Bid-Ask Spreads and Inventory Risk: Evidence from the FTSE 100 Index Options Market

Rüdiger Fahlenbrach*
The Wharton School
University of Pennsylvania
Philadelphia, PA 19104-6367
E-mail: rfahlenb@wharton.upenn.edu

Patrik Sandås
The Wharton School
University of Pennsylvania
Philadelphia, PA 19104-6367
and
Center for Economic Policy Research
E-mail: sandas@wharton.upenn.edu

First Draft: November 15, 2002
Current Draft: March 14, 2003

*We thank the London International Financial Futures and Options Exchange for providing the sample and the Rodney L. White Center for financial support and Burton Hollifield and seminar participants at the American University, SUNY-Binghamton, and the Wharton Brown Bag seminar for helpful comments.

Abstract

We study the determinants of the bid-ask spreads for index options using a sample that consists of all trades and quotes for the European style options and the futures on the FTSE 100 stock index from August 2001 to July 2002. We compute two measures that show that bid-ask spreads in our sample are economically large. First, we show that the reservation spreads for a liquidity provider who hedges the delta risk of a bought or sold index option by trading index futures is on average only 47% of the observed spread. Second, we show that, on average, the spread for a synthetic index futures contract is 8-14 times the spread for the actual index futures contract even after adjusting for differences in trading volume and contract size. We estimate fixed effects regressions for a panel of options and show that 42% of the variation in the daily average spreads is driven by order processing and inventory costs. Our results show that inventory risk is an important determinant of spreads for index options, but that the compensation for inventory risk together with order processing costs do not explain all systematic variation in the observed spreads. We reject the null that all fixed effects are equal to zero for both calls and puts. The fixed effects vary systematically with option characteristics; for example, out-of-the-money options have wider spreads and in-the-money options have narrower spreads than predicted by our regressions. One interpretation is that the spreads include a mark-up above the marginal costs of providing liquidity for at least some of the options in our sample.

Keywords: Index options; Panel Data; Bid-ask spread; Market making; Inventory models

JEL codes: G13; G14; C33

1 Introduction

Competition and risk-sharing among multiple liquidity providers determine the bid-ask spreads in many financial markets. When bid-ask spreads are determined competitively we associate variation in the spreads with variation in the marginal costs of providing liquidity. The marginal cost of providing liquidity depends on the adverse selection cost, the inventory cost, and the order processing cost. Index options offer an ideal setting for isolating the importance of the last two costs. First, adverse selection costs are of secondary importance since few traders can be expected to be privately informed about the value of the index. Second, options with multiple strike prices and times to maturity and therefore different sensitivities to the underlying index are traded at the same time; the marginal contribution to inventory risk varies across options. Third, the index option contracts are typically among the most actively traded contracts on any options exchange which implies that it is reasonable to assume bid-ask spreads are determined competitively. In such an environment we would predict that the bid-ask spreads depend only on the order processing and inventory costs. But are the observed bid-ask spreads consistent with the above predictions?

We use a sample from the London International Financial Futures and Options Exchange (LIFFE) to address the question. Our sample contains all quotes and trades for FTSE 100 (Financial Times Stock Exchange 100 stock index) European style index options and FTSE 100 index futures between August 2001 and July 2002. Consistent with earlier studies we find that average daily bid-ask spreads, computed from the tick data, vary systematically with strike prices and times to maturity. For call options, the average spreads range from £7.5 (9.2%) to £13.7 (3.4%) for options with less than one month to more than nine months to maturity and strike prices close to the current index value, and from £6.8 (26.0%) to £9.2 (4.7%) for options with one to two months to maturity and strike prices from 7% above to 2% below the current index value.

Are the average spreads for the index options economically large? One plausible determinant of the spreads for any derivative security is the spread of the underlying security. Using this link, we construct two empirical measures to compare the spreads for the options and the futures on the same underlying stock index.

The first measure is based on the reservation spread for a liquidity provider who hedges the delta risk of a bought or sold index option by trading index futures once every day until the option's maturity date. The second measure is the implied spread for a synthetic index futures contract computed using the put-call parity, the bid and ask quotes for options with the same strike price and maturity, and an adjustment for differences in contract size and trading volume. The results show that index option spreads are large given the spreads for the index futures.

The average reservation spreads computed using 10,000 simulations of the paths for the index and the hedging policy are 47% of the observed spreads. On average, the implied spread for the synthetic index futures is 8-14 times the actual spread for the index futures. The first result shows that the expected costs of daily delta hedging–a common hedging policy in practice–only explain a fraction of the observed spreads. The second result shows that the difference in trading volume for index options and index futures cannot explain much of the difference in the spreads for positions with identical exposure to the underlying index.

One plausible explanation for the above results is that the index option spreads contain compensation for costs of providing liquidity that do not affect liquidity providers in the index futures. Consider the following example. A liquidity provider writes a naked call and obtains a delta-neutral position by buying delta units of the index futures. The position is still exposed to gamma risk—the risk of changes in delta. To hedge the gamma risk, the liquidity provider needs to trade in other options since index futures have a gamma of zero. But hedging by trading in the futures and in some options is more costly than trading in the futures alone. It is possible that the liquidity provider chooses not to fully hedge the naked call and to bear some of the delta or gamma risk. Under any hedging policy, it is reasonable to expect the liquidity provider to demand greater compensation for providing liquidity in index options with a larger delta or gamma; since the underlying index is common to all options, the unit costs of hedging delta or gamma risk and the unit compensations for carrying delta or gamma risk are constant across options.

To determine whether compensation for hedging costs and inventory risk can explain the variation in the spreads, we construct a panel of options by sorting the options into a total of 96 strike price and time to maturity categories. By applying panel data methods we can identify systematic differences in the cross-section of spreads. We estimate a cost function for liquidity provision using fixed effects regressions. The different marginal contributions to inventory risk are captured by the options' greeks as well as their implied volatilities. In addition, we include the previous day's option price to capture a potential variable component in the order processing costs and we include trading volume, number of trades, and the open interest as of the previous day's close to capture differences in trading activity which directly affect the inventory costs.

The estimated cost function explains 42% of the variation in the daily average spreads in our panel. The estimated coefficients have the expected signs in most specifications; options with larger deltas, gammas, and vegas and options with greater implied volatility have wider spreads. Overall, we find both economically and statistically significant evidence that bid-ask spreads are determined by order processing and inventory costs.

But our empirical results also suggest that order processing and inventory costs cannot explain all systematic variation in the spreads. In our fixed effects regressions, we strongly reject the null hypothesis of all fixed effects jointly being equal to zero. The rejections imply that the estimated cost function fails to capture some systematic variation in the bid-ask spreads across option categories. The rejections also imply that pooling options with different strike prices and times to maturity into a single regression is likely to produce biased inference.

Our estimated fixed effects are strongly negatively correlated with the predicted spread; the correlation coefficient for the fixed effects and the predicted spread is -0.6 for calls and -0.7 for puts. The observed spreads are narrower than the predicted spreads for in-the-money and longer time to maturity options and wider than the predicted spreads for out-of-the-money and shorter time to maturity options. One interpretation is that we have omitted an inventory risk variable that is correlated with strike prices and time to maturity. An alternative interpretation is that some of the spreads include a mark-up and do not equal the marginal cost of providing liquidity, casting some doubt on our assumption that the spreads are competitively determined.

The empirical evidence on the importance of hedging costs for option spreads is mixed. Jameson and Wilhelm (1992) find that the magnitude of delta has a negative impact on dollar spreads for equity options, and George and Longstaff (1993) find that the squared value of delta has a negative impact on dollar spreads for index options. De Fontnouvelle, Fishe, and Harris (2000) find the opposite delta effect in their sample of equity options. The options' gamma has an insignificant effect on spreads in de Fontnouvelle, Fishe, and Harris (2000), but a positive effect on spreads in Jameson and Wilhelm (1992). Of the three studies only Jameson and Wilhelm (1992) include vega and their results suggest that options with larger vegas have wider spreads. Our results differ from the existing literature in that delta, gamma, and vega risk are jointly significant in absolute spread regressions, both economically and statistically. Accounting for the panel nature of option data seems to be important for the inference about the effect of inventory risk on the spreads.

The empirical market microstructure literature provides several methods for decomposing the bid-ask spread into components related to adverse selection costs, inventory costs, and order processing costs. Huang and Stoll (1997) report that a two-way decomposition assigns on average 11% of the spread to the sum of adverse selection and inventory costs for a sample of large stocks. They also report results for a modified three-way decomposition which attributes on average 28.7% to inventory costs. Our regression results provide a different method for assessing the relative importance of order processing and inventory costs and assigns a higher fraction (42%) of the variation in spreads to order processing and inventory costs.

Many recent studies examine the effect of informed trading on stock option spreads.[1] But trading of security baskets or index-linked securities, like the index options we study, differ in important ways from trading of individual stocks or options on individual stocks. Theoretically, Subrahmanyam (1991) and Gorton and Pennacchi (1993) show that adverse selection costs can be minimized by trading in security baskets or index-linked securities because orders submitted by informed traders in the individual securities tend to offset each other for the security basket or index-linked security. Empirically, Boehmer and Boehmer (2002) study exchange traded funds—an increasingly popular type of index-linked securities—and show that spreads are small and that the fraction of the spreads attributed to adverse selection or inventory costs is only 3-8%. Both the theoretical arguments and the empirical results above motivate our decision to abstract from adverse selection costs.

The remainder of the paper is organized as follows. In Section 2, we describe the market we

[1] Examples include theoretical models developed by Back (1993), Biais and Hillion (1994), and Easley, O'Hara, and Srinivas (1998) and empirical studies by Berkman (1996), Easley, O'Hara, and Srinivas (1998), and Vijh (1990).

study and report some descriptive statistics for our sample. In section 3, we develop two measures for comparing the spreads in the index options and the index futures market and we discuss the explanatory variables and the setup of our panel regressions. In section 4 we present and discuss the empirical results. Section 5 concludes.

2 Description of the Market and the Sample

2.1 The Market

In May 1998, Liffe's members voted to replace open outcry floor trading with an electronic trading system based on a central limit order book. By May 2000, all trading in financial contracts on Liffe was taking place on an electronic limit order book system called Liffe Connect.[2]

In the first half of 2001, the European style FTSE 100 index option contract had an average monthly volume of 1.25 million contracts and an average monthly open interest of over 1.4 million contracts, and it ranks fourth worldwide after the S&P 100 and S&P 500 contracts of the CBOE and the ODAX contract of the Eurex.

FTSE 100 index options are quoted in index points. Each contract is valued at £10 per index point. The tick size is 0.5 index points or £5. When we report absolute spreads in pounds we ignore the index-point multiplier of ten. Strike prices are multiples of 50 index points for contracts with 90 days or less to expiration and multiples of 100 index points for contracts with more than 90 days to expiration.

