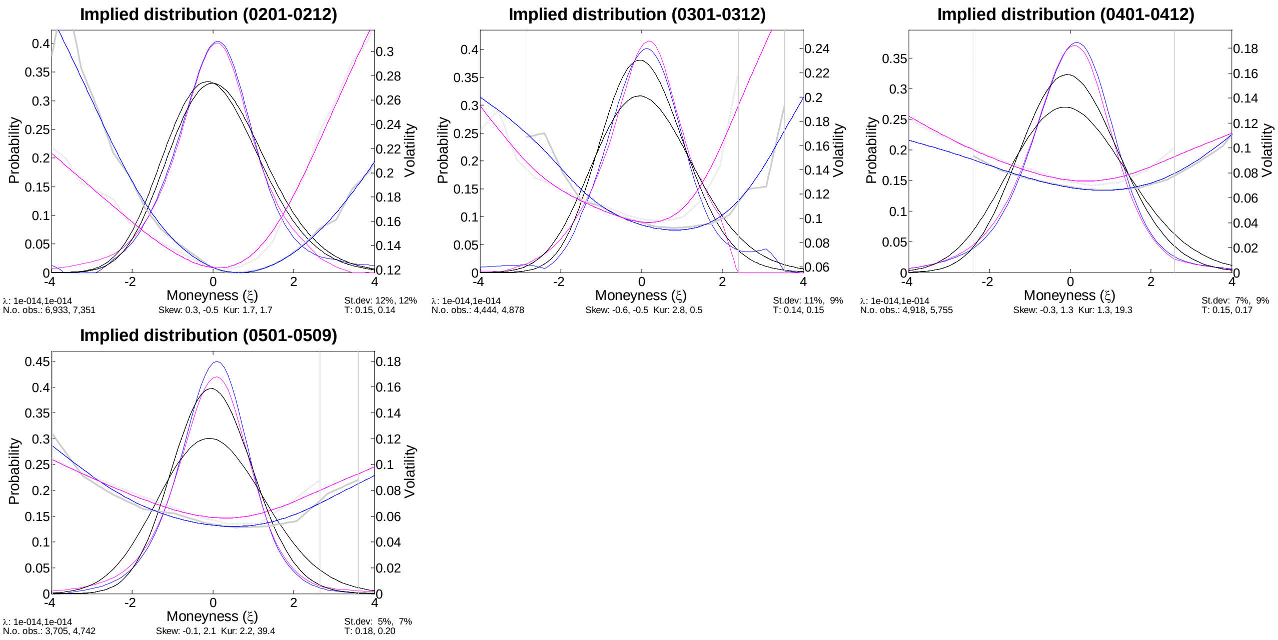


## Figure 9. Term structure of volatility smile.

*Note:* Term structure of volatility smile from aggregated data. Data is taken from the upper quartile of trading volume. x-axis is logarithmic.