A new expiration month starts trading with a minimum of eleven strike prices, set around the current index value. Additional strike prices are opened for trading after the index rises above the second highest, or falls below the second lowest, strike price that is currently open.

The standard expiration months are March, June, September and December. Additional expiration months are introduced so that contracts that expire during the three nearest calendar months are always available. The FTSE 100 index futures contract is traded with expiration months that match the three nearest standard expiration months above.

[2]In January 2002, Liffe became a part of Euronext to form a market called Euronext.liffe. Euronext was formed by the merger of the Amsterdam, Brussels and Paris cash and derivatives exchanges in September 2000. All derivatives trading on the different Euronext markets are scheduled to take place on the Liffe Connect trading system in 2003.

The trading hours are 8:00 a.m. to 4:30 p.m. Trader submit orders electronically to the central limit order book. Incoming market orders are automatically matched with orders in the order book to produce trades. Orders are given priority according to price and orders at the same price are filled in proportion to the volume. The exception to the pro rata allocation is that the order with the highest time priority receives a fill that is subject to a minimum and a maximum volume condition. Orders can be canceled at any time and all orders except good 'til cancel orders are canceled automatically at the market close.

Only exchange members can submit orders directly to Liffe Connect. There are no designated market makers with special quoting obligations or privileges in the FTSE 100 index options. In October 2002, there were 143 public order members. Public order members can trade on their own account or on behalf of their customers. There were also 60 non-public order members. Non-public order members can only trade on their own account or on behalf of other members as brokers. New members are accepted provided they meet a set of membership criteria.

Customers may submit orders by calling a member firm, but they may also submit orders without any direct intervention by a member. Customers may set up order routing software to electronically route orders via a member directly to Liffe Connect. Traders may also use software, so called automated price injection models, that automatically generates order submissions.

Information on the best quotes and depths as well as quotes and depth away from the best quotes is distributed in real-time via Liffe Connect to the members' computer screens. Customers who are not members may obtain the same information from quote vendors. In contrast, no information on the identity of members submitting orders is distributed. Trading is anonymous both pre and post-trade.

The buyer and the seller pay a fixed cost of £0.25 per trade. Order submissions and cancelations are free, but there is a fixed fee per message for the automated price injection models which depends on the frequency of usage.

The market for FTSE 100 index options has a large number of market participants who may provide liquidity and therefore act as de facto market makers. Customers who are not members can also make a market relatively easily without any intervention by members. The computerized trading system makes it easy to obtain information about spreads and depths in different contracts and to submit new orders. The computerized trading algorithm ensures that quotes are binding and price priority is strictly enforced. Based on the above features a reasonable starting point for examining the spreads for index options is to assume that competition drives the average spreads to the marginal costs of providing liquidity, which may reflect the compensation liquidity providers demand for order processing and inventory costs.

2.2 The Sample

We obtained our sample from Liffe's market data services. It consists of time-series of all quotes and trades for all the European style FTSE 100 index option contracts and the FTSE 100 index futures contracts. We also obtained the daily time-series of closing prices, open interest, and implied volatilities for all contracts. Our sample period covers 242 trading days between August 1, 2001 and July 30, 2002. A few days were dropped from our sample because the quote and trade records are missing for more than three hours.

For each contract, the sample consists of a time-series of the best bid and ask quotes with the corresponding total depths. A new observation is generated in the quote series for every change in either the best bid or ask quote or in the best bid or ask depth. All quotes are binding since trading is computerized.

We construct a proxy for the risk free interest rate using the overnight, one week, one month, three months and one year Sterling London Interbank Offered Rates (Libor), provided by Datastream. We construct a proxy for the interest rates for maturities other than the ones above via linear interpolation.

We obtain daily dividend yields on the FTSE 100 index from Datastream. Our proxy for the dividend yield implicitly assumes market participants expect that the current dividend yield is an unbiased forecast of the dividend yield during the rest of the life of the option. Given our sample period it is not feasible to follow the alternative approach of using the actual dividends as a proxy for the market participants' expectations.

We construct a daily panel by sorting call and put options into different categories according to the number of days to maturity and the difference between the strike price and the underlying index value; K-S. We define the categories in terms of K-S rather than the moneyness ratio K/S to ensure that at most one contract is included in each category. We refer to the K-S categories as moneyness categories. For calls and puts the panel consists of ten moneyness categories and six maturity categories. The panel covers all times to maturity but only a subset of the strike prices. By restricting the panel to ten moneyness categories we ensure that we have a sufficient number of observations per category.

The range of K/S is 0.98 to 1.07 for the call options and 0.93 to 1.02 for the put options. Both call and put options are sorted into one month, two months, three months, four to six months, seven to nine months, and 10 to 12 months to maturity categories. Table 1 reports the number of observations per category and the average number of days to maturity for each maturity category and the average of K/S for each moneyness category.

The top panels of Tables 2 and 3 report the time-series means of the daily average absolute spreads for call and put options. The smallest mean is £4.9 which is almost ten times the minimum tick size suggesting that price discreteness is not a key determinant of spreads. The spread increases across maturity categories for both calls and puts. Across moneyness categories the pattern is mixed; for maturity categories 1-3, the spreads decrease across moneyness categories for calls and puts with two exceptions for calls in maturity category 2 and 3 for puts in maturity category 1. The pairwise Wilcoxon tests of equal means reject for all adjacent pairs of means formed based on maturity but only for four of nine pairs based on moneyness (p-values less than 0.05).

The bottom panels of Tables 2 and 3 report the time-series means of the daily average relative spreads for call and put options.[3] Spreads decline monotonically across maturity categories with the exception of maturity category 6, and increase monotonically across moneyness categories for maturity categories 1-5, with the exception of puts in maturity category 3. The Wilcoxon tests reject the null of equal means for all adjacent pairs of means formed based on maturity at the 1% level or better with the exception of the pair formed for puts in maturity categories 5 and 6 where the p-value is 0.08. Across moneyness categories the null of equal means is also rejected with p-values less than 0.01 for all but two pairs; the exception are for call and puts in moneyness

[3]The average relative spreads for the FTSE 100 stocks is 0.54% according to Naik and Yadav (1999).

categories 1 and 2.

Overall, the summary statistics in Tables 2 and 3 provide strong evidence of variation in the spreads. Our nonparametric tests of pairs of adjacent means are conservative; tests using pairs of means farther apart reject the null consistently. The general patterns in the absolute and relative spreads across categories are qualitatively similar to the ones reported in Bakshi, Cao, and Chen (2000) for S&P 500 options and Brandt and Wu (2002) for FTSE 100 options. While our definition for the categories differ somewhat from theirs, the patterns we report also appear to be quantitatively consistent with the ones reported in Brandt and Wu (2002) for an earlier sample period.

While the absolute spreads are approximately the same for call and put options in the same category, put options in maturity category 1 have lower relative spreads than call options in the same category with the exception of category 2.

Table 4 reports the time-series average of the implied volatilities calculated from the closing prices of the call and put options in each category. The volatility smirk pattern across moneyness classes in our sample is typical for index options; see, for example, Derman and Kani (1994), Rubinstein (1994), Dumas, Fleming and Whaley (1998), or Brandt and Wu (2002).

Put options in the first maturity category have implied volatilities that increase monotonically from 22.8% for in-the-money options to 31.2% for deep out-of-the-money options. The range decreases across maturity categories and is less than two percent for maturity category 6. Call options in the first maturity category have implied volatilities that decrease monotonically from 24.9% for in-the-money options to 22.7% for deep out-of-the-money options. The range is larger for longer maturities; for maturity category 5 the range is 3.4% and for category 6 it is 2.8%.

Table 5 reports the time-series average of the daily total number of contracts traded in each call and put option category. The total trading volume includes regular, combination, and block trades. Block trades are rare in our sample, but combination trades constitute a substantial fraction of the total trading volume. The trading volume is highest for the at-the-money and out-of-the-money call and put options in the first maturity category. The next highest trading volume is also for options in the same moneyness categories in maturity category four. On average the trading volume for maturity category 1 is more than double the next maturity category.

The top graph in Figure 1 plots the daily closing values for the FTSE 100 index and the bottom graph plots the average implied volatility for options with a strike price within 75 index points of the index value and a time to maturity of 2 to 4 months. The index return over the sample is approximately -25%; the highest index closing value occurs close to the beginning of the sample and the lowest value occurs a few days before the end of the sample. The implied volatility for the at-the-money options is negatively correlated with the level of the index; there is a sharp increase in the implied volatility following September 11th, 2001, and likewise from June 2002 to the end of the sample. Dumas, Fleming, and Whaley (1998) also report a negative correlation between the implied volatility and the index level for the the S&P 500 index.

The top graph in Figure 2 plots the daily average absolute spread measured in pounds and the bottom graph plots the average relative spread for all options with a strike price within 75 index points of the current index value and a time to maturity of 61-120 days. The spreads fluctuate over the sample with some of the widest spread approximately coinciding with the peaks of the implied volatility in Figure 1.

Table 6 reports summary statistics for the spreads and the trading activity in FTSE 100 index futures. The mean spread ranges from £1.4 for short-term contracts to £6.4 for contracts with six to nine months to maturity. Quotes are updated frequently for all three maturities, but the trading activity is extremely skewed with most of the activity in the short-term contracts. The mean number of contracts traded for long maturities is 601 which is comparable to many of the active index option contracts.

The descriptive statistics show that on a typical day options with different strike prices and times to maturity are traded. Both the absolute and relative spreads vary systematically with the strike prices and times to maturity. In addition, the spreads vary over time. Our subsequent analysis focuses on whether or not the observed spreads are large given the spreads we observe in the closely related index futures market and to what extent the variation in the spreads can be attributed to order processing and inventory costs.

3 Methodology

Our methodology is designed to address two questions: Are the observed bid-ask spreads of index options economically large? If so, to what extent can the spreads be rationalized by order processing and inventory costs? The first two subsections describe two empirical measures that exploit the link between the spreads for index options and the spreads for index futures to address the first question. The third subsection presents the panel data methods we apply to address the second question.

3.1 Simulated Spreads for a Delta-Hedged Liquidity Provider

The reservation spread for a liquidity provider who buys or sells one option and then hedges the resulting delta risk every day by trading in the index futures provides one benchmark for the size of the observed spread.

A liquidity provider can hedge the risk of an option that is associated with movements of the index (the option's delta) by taking an offsetting position in an appropriate number of index futures contracts. As the option's delta is not constant but varies with the underlying index level, the liquidity provider should continuously review and modify her hedging decisions. In the presence of transaction costs, such a continuous revision is not feasible. A liquidity provider needs to choose a discrete time interval to hedge the delta risk. Often, market participants choose to have delta-neutral positions overnight.

We simulate the costs of a simple daily delta-hedging program for a single long and a single short option contract in the presence of transaction costs in the index futures market. Each day, the liquidity provider recalculates the delta of the option and adjusts her futures position accordingly. If she trades in the futures contract, she is subject to the bid-ask spread and incurs an order processing cost. We assume that her futures account is marked to market each day and earns the prevailing risk-free interest rate. We then take the difference of the hedging costs for a long position and a short position in the same option contract to be a measure of the reservation bid-ask spread that would arise if hedging the individual option's delta risk was the only cost to providing liquidity.

Our spread estimates are conservative for at least two reasons. First, the simulations assume that the liquidity provider does not aggregate offsetting option positions to reduce the overall hedging costs. Second, we assume that the liquidity provider trades in futures contracts that match the maturity of the option. But it is likely to be cheaper, especially for long-term options, to use a roll-over hedging program with short-term futures contracts that have lower transaction costs.

We assume that the daily movement of the FTSE 100 index can be approximated by geometric Brownian motion. We calculate delta from the Black-Scholes option price formula; see (A4) and (A9) in the appendix. We calculate a theoretical futures price and assume that this price coincides with the midpoint of ask and bid prices.

We calibrate our model using sample information. We use a dividend yield of 2.2%, an annualized risk-free interest rate of 2%, a volatility of the index returns of 21.1%, an expected index return of 6%, and an initial index level of 4,400. The expected index return is, of course, not based on the index return over our sample period. We assume that the transaction costs for trading in the futures market consist of the empirically observed futures bid-ask spreads (see Table 6) and of an order processing cost of £1 per traded futures contract. We ignore divisibility issues and initial and maintenance margins in all simulations.

We simulate bid-ask spreads for a wide cross-section of time to maturity and moneyness categories that coincide with some of the categories of Tables 2 and 3. We simulate call and put options for 3% and 6% out-of-the-money options as well as for at-the-money and 2% in-the money options. Each moneyness category is simulated for 3 different time to maturities for a total of 12 different resrvation bid-ask spreads. Each simulation is run 10,000 times.

3.2 Spreads for Synthetic and Actual Index Futures

The measure of the spread developed above depends on the assumed stochastic process followed by the underlying index and the assumed option pricing model. The measure developed below does not require any such assumptions.

The underlying asset for the index future and the index option contracts is the same. We can therefore compare the spreads in the two markets for a position that has the same exposure to the underlying index and therefore is insensitive to assumptions about the stochastic process or the option pricing model.

Denote the underlying cash index by S and the futures contract by F. Let C and P denote the prices of a call and a put option contract with strike price K. Let r_l denote the lending rate and r_b the borrowing rate. We indicate whether a given price is a bid or ask quote by a different superscript. Assume all contracts have T periods to maturity.

Suppose a market maker engages in the simultaneous purchase and sale of a futures contract. There are no initial cash-flows; at expiration, the market maker obtains a profit of $F^{ask} - S_T + S_T - F^{bid}$. The implied bid-ask spread for the index futures position the market maker can realize today is:

$$e^{-r_b T}\left(F^{ask} - F^{bid}\right). \tag{1}$$

Suppose further a market maker can, with guaranteed execution, buy and sell a synthetic index position through the option and bond markets. A bought or sold synthetic index futures position is constructed using the put-call parity. He immediately realizes a synthetic bid-ask spread of

$$\left[C^{ask} - P^{bid} + e^{-r_l T}K - C^{bid} + P^{ask} - e^{-r_b T}K\right].$$

We need to adjust for the difference in the contract size to make the overall exposure comparable; the futures contract is quoted in £25 per index point whereas the index options are quoted in £10 per index point. To make the adjustment we multiply the required option positions by 2.5.

The bid-ask spread for the synthetic index position is equal to

$$2.5\left[C^{ask} - P^{bid} + e^{-r_l T}K - C^{bid} + P^{ask} - e^{-r_b T}K\right],$$

which simplifies to

$$2.5\left[(C^{ask} - C^{bid}) + (P^{ask} - P^{bid})\right], \tag{2}$$

if we ignore the bid-ask spread in the money market.

Each leg of both positions generates an identical exposure to the index at the maturity date. We therefore use equations (1) and (2) as a way to assess the size of the index option spreads. The comparison is based on hypothetical order flow because there is no guarantee of a simultaneous execution of all necessary index options or index futures.

It is important to note that we do not recur to no-arbitrage arguments to link futures prices to synthetic index prices from options. Such arguments are difficult to make since with transaction costs, we only obtain lower and upper boundaries for both markets. The objective of our comparison is different. We want to compare the gross profit from providing liquidity for hypothetical trades in two different contracts with the same exposure to the underlying index.

3.3 Panel Regressions

3.3.1 Explanatory Variables

Our empirical approach exploits the fact that multiple option contracts with different sensitivities to the underlying index are traded at the same time. Since the costs of hedging the different sensitivities must be the same for all contracts we can measure the sensitivities and determine whether the bid-ask spreads are increasing in the sensitivities. This is our empirical strategy.

For completeness we review the standard Taylor series expansion of the value of an option position and the definitions of the so called greeks that measure the risk characteristics (see, for example, Hull (2000) Ch. 13). We will use the greeks calculated for each option as proxies for the cross-sectional differences in the marginal contributions of different contracts to the market makers' hedging costs and inventory risks.

For simplicity we consider the arguments for a single call option contract with a price C_t at time t which depends on the underlying index value S_t, a volatility of the index σ_t^C, and the time to expiration $T - t$. Ignoring higher order terms, the change in the price of the option over a short interval of time dt is given by :

$$\begin{aligned} dC_t &= \frac{\partial C_t}{\partial S_t} dS + \frac{1}{2} \frac{\partial^2 C_t}{\partial S_t^2} dS^2 + \frac{\partial C_t}{\partial \sigma_t^C} d\sigma + \frac{\partial C_t}{\partial (T-t)} dt \\ &= \Delta_t^C dS + \frac{1}{2} \Gamma_t^C dS^2 + \mathcal{V}_t^C d\sigma + \Theta_t^C dt. \end{aligned} \tag{3}$$

The definitions of the greek letters, $\Delta_t^C, \Gamma_t^C, \mathcal{V}_t^C, \Theta_t^C$, based on the Black-Scholes option values are provided in appendix A. An analogous relationship hold for put options and we denote the greek letter for put options with a P superscript.

In calculating the greeks based on the Black-Scholes prices with different volatilities for different options we are implicitly making an assumption about the correct pricing model. The Black-Scholes model with an ad hoc formulation for the volatility is a popular pricing approach that can match the implied volatility smirk. Dumas, Fleming, and Whaley (1998) and Brandt and Wu (2002) show that using an ad hoc method which smoothes the Black-Scholes implied volatilities across strike prices and times to maturity generates smaller mean-squared pricing errors than a method based on a deterministic volatility function or implied tree approach. Based on their evidence we decided to use Black-Scholes with the ad hoc formulation for the volatility. We assume that the greeks calculated for option prices with or without smoothing the implied volatilities do not differ systematically, and therefore we follow the simpler method of no smoothing.

In managing the risk of an option position a market maker must trade-off the different risk characteristics against the costs of hedging them. By either buying or selling an appropriate number of index futures contracts a position can be made delta-neutral; eliminating the first term in equation 3.

Of course, if the liquidity providers can make an offsetting trade shortly after making the initial trade all risks can be exactly offset and the liquidity provider earns the bid-ask spread. In practice, such perfectly offsetting trades are hard to obtain and therefore liquidity providers monitor the delta, gamma, and vega of their positions; see, for example, the discussion in Hull (2000) Ch. 13.

A market maker who follows a discrete rebalancing strategy is exposed to changes in the option's delta measured by the option's gamma. In order to eliminate the gamma risk, another option must be traded, because the index futures contract has a gamma of zero. The gamma risk is particularly important for at-the-money options with a short time to maturity.

Hedging the risk of changes in the volatility or the vega risk requires a position in another option contract. In practice, market makers monitor the delta, gamma, and vega of their positions but seldom completely hedge the risks because the costs of doing so are too large (see, for example, the discussions in Hull (2000) Ch. 13.10 or Baird (1993)).

Since the underlying asset is the same for all options it is reasonable to assume that the per unit cost of hedging delta, gamma, and vega risks are constant across different option contracts. The cheapest way to hedge delta risk is typically to trade in the index futures contract with the shortest time to maturity. For gamma and vega risk we can think of the costs being determined by the options that are the cheapest to trade per unit of gamma or vega risk. It then follows that in the cross-section hedging costs are increasing in the magnitudes of the option's delta, gamma, and vega. As long as market makers demand more compensation for carrying positions with greater exposure to delta, gamma, or vega risk it follows that their reservation spreads should be increasing in the magnitudes of delta, gamma, and vega.

Theta differs from the other greeks in that it measures the deterministic time decay. Nevertheless, market makers are reluctant to follow trading strategies that leave them long in contracts with large negative thetas (Baird (1993)). Ex ante a market maker does not know whether the next trade will increase or decrease the overall theta of his position; a risk averse market maker may therefore quote wider spreads in options with more negative thetas.

The trade-offs between the greeks and the hedging costs depend on the trading activity in the option contracts. In more actively traded contracts it is easier to eliminate inventory risk simply through trading; the exposure arising from selling a certain option contract is offset by shortly afterwards buying the same contract. Based on inventory models like Ho and Stoll (1983), we expect spreads to be wider in option contracts where the time between order arrivals is longer. We use the option's trading volume, number of trades, as well as the open interest to proxy for variation in the trading activity between contracts. The first two are directly related to trading activity. Options with a larger open interest may also be more actively traded because there are more traders who already hold contracts and who may initiate trades to roll over an existing position to a different strike price or time to maturity option.

An alternative to trading in the same contract is for the market maker to trade in a synthetic contract and the underlying: a so called conversion which involves a long futures contract, a short call, and a long put with the same strikes or a so called reversal which involves a short futures contract and a long call, and a short put with the same strikes. A trading strategy based on conversions and reversals can circumvent many of the inventory risk control problems discussed above. In practice, the usefulness of the strategy is limited by the fact that in-the-money contracts are often very inactive (Baird (1993)).

The above discussion focused on the differences in the hedging costs and risks across contracts.

Over time, changes in the volatility of the underlying asset cause risk averse market makers to post wider spreads (see Ho and Stoll (1981, 1983)). We include each option's implied volatility to capture such effects. Of course, in the cross-section it is inconsistent to have multiple levels of implied volatilities for one underlying asset. But for each contract it may still be rational for a liquidity provider to behave as if each option is on a different asset with a unique volatility.

We include the option price in all regressions to allow for variable order processing costs; additionally, the inclusion of the option price allows us to identify the impact of changes in delta and vega risk controlling for the price of the option.