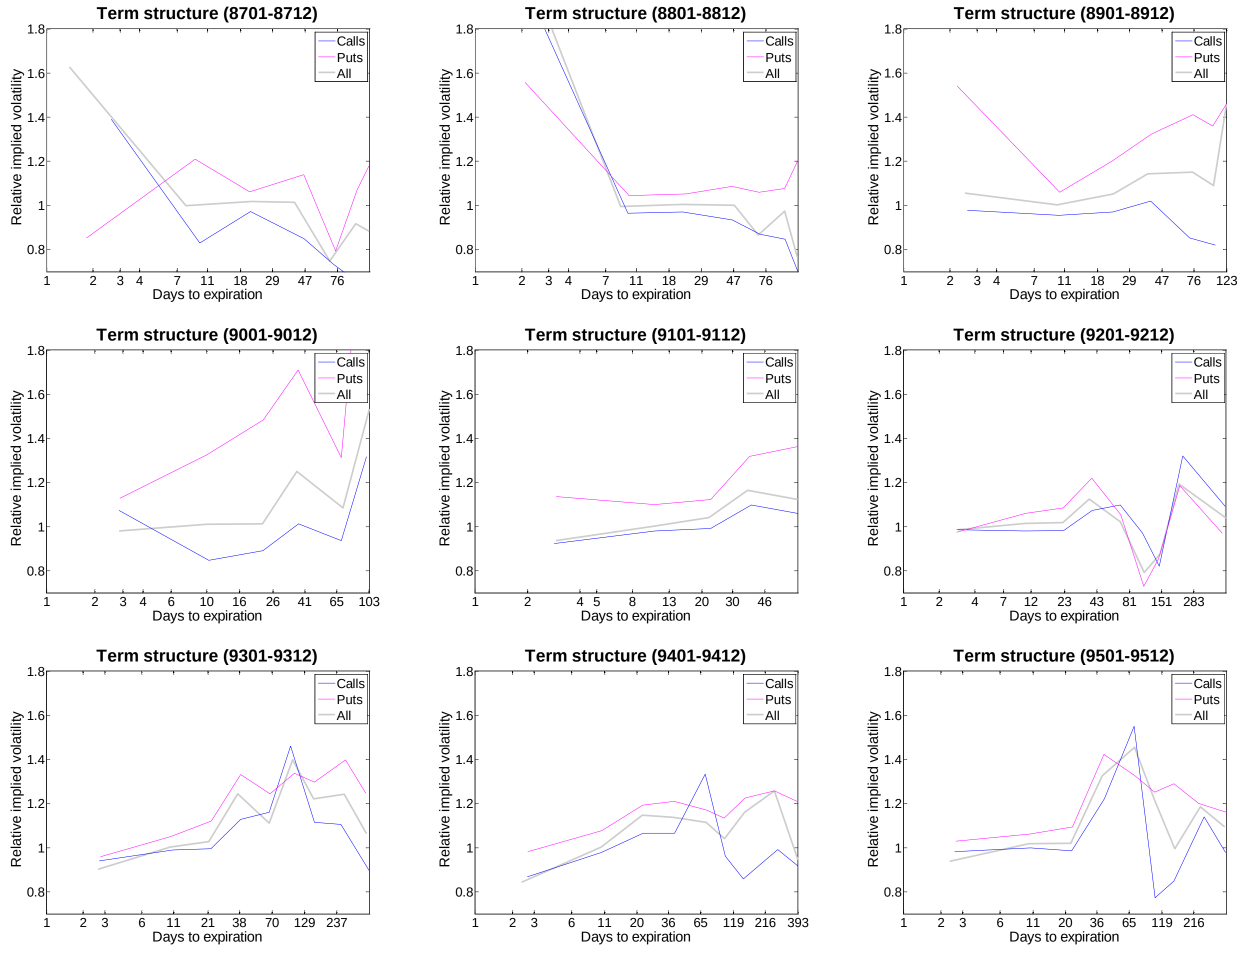

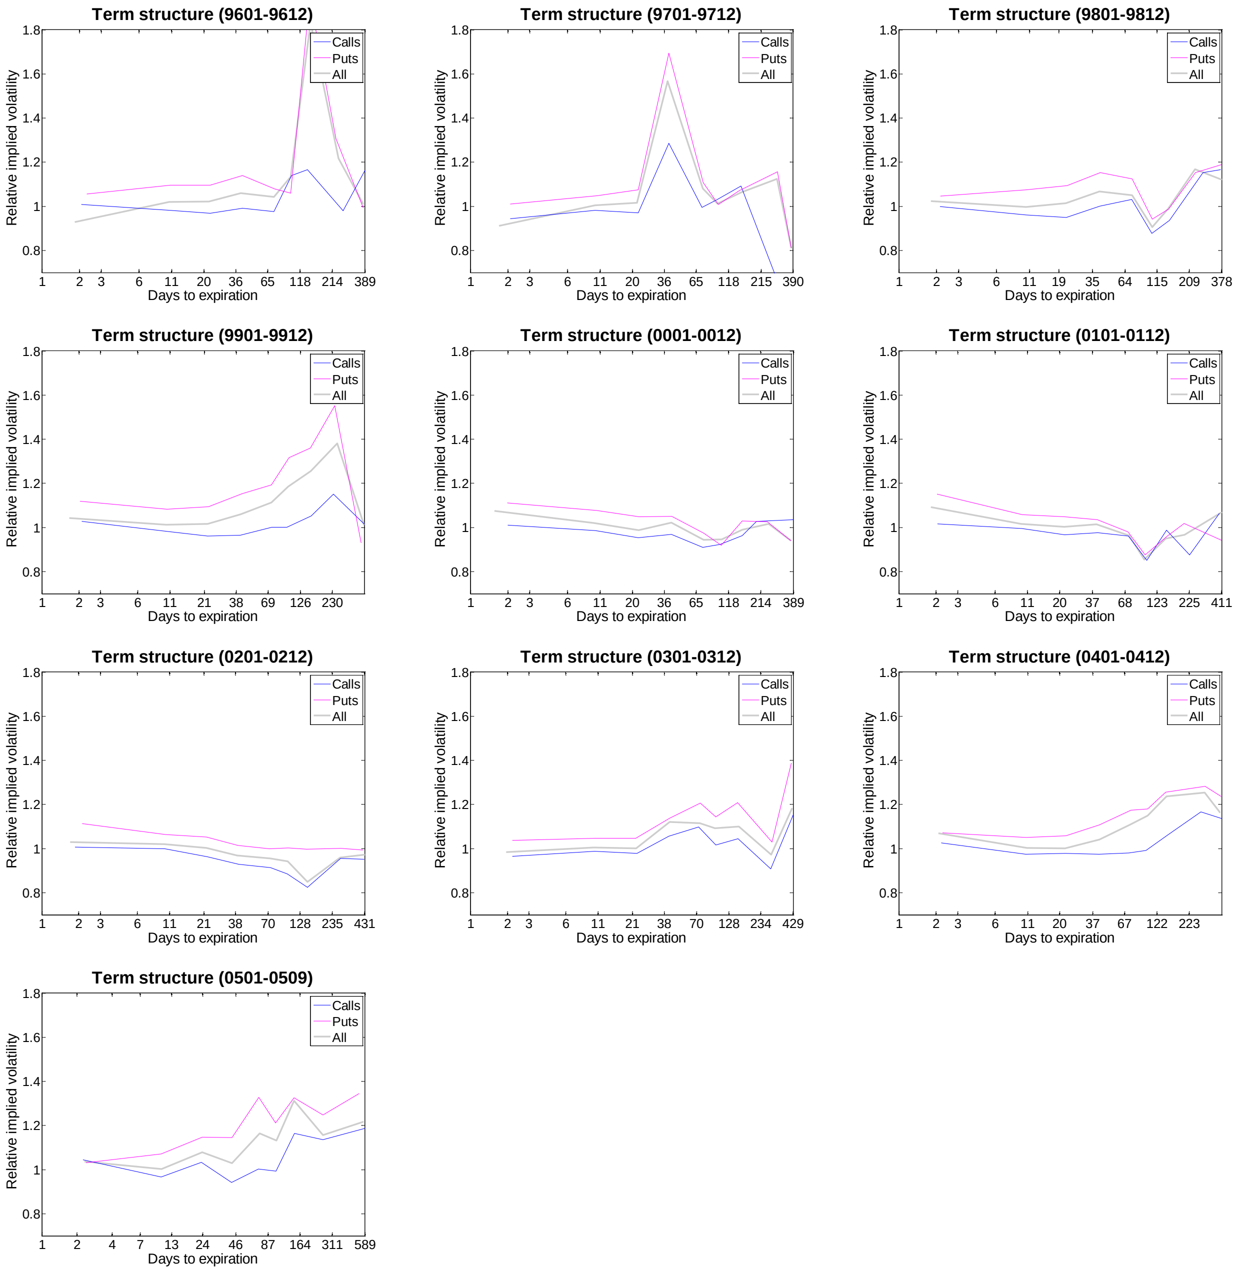


## Figure 10. Volatility surface.

*Note:* Top: Volatility surface, the implied volatility smile separated in time-to-expiration groups, in narrow range, between –5 and 5, to magnify the part of moneyness that matters for the implied distribution. Bottom: Volatility surface in a wider range, between –10 and 10. The legend gives the calendar days to expiration span for the group, with the average number of days in parenthesis. The footer gives the number of observations per term-group. Observation groups with less than ten observations have been removed.