3.3.2 Implementation

Our empirical approach differs from the empirical approaches taken in earlier studies in that we treat the cross-section and time-series of option spreads as a panel.

We reduce the number of categories from 60 to 48 by combining the moneyness categories for the in-the-money and out-of-the-money options in maturity categories 4-6; we combine moneyness categories 1 and 2, 5 and 6, 7 and 8, and 9 and 10. The sparser moneyness grid for longer maturities matches the rules for introducing new strike prices; for longer maturities strike prices are usually introduced in multiples of 100 index points whereas multiples of 50 points are used for shorter maturities.

Our categories are constructed so that on any day each category contains at most one option. Over the sample, the average number of observations per category in the final panels is 176 for call options and 182 for put options. The combined panel has 96 categories and and average number of 179 observations.

We estimate the following fixed effects model for the absolute spreads for call options[4]:

$$\text{Spread}_{i,t}^{C} = X_{i,t}^{C}\beta + u_i + e_{i,t} \quad i = 1, \dots, 48 \quad t = 1, \dots, 242, \tag{4}$$

where i indexes the maturity and moneyness categories and t indexes the trading days. For each trading day t and category i the vector of explanatory variables $X_{i,t}$ is defined as

$$X_{i,t}^{C} = \left[\Delta_{i,t-1}^{C}\ \Gamma_{i,t-1}^{C}\ \mathcal{V}_{i,t-1}^{C}\ \Theta_{i,t-1}^{C}\ \sigma_{i,t-1}^{C}\ Volm_{i,t-1}^{C}\ Ntrd_{i,t-1}^{C}\ OI_{i,t-1}^{C}\ Prce_{i,t-1}^{C}\right], \tag{5}$$

[4]See, for example, Greene (2000) Ch. 14 for an overview of panel data models.

where σ_{t-1}^{C} denotes the implied volatility that is obtained by solving the Black-Scholes formula for the volatility that sets the theoretical price equal to the market price for the call option in category i, $Volm_{i,t-1}^{C}$ is the daily trading volume for call options in category i at the end of day $t-1$, $Ntrd_{i,t-1}^{C}$ is the number of trades for call options in category i for day $t-1$, and $OI_{i,t-1}^{c}$ denotes the open interest for call options in category i at the end of day $t-1$.

The corresponding fixed effects model and vector of explanatory variables for the put options are given by:

$$\text{Spread}_{i,t}^{P} = X_{i,t}^{P}\beta + u_i + e_{i,t} \quad i = 1, \ldots, 48 \quad t = 1, \ldots, 242, \tag{6}$$

and

$$X_{i,t}^{P} = \left[|\Delta_{i,t-1}^{P}| \;\; \Gamma_{i,t-1}^{P} \;\; \mathcal{V}_{i,t-1}^{P} \;\; \Theta_{i,t-1}^{P} \;\; \sigma_{i,t-1}^{P} \;\; Volm_{i,t-1}^{P} \;\; Ntrd_{i,t-1}^{P} \;\; OI_{i,t-1}^{P} \;\; Prce_{i,t-1}^{P} \right]. \tag{7}$$

We compute the greeks for each option according to equations (A4)-(A12) based on the closing index value, dividend yield, risk free interest rate, and implied volatility at $t-1$. The trading volume, open interest and implied volatilities are obtained directly from the daily closing observations.

The fixed effects regressions are estimated using OLS. To test for the significance of the group effects (the u_i's) we use an F-test, see, for example, Greene (2000), section 14.3.1. We report three different R^2 values for the models; R^2 between refers to the $\left(corr(\overline{X}_i\hat{\beta}, \overline{\text{Spread}_i})\right)^2$; R^2 within refers to the $\left(corr((X_{i,t} - \overline{X}_i)\hat{\beta}, \text{Spread}_{i,t} - \overline{\text{Spread}_i})\right)^2$; and R^2 overall refers to the $\left(corr(X_{i,t}\hat{\beta}, \text{Spread}_{i,t})\right)^2$. The within R^2 tells us how well the cost function can explain the observed spreads once the observations in each group have been demeaned. The overall R^2 tells us how well the cost function and the fixed effects explain the observed spreads.

The panel regressions above are estimated for an unbalanced panel; we have days with missing observations for some of the categories. Robustness checks that require, say, the estimation of residual autocorrelations are complicated by the missing observations. So rather than perform the robustness checks on the unbalanced panels we examine the robustness on a smaller balanced panel. The advantage of the simpler methodology outweighs the disadvantage of less cross-sectional variation in the balanced panel. We use the combined sample of call and put options in order to have a larger number of categories.

We estimate the balanced panel regressions using GLS; we replace $u_i + e_{i,t}$ in equations (4) and (6) with $\epsilon_{i,t}$. We estimate three different specifications: the first specification allows for a single autocorrelation coefficient for the residuals in all groups but no heteroscedasticity; the second specification allows for different autocorrelation coefficients for each group and no heteroscedasticity; and the third specification allows for both group specific autocorrelation and heteroscedasticity in the residuals.

4 Empirical Results

4.1 Simulated Reservation Spreads

Table 7 presents the reservation spreads for a liquidity provider who hedges daily the delta risk of a bought or sold index option by trading in the index futures. The top four rows of each panel report the average reservation spreads for an option in a given moneyness and maturity category across the 10,000 simulations. The bottom four rows of each panel report the ratios of the observed spreads, reported in Tables 2 and 3, to the simulated reservation spreads.

The variation in the simulated spreads across moneyness and maturity categories match that of the observed spreads. The spreads are increasing in time to maturity and are decreasing the further the option moves out of the money.

For the short-term options in maturity category 2, the observed bid-ask spreads are substantially larger than the simulated bid-ask spreads. The ratio of observed to simulated bid-ask spreads for the in the money put and call option is more than 3 to 1, and the ratio for the 6% out of the money options is more than 4 to 1. While the ratios decrease with longer maturity, the observed bid-ask spreads are always larger than the simulated spreads. The decline in the ratios is partly due to our conservative assumption about which index futures contract is used for hedging. If all hedging was done in the closest to maturity index futures and rolled over we expect larger ratios for the longer time to maturity options even after imposing some rollover cost. Overall, the results imply that expected costs of hedging delta-risk alone fail to explain the observed bid-ask spreads.

4.2 Spreads for Synthetic and Actual Index Futures

The first row of Table 8 reports the mean bid-ask spread for the futures contracts for three different time to maturity categories. We have omitted the discounting in equation (1). Discounting would make the futures bid-ask spread smaller. The second row reports the bid-ask spread for the synthetic index futures position ignoring the effects of different borrowing and lending rates.

The mean synthetic spread is about 27 times the mean actual spread for the shortest maturity, the ratio for the longest time to maturity category is about nine to one.

The narrower spreads in the index futures market may be explained by greater trading activity; the near month contract has a daily trading volume of over fifty-thousand contracts compared to less than one thousand for the most active option contract. But trading activity alone cannot explain the difference.[5] To quantify the effect of differences in trading volume we estimated a regression of option spreads on the logarithm of trading volume. We use the estimated slope coefficient on trading volume to predict the different option spreads for trading volume that match the index futures. The volume adjusted spreads and the ratio of the implied synthetic spreads and the actual spreads are reported on the last two rows.

Both measures of the size of the bid-ask spread therefore imply that option bid-ask spreads are economically large. One interpretation of our results is that inventory risk may cause wider spreads than our adjustment for expected cost of delta-hedging and differences in trading volume would imply. In our subsequent empirical work we try to determine how much of the spreads can be explained by standard order processing costs and more detailed proxies for inventory risk.

4.3 Panel Regression Results

Table 9 contains descriptive statistics for the explanatory variables we use in the fixed effects regressions. Note that the absolute delta of the put options is smaller than that of the call options; the put options of our sample are more out-of-the-money than the call options. Due to the higher implied volatility of the put options relative to the call options, the average price of put options is higher than that of call options. Put options are traded in higher volumes than call options and

[5] The spreads and trading volume are, of course, interdependent. Since our goal is to obtain an approximate correction for difference in trading volume we ignore this issue.

have higher open interest.

Table 11 reports the results of the fixed effects models in equations (4) and (6) applied to the call and put option panels for the absolute spreads. The first column reports results for the call options, the second column reports results for the put options and the final column reports results for the combined panel of call and put options. For each variable, the table contains the parameter estimate, and below the parameter estimate the corresponding t-statistic.

For both calls and puts the parameter on absolute delta is positive and statistically significant which is consistent with our predictions. A positive coefficient implies that absolute spreads are larger for contracts whose prices are more sensitive to changes in the index. The magnitude of the parameter is large: A one standard deviation increase in the call's delta increases the absolute spreads by £0.75. The corresponding change in the spread for put options is £0.56.

The parameters on gamma are positive and significantly different from zero for call options; call options that exhibit greater convexity have wider spreads. The parameters in the call equation are seven times larger than the ones in the put equation suggesting that call option spreads are more sensitive to the discrete rebalancing risks. While call options are traded less frequently, which would tend to force market makers to hold call option positions longer holding everything else equal, the difference seems too large to be explained by that.

The parameter estimates for vega are all positive and significantly different from zero; options that are more sensitive to changes in the volatility have wider spreads. Put option spreads are more sensitive to changes in volatility than spreads of call options. A one standard deviation higher vega increases the absolute spreads of put options by £1.23.

The parameter on theta for the call option spread is negative and significantly different from zero; since theta is negative for all observations the result implies that call options with greater time decay have wider spreads. For put options, the pattern is reversed; put options with a smaller time decay have wider spreads. For call and put options together we obtain a negative coefficient on theta.

The implied volatility parameters are positive and significant; all spreads widen when the level of implied volatility increases; the effect is consistent with Proposition 1 in Ho and Stoll (1983).

A one standard deviation increase in the implied volatility is associated with a £1.29 increase in bid-ask spreads for calls and a £1.01 increase in the bid-ask spreads for puts.

Three explanatory variables proxy for the trading activity: trading volume, number of trades, and open interest. While all of these variables enter with negative coefficients as an inventory model would predict, the trading volume seems to have the smallest effect on absolute spreads. The number of trades has negative and statistically significant coefficients. The higher the trading frequency, the more likely it is for a market maker to find an offsetting counterparty. The parameter on open interest is negative and significantly different from zero; options with larger open interest have narrower spreads. Increasing the open interest by one standard deviation decreases the absolute spread by £0.16 to £0.22.

The option's price has a positive and statistically significant impact on absolute spreads consistent with a variable order processing cost. A one standard deviation higher option price increases absolute spreads by approximately £3. The estimated constant is £1.9 for call options and £1.2 for put options, and for the combined regression the constant is £1.1. The constant is larger than the fixed order fee for two transactions of £0.50. The larger fixed costs imply that there must be either substantial additional order processing costs per order or positive average profit margins or some combination of the two.

The table reports F-tests for the null hypotheses of all parameters being jointly equal to zero below the top panel. The tests reject the null for all three specifications.

The overall R^2's are between 52.6% and 54.7%. The within-R^2 is 43% for call option spreads and 41% for put option spreads. The between-R^2, obtained from a regression on group means which removes time-series variation, is over 85% for call spreads and over 90% for put spreads.

The difference between the overall R^2 and the within R^2 implies that the fixed effects regression captures information that would be ignored in a pooled cross-section and time-series regression. The F-tests reported in Table 12 reject the null hypotheses that the fixed effects are jointly equal to zero for the call and put options. The test implies that a pooled cross-section and time-series regression for the bid-ask spread leads to biased inference.

The fixed effects in Table 12 vary systematically across strike prices and times to maturity.

Recall from Tables 2 and 3 that the bid-ask spreads vary systematically across strike prices and times to maturity. The correlation between fixed effects and the predicted spread is -0.61 for call and -0.70 for put options. The negative correlation implies that the model underestimates the bid-ask spreads for out-of-the-money options and overestimates the bid-ask spreads for in-the-money options; order processing cost and inventory risk cannot explain all the observed variation in bid-ask spreads for either call or put options.

Table 13 reports the estimation results for alternative specifications of the panel regressions that allow for autocorrelation and heteroscedasticity in the residuals. The signs of the parameters are identical to the ones reported in Table 11. The magnitudes of the parameters are comparable for all variables but delta; for delta the estimated parameters are approximately 1.3 versus 4.5 in the earlier regression. Since the parameter estimates are similar across the three specifications, the results do not appear to be driven by autocorrelation or heteroscedasticity alone. In constructing the balanced panel we have focused on a narrower range of strike prices and times to maturity to ensure that the panel is balanced. A consequence is that we have less dispersion in the explanatory variables, which may explain the change in the magnitude of the parameter for delta. Overall, the results for the alternative specifications show that our key findings reported in Table 11 are robust to autocorrelation and heteroscedasticity.

5 Conclusions

We study the determinants of the bid-ask spreads for index options using a sample that consists of all trades and quotes for the European style options and the futures on the FTSE 100 stock index from August 2001 to July 2002. We address two main questions: Are the observed spreads for index options economically large? If so, to what extent can order processing costs and inventory costs rationalize the observed spreads?

We compute two measures of the size of the spreads that exploit the idea that the spreads for a derivative asset are partly determined by the spreads for the underlying asset. The first measure is the reservation spread for a liquidity provider who quotes a spread that covers the expected costs of delta-hedging a bought or sold index option by trading in the index futures once a day until the option's maturity date. The second measure is the implied spread for a synthetic index futures contract constructed to have the same exposure to the underlying index as the actual index futures. The average reservation spreads are approximately twice the observed spreads. The synthetic spreads are on average 8-14 times the actual spreads even after adjusting for difference in trading volume. Overall, the index options spreads appear to be economically large underscoring the need to better understand how the spreads are determined.

To determine to what extent inventory risk determines the bid-ask spreads we estimate fixed effects regressions with explanatory variables that proxy for differences in the options' marginal contributions to inventory risk. The within R^2 is 42% in our fixed effects regressions implying that inventory risk and order processing cost explain a substantial fraction of the variation in the average daily bid-ask spreads. Nevertheless, we strongly reject the null that the estimated fixed effects are equal to zero both for call and put options separately and jointly; order processing and inventory costs cannot explain all variation in the observed bid-ask spreads.

The estimated fixed effects account for approximately 23% of the unexplained variation in the spreads. The estimated fixed effects are strongly negatively correlated with the predicted bid-ask spreads, which implies that according to the estimated cost function the bid-ask spreads are too wide for out-of-the-money options and the bid-ask spreads are too narrow for in-the-money options. One interpretation is that the fixed effects capture the impact of an omitted inventory risk variable that is correlated with strike prices and times to maturity. An alternative interpretation is that our variables capture all inventory risk related effects and that the bid-ask spreads for out-of-the-money options are non-competitive.

Recent studies by de Fontnouvelle, Fishe, and Harris (2000) and Mayhew (2002), revisit the question of the competitiveness of equity option market making first studied by Neal (1987, 1992). The new evidence shows that spreads are narrower for options traded on multiple exchanges than for options traded on a single exchange even though on every exchange multiple liquidity providers compete in setting bid and ask quotes. The index options that we study also trade at a single exchange and feature a large number of potential liquidity providers. It seems reasonable that the same economic forces that lead to imperfect competition in the equity option markets may be relevant for index options. To determine whether or not imperfect competition can explain some fraction of the observed spreads one would need a more detailed model of the quote setting behavior that can generate testable restrictions for the options spreads. Theoretical and empirical work on such a question is a useful direction for future work.

References

Back, K., 1993, Asymmetric Information and Options, *Review of Financial Studies* 6, 435-472.

Baird, A., 1993, *Option Market Making: Trading and Risk Analysis for the Financial and Commodity Option Markets*, Wiley, New York, N.Y.

Bakshi, G., C. Cao, and Z. Chen, 2000, Do Call Prices and the Underlying Stock Always Move in the Same Direction?, *Review of Financial Studies* 13, 549-584.

Berkman, H., 1996, Large Option Trades, Market Makers, and Limit Orders, *Review of Financial Studies*, 9, 977-1002.

Biais, B. and P. Hillion, 1994, Insider and Liquidity Trading in Stock and Options, *Review of Financial Studies*, 7, 743-780.

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

Boehmer, B and E. Boehmer, 2002, Trading your neighbor's ETFs: Competition or fragmentation?, working paper, New York Stock Exchange.

Brandt, M., and T. Wu, 2002, Cross-sectional Tests of Deterministic Volatility Functions, *Journal of Empirical Finance* 9, 525-550.

De Fontnouvelle, P., R.P. Fishe, and J. Harris, 2000, The Behavior of Bid-Ask Spreads and Volume During the Competition for Option Listings in 1999, forthcoming, *Journal of Finance*.

Derman, E. and I. Kani, 1994, The Volatility Smile and Its Implied Tree, Quantitative Strategies Research Notes, Goldman Sachs, New York.

Dumas, B., J. Fleming, and R. Whaley, 1998, Implied Volatility Functions: Empirical Tests, *Journal of Finance,* 53, 2059 - 2106.

Easley, D., M. O'Hara, and P.S. Srinivas, 1998, Option Volume and Stock Prices: Evidence from Where Informed Traders Trade, *Journal of Finance*, 53, 431-465.

George, T. and F. Longstaff, 1993, Bid Ask Spreads and Trading Activity in the S&P 100 Index Options Market, *Journal of Financial and Quantitative Analysis*, 28, 381-397.

Gorton, G. and G. Pennacchi, 1993, Security Baskets and Index-linked Securities, *Journal of Business*, 66, 1-27.

Greene, W., 2000, *Econometric Analysis*, Prentice Hall, Upper Saddle River, NJ.

Ho, T. and H. Stoll, 1981, Optimal Dealer Pricing under Transactions and Return Uncertainty, *Journal of Financial Economics*, 9, 47-73.

Ho, T. and H. Stoll, 1983, The Dynamics of Dealer Markets under Competition, *Journal of Finance*, 38, 1053-1074.

Huang, R. and H. Stoll, 1997, The Components of the Bid-Ask Spread: A General Approach, *Review of Financial Studies* 10, 995-1034.

Hull, J., 2000, *Options Futures and Other Derivatives*, Prentice-Hall, Upper Saddle River, N.J.

Jameson, M. and W. Wilhelm, 1992, Market Making in Options Markets and the Costs of Discrete Hedge Rebalancing, *Journal of Finance*, 47, 765-779.

LIFFE trading statistics, downloaded from http://www.liffe.com/liffedata/contracts/monthly.csv on August 15 2002.

LIFFE: How the Market Works, a guide to LIFFE's electronic market, April 2002, downloaded from http://www.liffe.com on October 30, 2002.

Mayhew, S., 2002, Competition, Market Structure and Bid-Ask Spreads in Stock Option Markets, *Journal of Finance*, 57, 931-958.

Naik, N. and P. Yadav, 1999, The Effects of Market Reform on Trading Costs of Public Investors: Evidence from the London Stock Exchange, working paper, London Business School.

Neal, R., 1987, Potential Competition and Actual Competition in Equity Options, *Journal of Finance*, 42, 511-531.

Neal, R., 1992, A Comparison of Transaction Costs in Specialist and Open-outcry Markets, *Journal of Business*, 65, 317-334.

Rubinstein, M., 1994, Implied Binomial Trees, *Journal of Finance*, 49, pp. 771–818.

Subrahmanyam, A., 1991, A Theory of Trading in Stock Index Futures, *Review of Financial Studies*, 4, 17-51.

Vijh, A., 1990, Liquidity of the CBOE Equity Options, *Journal of Finance*, 45, 1157-1179.

A Definition of the Greek Letters

For completeness we provide the Black-Schoels values for call and put options as well as the definitions of the greek letters that we use.

The Black-Scholes value of the call is given by:

$$C_t = S_t e^{-q(T-t)} N(d_1) - K e^{-r(T-t)} N(d_2), \tag{A1}$$
$$\tag{A2}$$

where

$$d_1 = \frac{ln(\frac{S_t}{K}) + (r - q + \frac{\sigma_t^2}{2})(T-t)}{\sigma_t^C \sqrt{T-t}} \quad \text{and} \quad d_2 = d_1 - \sigma_t \sqrt{T-t}. \tag{A3}$$

where K denotes the strike price, r denotes the risk-free interest rate, q denotes the continuous dividend yield, and $N(x)$ denotes the cumulative standard normal distribution function evaluated at x. The greeks in equation 3 are then given by:

$$\Delta_t^C = \frac{\partial C_t}{\partial S_t} = e^{-q(T-t)} N(d_1) \tag{A4}$$

$$\Gamma_t^C = \frac{\partial^2 C_t}{\partial S_t^2} = \frac{N'(d_1) e^{-q(T-t)}}{S_t \sigma_t \sqrt{T-t}} \tag{A5}$$

$$\mathcal{V}_t^C = \frac{\partial C_t}{\partial \sigma_t^C} = S_t \sqrt{T-t} N'(d_1) e^{-q(T-t)} \tag{A6}$$

$$\Theta_t^C = \frac{\partial C_t}{\partial (T-t)} = -\frac{S_t e^{-q(T-t)} N'(d_1) \sigma_t^C}{2\sqrt{T-t}} - rKe^{-r(T-t)} N(d_2) + qS_t N(d_1) e^{-q(T-t)} \tag{A7}$$

where $N'(x)$ is the probability density function of the standard normal distribution, evaluated at x.