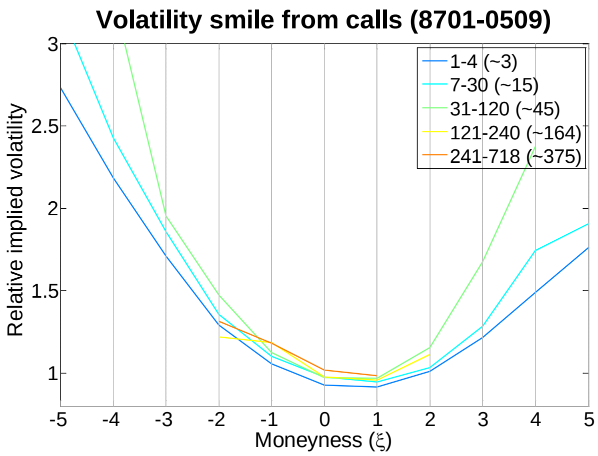


N.o. obs.: 5,283; 33,964; 40,525; 2,091; 1,431

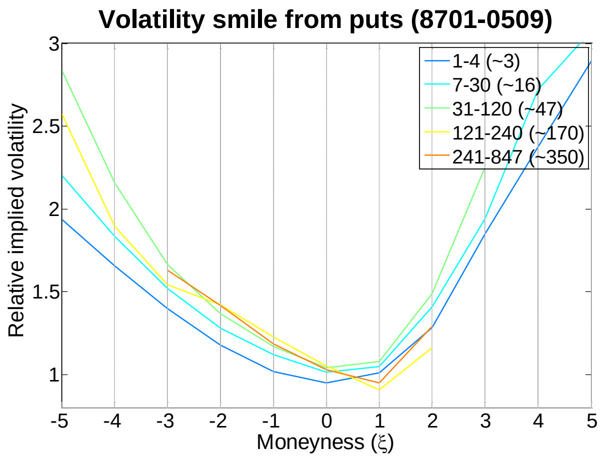


N.o. obs.: 5,419; 37,096; 45,020; 2,993; 1,921

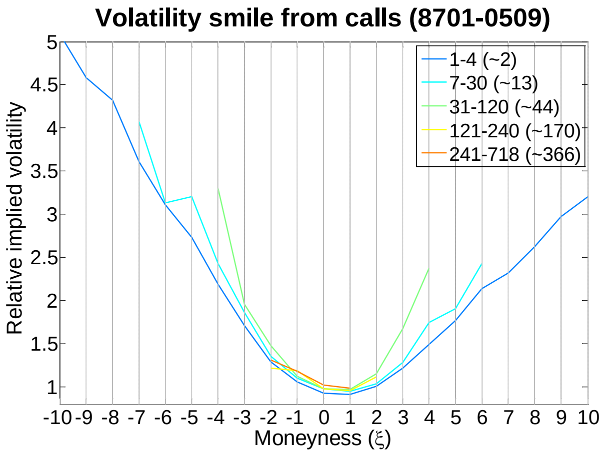


N.o. obs.: 5,822; 34,037; 40,556; 2,101; 1,432

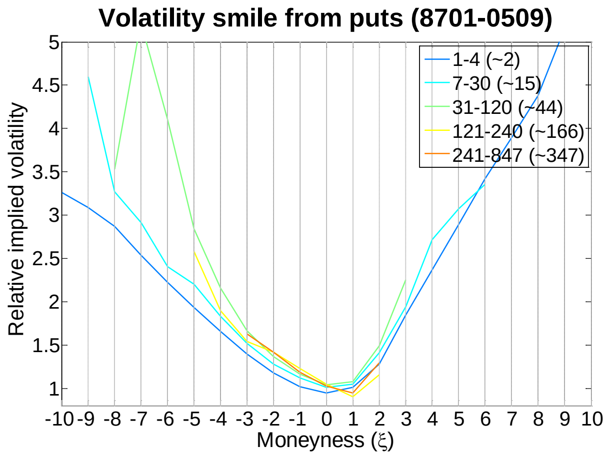


N.o. obs.: 6,141; 37,275; 45,094; 3,006; 1,923