The corresponding Black-Scholes value of the put option P_t is given by:

$$P_t = Ke^{-r(T-t)} N(-d_2) - S_t e^{-q(T-t)} N(-d_1), \tag{A8}$$

and the the greeks for the put option are given by:

$$\Delta_t^P = \frac{\partial P_t}{\partial S_t} = e^{-q(T-t)} (N(d_1) - 1) \tag{A9}$$

$$\Gamma_t^P = \frac{\partial^2 P_t}{\partial S_t^2} = \frac{N'(d_1) e^{-q(T-t)}}{S_t \sigma_t^P \sqrt{T-t}}, \tag{A10}$$

$$\mathcal{V}_t^P = \frac{\partial P_t}{\partial \sigma_t^P} = S_t \sqrt{T-t} N'(d_1) e^{-q(T-t)} \tag{A11}$$

$$\Theta_t^P = \frac{\partial P_t}{\partial (T-t)} = -\frac{S_t e^{-q(T-t)} N'(d_1) \sigma_t^P}{2\sqrt{T-t}} + rKe^{-r(T-t)} N(-d_2) - qS_t N(-d_1) e^{-q(T-t)} \tag{A12}$$

Table 1: **Number of Daily Observations**

Call Options

Maturity Category			1	2	3	4	5	6
Days to Maturity			2-30	31-60	61-90	91-180	181-270	271-360
Average days to maturity			17.06	46.65	75.88	129.30	222.03	310.36
Moneyness Category	K - S	Average K/S						
1	-125 to -75	0.98	117	98	60	57	36	14
2	-75 to -25	0.99	235	209	155	100	78	48
3	-25 to 25	1.00	241	235	186	131	90	48
4	25 to 75	1.01	242	239	204	127	94	57
5	75 to 125	1.02	242	240	197	135	90	50
6	125 to 175	1.03	240	243	212	122	81	56
7	175 to 225	1.04	240	240	203	125	84	46
8	225 to 275	1.05	231	238	199	121	75	48
9	275 to 325	1.06	212	235	188	128	86	47
10	325 to 375	1.07	204	231	199	119	76	48

Put Options

Maturity Category			1	2	3	4	5	6
Days to Maturity			2-30	31-60	61-90	91-180	181-270	271-360
Average days to maturity			16.76	46.55	76.08	129.93	221.66	311.55
Moneyness Category	K - S	Average K/S						
1	75 to 125	1.02	112	106	78	59	41	27
2	25 to 75	1.01	240	227	195	123	86	60
3	-25 to 25	1.00	241	241	209	141	105	69
4	-75 to -25	0.99	242	240	218	121	94	71
5	-125 to -75	0.98	242	238	199	133	86	63
6	-175 to -125	0.97	241	237	201	122	78	63
7	-225 to -175	0.96	242	231	188	132	88	59
8	-275 to -225	0.95	239	227	187	121	80	57
9	-325 to -275	0.94	237	227	189	132	83	50
10	-375 to -325	0.93	229	230	180	117	81	53

The table reports the number of daily observations of quoted option prices sorted into ten moneyness and six maturity categories. Each row represents a moneyness category defined by an interval for the difference between the strike price K and the underlying index value S. Each column represents a maturity category defined by a range for the option's number of days to maturity.

Table 2: **Average Daily Spreads - Call Options**

Moneyness Category	Absolute Spreads [£] Maturity Category 1	2	3	4	5	6	Average of 1-6	Wilcoxon Test p-value
1	9.0	9.7	10.7	11.8	12.2	12.8	10.3	—
2	8.0	9.4	10.9	11.7	11.6	11.9	9.9	0.01
3	7.5	8.8	10.1	10.9	12.5	13.7	9.6	0.08
4	6.5	8.2	9.5	10.6	11.3	12.0	8.9	0.00
5	5.9	7.8	9.3	10.6	12.0	14.1	8.7	0.04
6	5.4	7.4	9.0	10.4	10.9	12.4	8.2	0.00
7	5.3	7.2	8.6	9.8	12.2	14.6	8.2	0.32
8	5.1	7.1	8.3	9.6	11.0	12.0	7.8	0.07
9	5.1	7.0	8.0	9.5	11.8	14.4	8.0	0.48
10	4.9	6.8	7.8	9.4	10.6	11.9	7.5	0.08
Average of 1-10	6.1	7.8	9.0	10.3	11.6	13.0		
Wilcoxon Test p-value	—	0.00	0.00	0.00	0.00	0.00		

Moneyness Category	Relative spreads [%] Maturity Category 1	2	3	4	5	6	Average of 1-6	Wilcoxon Test p-value
1	6.7	4.7	4.5	3.8	3.1	3.0	5.0	—
2	7.4	5.2	4.9	4.3	3.2	2.8	5.3	0.04
3	9.2	5.6	5.1	4.3	3.6	3.4	5.9	0.00
4	12.2	6.1	5.5	4.7	3.7	3.3	6.9	0.00
5	17.9	7.1	6.0	5.1	4.0	4.0	8.9	0.00
6	27.0	8.5	7.0	5.8	4.1	4.0	11.8	0.00
7	39.1	10.7	8.1	6.3	4.9	5.0	16.0	0.00
8	48.9	14.2	9.8	7.1	5.2	4.6	19.8	0.00
9	58.2	19.4	11.8	8.1	5.9	5.8	23.3	0.00
10	68.6	26.0	14.3	9.7	6.3	5.4	28.2	0.00
Average of 1-10	29.7	11.2	8.0	6.0	4.4	4.2		
Wilcoxon Test p-value	—	0.00	0.00	0.00	0.00	0.00		

The table reports the time series average daily absolute and relative spreads for the call options by maturity and moneyness category. The second row of the last column in each panel reports the p-value for a nonparametric Wilcoxon test of equal means for moneyness categories 1 and 2 and the rows below report the corresponding p-values for tests for moneyness category pairs i and i+1 ($2 \leq i \leq 9$). The last row of each panel reports the corresponding p-values for tests of equal means for maturity categories pairs j and j+1 ($1 \leq j \leq 5$).

Table 3: **Average Daily Spreads - Put Options**

Moneyness Category	Absolute Spreads [£] Maturity Category 1	2	3	4	5	6	Average of 1-6	Wilcoxon Test p-value
1	8.8	9.5	11.4	10.9	11.4	14.5	10.3	—
2	8.2	9.0	9.9	10.7	11.2	12.4	9.6	0.00
3	7.4	8.2	9.3	10.2	12.0	13.8	9.3	0.00
4	6.6	7.8	9.0	10.3	10.9	11.8	8.6	0.00
5	6.0	7.5	8.8	10.2	11.1	13.6	8.5	0.04
6	5.9	7.2	8.8	10.3	11.0	12.7	8.2	0.15
7	5.8	7.2	8.6	9.4	11.4	13.7	8.2	0.38
8	5.7	7.2	8.4	10.1	10.5	11.9	8.0	0.37
9	5.8	7.1	8.4	9.2	11.5	13.0	8.0	0.98
10	5.6	6.9	8.3	9.5	10.5	11.6	7.8	0.31
Average of 1-10	6.5	7.7	9.0	10.0	11.2	12.8		
Wilcoxon Test p-value	—	0.00	0.00	0.00	0.00	0.00		

Moneyness Category	Relative spreads [%] Maturity Category 1	2	3	4	5	6	Average of 1-6	Wilcoxon Test p-value
1	6.6	4.9	4.9	4.0	3.3	3.9	5.0	—
2	7.4	5.1	4.6	4.1	3.5	3.5	5.2	0.51
3	8.7	5.3	4.8	4.3	3.9	4.0	5.6	0.00
4	10.6	5.7	5.1	4.7	3.9	3.7	6.3	0.00
5	14.0	6.4	5.7	5.0	4.2	4.5	7.7	0.00
6	19.0	7.2	6.4	5.6	4.6	4.6	9.5	0.00
7	25.2	8.3	7.0	5.7	5.0	5.1	11.5	0.00
8	31.7	10.0	8.0	6.7	5.1	5.0	14.1	0.00
9	38.3	11.7	8.8	6.7	5.7	5.8	16.4	0.00
10	42.5	13.1	9.7	7.6	5.9	5.5	18.1	0.00
Average of 1-10	21.0	7.9	6.5	5.5	4.5	4.6		
Wilcoxon Test p-value	—	0.00	0.00	0.00	0.00	0.08		

The table reports the time series average daily absolute and relative spreads for the put options by maturity and moneyness category. The second row of the last column in each panel reports the p-value for a nonparametric Wilcoxon test of equal means for moneyness categories 1 and 2 and the rows below report the corresponding p-values for tests for moneyness category pairs i and i+1 ($2 \leq i \leq 9$). The last row of each panel reports the corresponding p-values for tests of equal means maturity categories for pairs j and j+1 ($1 \leq j \leq 5$).

Table 4: **Average Daily Implied Volatility**

Call Options [%]

Moneyness Category	Maturity Category 1	2	3	4	5	6	Average of 1-6
1	24.9	23.5	21.8	22.7	22.3	21.1	22.7
2	24.3	23.0	22.2	21.7	20.9	20.4	22.1
3	24.1	22.7	21.6	21.3	21.0	20.3	21.8
4	23.3	22.4	21.7	21.5	20.7	20.3	21.7
5	22.8	22.0	21.2	20.9	20.6	20.4	21.3
6	22.4	21.5	21.3	20.6	20.1	19.5	20.9
7	22.2	21.1	20.6	20.1	19.9	19.4	20.6
8	22.1	20.6	20.2	20.3	19.4	18.7	20.2
9	22.4	20.4	19.6	19.7	19.5	18.9	20.1
10	22.7	20.1	19.7	19.5	18.9	18.3	19.9
Average of 1-10	23.1	21.7	21.0	20.8	20.3	19.7	21.1

Put Options [%]

Moneyness Category	Maturity Category 1	2	3	4	5	6	Average of 1-6
1	22.8	22.0	21.3	20.7	20.3	19.9	21.2
2	23.2	22.2	21.7	21.4	20.7	19.9	21.5
3	23.9	22.7	21.9	21.5	21.2	20.7	22.0
4	24.6	23.3	22.6	22.0	21.1	20.7	22.4
5	25.5	23.8	23.0	22.3	21.7	21.1	22.9
6	26.5	24.3	23.3	22.7	21.6	21.2	23.3
7	27.6	24.8	23.8	22.9	22.3	21.8	23.9
8	28.8	25.3	24.0	23.6	22.2	21.1	24.2
9	29.9	26.0	24.7	23.8	22.7	21.5	24.8
10	31.2	26.8	25.3	24.3	22.9	21.7	25.4
Average of 1-10	26.4	24.1	23.2	22.5	21.7	21.0	23.1

The table reports the time series average of daily implied volatilities for the call and put options by maturity and moneyness category. The implied volatility for a contract is calculated each day by solving the Black-Scholes formula for the volatility that sets the theoretical price equal to the market price. The seventh column and the eleventh row in each panel reports the average implied volatility across the maturity categories and across the moneyness categories.

Table 5: **Average Daily Trading Volume**

Moneyness Category	Call Options Maturity Category 1	2	3	4	5	6	Average of 1-6
1	211	129	117	95	131	57	123
2	397	188	156	345	251	329	278
3	604	257	176	319	395	199	325
4	675	272	249	549	270	318	389
5	714	288	176	315	223	365	347
6	637	249	152	270	208	86	267
7	616	353	184	251	184	148	289
8	519	230	194	334	261	108	274
9	743	308	187	272	347	90	325
10	424	282	219	380	319	70	282
Average of 1-10	554	256	181	313	259	177	290

Moneyness Category	Put Options Maturity Category 1	2	3	4	5	6	Average of 1-6
1	363	141	79	312	235	306	239
2	338	152	80	288	306	204	228
3	651	299	156	354	236	107	301
4	741	243	154	318	273	212	324
5	745	265	176	300	290	130	318
6	718	386	119	451	236	79	332
7	814	328	271	457	260	77	368
8	836	258	221	374	353	197	373
9	680	344	211	382	370	145	355
10	691	241	154	353	190	78	285
Average of 1-10	658	266	162	359	275	154	312

The table reports the time series average total daily volume for the call and put options by maturity and moneyness category measured in number of contracts. The volume includes regular, combination and block trades. The seventh column and the eleventh row in each panel reports the average trading volume across the maturity categories and across the moneyness categories.

Table 6: **Index Futures - Summary Statistics**

	Days to Maturity		
	0–90	91–180	181–270
Mean number of days to maturity	43.6	134.4	226.1
Mean daily spread [£]	1.4	2.8	6.4
	(0.3)	(1.2)	(4.5)
Mean daily trading volume	50,117	7,927	601
	(27,195)	(17,941)	(1,026)

The table reports summary statistics for the FTSE 100 index futures contracts with different times to maturity. The mean number of days to maturity is reported in the first row. The mean absolute daily spreads are reported on the second row with the standard deviations in parentheses. The mean daily trading volume with standard deviations in parentheses are reported on rows four and five.

Table 7: **Simulated Reservation Spreads for a Delta-Hedged Liquidity Provider**

	Moneyness Category	Call Options [£] Maturity Category 2	4	5	Average across maturities
Simulated Spread	1	2.83	6.31	10.65	5.60
	3	2.81	6.32	10.90	6.28
	6	2.33	5.95	10.55	6.68
	9	1.66	5.24	9.89	6.60
Average across moneyness		2.41	5.96	10.50	6.29
Ratio Observed/Simulated	1	4.22	2.25	1.23	2.57
	3	3.18	1.83	1.18	2.06
	6	3.13	1.65	1.00	1.93
	9	3.43	1.51	1.11	2.01
Average across moneyness		3.49	1.81	1.13	2.14

	Moneyness Category	Put Options [£] Maturity Category 2	4	5	Average across maturities
Simulated Spread	1	2.86	6.30	10.78	5.17
	3	2.78	6.24	10.75	6.10
	6	2.29	5.73	10.27	6.59
	9	1.48	4.86	9.18	6.65
Average across moneyness		2.35	5.78	10.25	6.13
Ratio Observed/Simulated	1	4.80	1.95	1.25	2.67
	3	3.14	1.80	1.07	2.00
	6	2.95	1.63	1.12	1.90
	9	3.32	1.73	1.06	2.04
Average across moneyness		3.55	1.78	1.12	2.15

The top four rows of each panel report the reservation spreads for a liquidity provider who trades daily in index futures to hedge the delta risk of an initial option purchase or sale. The reservation spread is computed by simulating 10,000 paths for the underlying index. We calculate the costs of daily delta-hedging a long and short option position and calculate the reservation spread by subtracting the average cost for the hedged short position from the average cost for the hedged long position. We calibrate our simulation using a dividend yield of 2.2%, an annualized risk-free interest rate of 2%, a volatility of 21.1%, a long-term average expected index return of 6% and an initial index level of 4,400. We assume that the transaction costs for trading in the futures market consist of the empirically observed futures bid-ask spreads (see Table 6) and of an order processing cost of £1 per traded futures contract. The bottom four rows of each panel report the ratios of the observed spreads, obtained from the appropriate categories in Tables 2 and 3, and the simulated reservation spreads. The top panel reports results for call options, the bottom panel for put options by moneyness and maturity categories.

Table 8: **Spreads for the Actual versus the Synthetic Index**

	Mean Daily Spreads [£]		
	Days to Maturity		
	0-90	91-180	181-270
Actual index	1.42	2.75	6.44
Synthetic index	38.40	50.75	57.00
Adjusting for differences in trading volume			
Synthetic index	20.36	39.25	52.52
Ratio synthetic to actual	14.34	14.27	8.16

The first two rows of the table report the mean bid-ask spreads for the actual index futures contract and the synthetic index futures contract constructed as described in Equations (1) and (2). The third row report the bid-ask spread for the synthetic index futures contract after adjusting for the higher trading volume in the index futures. The adjustment was made using the slope coefficient from a regression of the option spreads on the logarithm of option trading volume and plugging in the observed volume for the corresponding futures contracts.

Table 9: **Descriptive Statistics**

	Call Options N=8460		Put Options N=8725	
Explanatory Variable	Mean	Std.Dev	Mean	Std.Dev
Bid-Ask Spread [£]	8.57	3.92	8.56	3.78
Delta	0.39	0.16	0.35	0.13
Gamma	0.09	0.05	0.08	0.05
Vega	8.02	4.52	7.99	4.38
Theta	2.66	3.48	1.72	1.51
Implied Volatility [%]	21.50	6.71	24.01	7.49
Trading Volume [100's contracts]	3.33	7.65	3.58	9.04
Number of Trades [100's]	2.76	4.94	2.69	4.97
Open Interest [100's of contracts]	43.89	51.87	47.57	54.44
Option Price	139.42	104.95	144.70	93.23

The table reports the descriptive statistics for the absolute spreads and the explanatory variables used in the fixed effects regressions.

Table 10: **Correlation Matrix**

	Delta	Gamma	Vega	Theta	Vlty	Volm	Ntrd	Open	Price
Bid-Ask Spread	0.50	-0.39	0.55	-0.25	0.41	-0.11	-0.30	-0.12	0.69
	0.33	-0.35	0.51	-0.04	0.31	-0.10	-0.21	-0.05	0.67
Delta		0.07	0.51	-0.44	0.19	-0.10	-0.20	-0.09	0.73
		0.33	0.32	0.14	-0.16	-0.11	-0.08	-0.09	0.56
Gamma			-0.52	0.15	-0.13	0.13	0.47	-0.01	-0.44
			-0.52	0.57	-0.19	0.11	0.42	-0.06	-0.41
Vega				-0.53	-0.15	-0.14	-0.44	0.07	0.81
				-0.56	-0.30	-0.17	-0.42	0.05	0.82
Theta					0.33	0.12	0.31	0.01	-0.39
					0.51	0.17	0.46	-0.03	-0.28
Volatility						0.04	0.11	-0.18	0.21
						0.11	0.19	-0.02	0.03
Trading Volume							0.30	0.26	-0.10
							0.35	0.30	-0.13
Number of Trades								0.13	-0.36
								0.13	-0.30
Open Interest									-0.03
									0.02

The table reports the correlation coefficients for the absolute spreads and the explanatory variables used in the fixed effects regressions; the top number is the correlation coefficient for the call option panel and the bottom number is the correlation coefficient for the put option panel.

Table 11: **Fixed Effects Regression Results**

Explanatory Variable	Calls	Puts	Calls & Puts
Delta	4.609	4.297	4.575
	4.929	4.848	7.319
Gamma	4.287	0.607	4.553
	3.682	0.425	5.242
Vega	0.097	0.281	0.188
	2.653	8.089	7.547
Theta	-0.048	0.312	-0.010
	-4.006	8.063	-0.939
Implied Volatility	0.192	0.136	0.183
	23.976	15.686	35.281
Trading Volume	-0.004	-0.005	-0.004
	-1.078	-1.504	-1.642
Number of Trades	-0.053	-0.029	-0.037
	-7.399	-4.178	-7.525
Open Interest	-0.003	-0.004	-0.004
	-5.736	-6.661	-9.116
Option Price	0.021	0.023	0.021
	24.331	24.932	33.993
Constant	1.904	1.221	1.100
	5.360	2.936	4.160
F-test of	706.53	675.35	1362.0
H_0 : all parameters =0	0.000	0.000	0.000
R^2 within	0.431	0.412	0.418
R^2 between	0.854	0.909	0.878
R^2 overall	0.547	0.526	0.537
$\frac{\sigma_u^2}{\sigma_e^2+\sigma_u^2}$	0.27	0.23	0.23
Number of observations	8460	8725	17185
Number of groups	48	48	96
Avg. num. of obs. per group	176.2	181.8	179.0

The table reports the parameter estimates for fixed effects regressions of the daily averages of the pound bid-ask spreads for FTSE 100 index options; equations (4) and (6). The t-statistics for each parameter estimate is reported below the parameter estimate. An F-test of the null of all parameters jointly being equal to zero with p-values is reported in the second panel. The next three lines report the within, between, and overall R^2's.

Table 12: **Estimated Fixed Effects**

	Call Options [£]					
Moneyness	Maturity Category					
Category	1	2	3	4	5	6
1	-1.50	-1.94	-1.30	-1.52	-3.25	-4.10
2	-1.61	-1.38	-0.77	-	-	-
3	-0.96	-1.14	-0.70	-1.30	-1.94	-1.94
4	-0.70	-0.93	-0.64	-0.82	-2.09	-2.83
5	-0.20	-0.56	-0.25	-0.05	-1.14	-0.68
6	0.27	-0.13	0.22	-	-	-
7	0.90	0.46	0.66	0.72	0.48	0.94
8	1.15	1.13	1.24	-	-	-
9	1.39	1.69	1.53	1.89	1.49	1.95
10	1.48	1.95	1.86	-	-	-

F-test, $H_0 : u_i = 0$ i=1,..,48:
F(47,8403)=10.91, p-value=0.000
Correlation between fixed effects and predicted spread:
$corr(u_i, X'\beta) = -0.61$

	Put Options [£]					
Moneyness	Maturity Category					
Category	1	2	3	4	5	6
1	-0.69	-0.90	-0.15	-2.04	-3.37	-3.02
2	-0.49	-0.86	-1.05	-	-	-
3	-0.15	-0.90	-0.96	-1.83	-2.17	-1.96
4	0.05	-0.67	-0.77	-1.04	-2.38	-3.12
5	0.32	-0.29	-0.23	-0.17	-1.49	-0.92
6	0.77	0.03	0.27	-	-	-
7	1.18	0.65	0.65	0.69	-0.41	-0.17
8	1.36	1.12	1.08	-	-	-
9	1.58	1.41	1.46	1.20	0.54	0.38
10	1.57	1.57	1.73	-	-	-

F-test, $H_0 : u_i = 0$ i=1,..,48:
F(47,8403)=10.06, p-value=0.000
Correlation between fixed effects and predicted spread:
$corr(u_i, X'\beta) = -0.70$

The table reports the estimated fixed effects, u_i, by maturity and moneyness category (48 categories). Beneath each panel of fixed effects we report an F-test of the null hypothesis that all fixed effects are jointly equal to zero and the correlation between the estimated fixed effect and the predicted spread, $X'\beta$.

Table 13: **Robust Panel Regressions Results**

	Specification:		
Explanatory Variable	I	II	III
Delta	1.332	1.388	1.347
	2.476	2.834	2.734
Gamma	1.597	1.572	1.584
	1.316	1.520	1.522
Vega	0.200	0.191	0.190
	7.591	7.154	7.240
Theta	-0.027	-0.029	-0.032
	-1.599	-2.077	-2.265
Implied Volatility	0.193	0.188	0.189
	24.025	26.298	26.097
Trading Volume	-0.008	-0.008	-0.008
	-1.867	-2.098	-2.329
Number of Trades	-0.052	-0.045	-0.044
	-6.820	-7.132	-6.989
Open Interest	-0.002	-0.003	-0.003
	-3.365	-3.738	-3.700
Option Price	0.016	0.016	0.016
	13.114	13.552	13.821
Constant	1.690	1.805	1.824
	4.571	5.340	5.410
χ^2_9 test of	3567	3703	3708
H_0 : all parameters = 0	0.000	0.000	0.000
Log-likelihood	-8220.96	-8085.32	-8075.80
Number of observations		3972	
Number of groups		16	
Number of time periods		237	

The table reports the parameter estimates for cross-section and time-series GLS regression applied to the balanced panel of call and put options. Three different specifications are estimated. Specification I allows for a common AR1 coefficient for all groups, but assumes $\epsilon_{i,t}$ is homoscedastic. Specification II allows for heteroscedasticity and a common AR1 coefficient for $\epsilon_{i,t}$. Specification III allows for group specific autocorrelation coefficients and heteroscedasticity. The t-statistics for each parameter estimate is reported below the parameter estimate. A chi-squared test of the null of all parameters jointly being equal to zero with p-values is reported in the second panel together with value of the log-likelihood function.

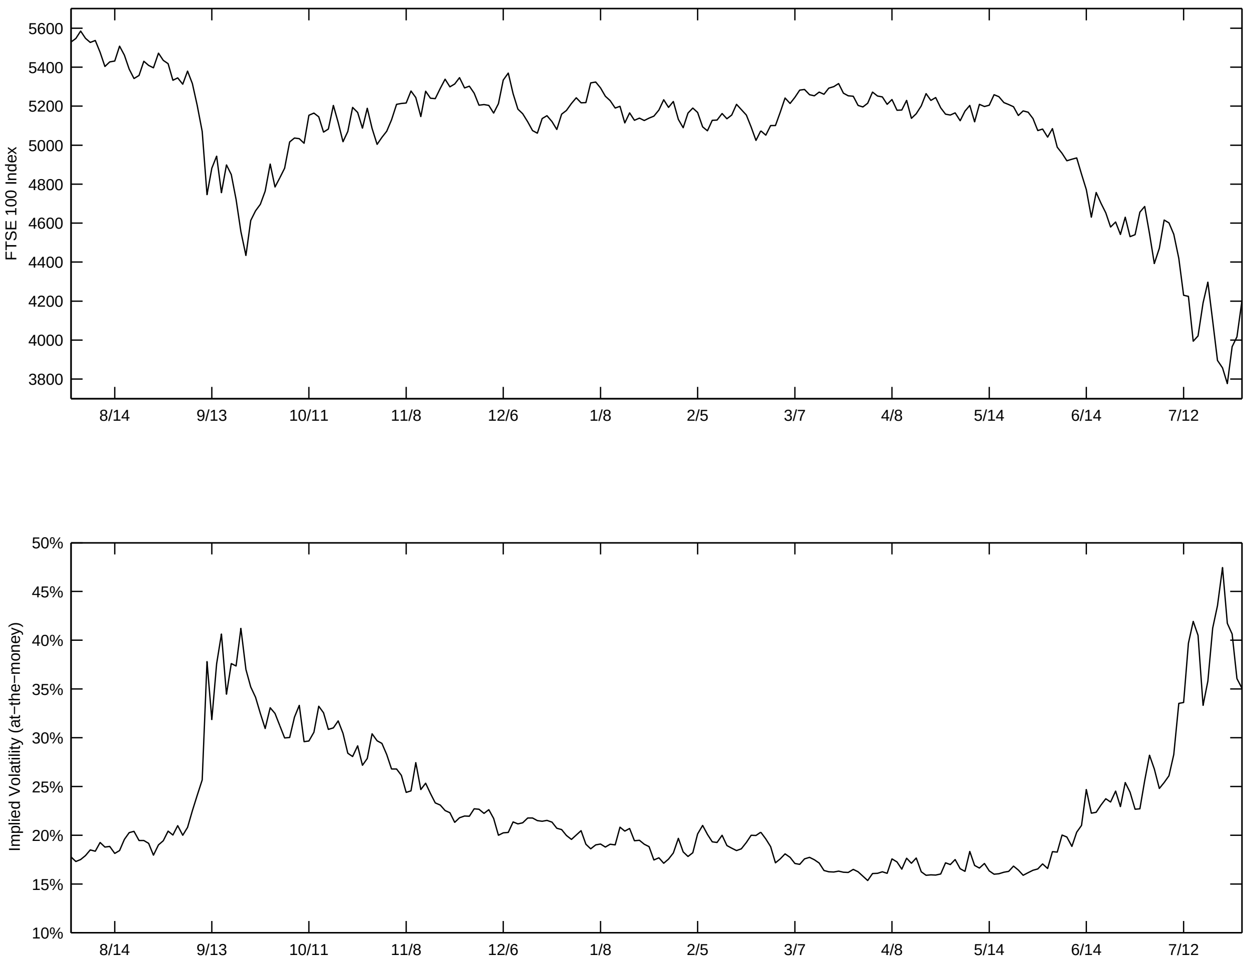

Figure 1: The top graph plots the time-series of the daily closing values of the FTSE 100 index. The bottom graph plots the average implied volatility calculated for call and put options with strike prices within 75 points of the index values and with 61-120 days to maturity. The sample period is August 1, 2001 to July 30, 2002.

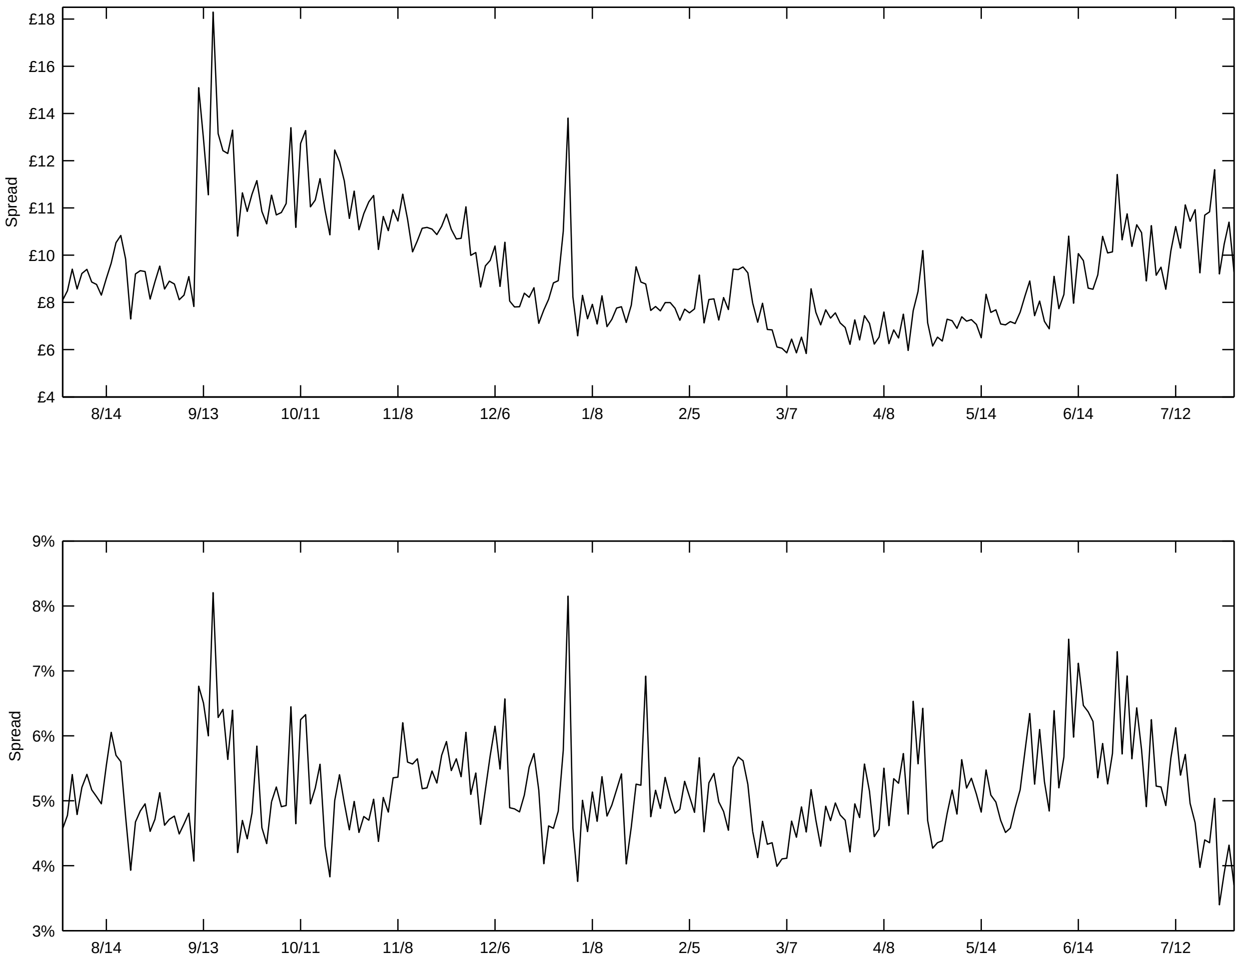

Figure 2: The top graph plots the daily average pound spread and the bottom graph plots the daily average relative spread for all contracts with strike prices within 75 index points of the current index level and with a time to maturity between 60 and 120 days. The sample period is August 1, 2001 to July 30, 2002.