ELSEVIER

Journal of Financial Markets 6 (2003) 573–605

Journal of FINANCIAL MARKETS

www.elsevier.com/locate/econbase

Local parametric analysis of derivatives pricing and hedging ☆

Peter Bossaerts[a,*], Pierre Hillion[b]

[a] California Institute of Technology, HSS-228-77, California Institute of Technology, Pasadena, CA 91125, USA

[b] INSEAD Asia Campus, 1 Ayer Rajah Avenue, Singapore 138676, Singapore

Abstract

A novel methodology for the analysis of derivatives pricing in incomplete markets is tested empirically. The methodology generates hedge ratios and derivatives prices. They are estimated from the correlation structure between the local co-movements of securities prices. First, the hedge ratios from a parsimonious complete-market model are estimated by fitting locally the changes in the derivatives and the underlying securities prices. Second, derivatives prices are obtained from the locally estimated hedge ratios. The methodology, referred to as local parametric estimation, is tested on a dataset of DAX index options and futures transactions from the computerized German Futures Exchange.

Keywords: Option; Local parametric estimation/hedging

1. Introduction

We present a novel methodology for the analysis of derivatives hedging and pricing in incomplete markets based on the optimal hedge estimation technique proposed by Bossaerts and Hillion (1997) and the hedge portfolio rebalancing analysis of Bossaerts and Werker (1997).

☆ This paper benefited from comments during seminars at UFSIA (Antwerp, Belgium), ECARE (Brussels, Belgium), the HFDF-II Conference in Zurich, and the Eighth Annual Derivatives Conference in Boston. We are grateful to the Deutsche Terminbörse for sending us the transaction data and to Wolfgang Bühler for the interbank deposit rates. We would like to thank an anonymous referee and the editor for constructive comments that led to major improvements and greatly improved the clarity of the paper.

*Corresponding author. Tel.: +1-626-395-4028; fax: +1-626-405-9841.

E-mail addresses: pbs@rioja.caltech.edu (P. Bossaerts), hillion@insead.edu (P. Hillion).

1386-4181/03/$ - see front matter
doi:10.1016/S1386-4181(03)00004-1

The approach suggested in this paper is based on Local Parametric Estimation (LPE), recently introduced in the statistics literature by Hjort and Glad (1995). Parametric methods for estimating non-linear models such as the generalized method of moments, non-linear least squares, or maximum likelihood are easy to apply and the results are easy to interpret. A drawback of parametric models is to impose too much structure on the data leading to inconsistent estimates and misleading inferences. In contrast, non-parametric techniques, such as kernels, nearest neighbors, among others, provide valid inferences under a broader class of structures. This comes at a cost. Non-parametric models suffer from two drawbacks, namely a slow rate of convergence and a smoothing bias.

LPE can be viewed as a semi-parametric approach that combines parametric and non-parametric methods. More specifically, LPE is a cross between the least squares estimator and the kernel density estimator. The traditional kernel density estimator of an unknown density is by construction completely non-parametric in the sense that it is impartial to special types of shapes of the underlying density. The idea of LPE is to cross an initial parametric density estimate with a kernel-type estimate, i.e., to encompass non-parametrically a parametric model and to estimate the local parameters using information from a neighborhood of the point of interest. LPE is a kernel smoother for non-parametric regression that uses prior information on regression shape in the form of a parametric model.

Semi-parametric estimators, such as LPE, offer three benefits. First, they work better than kernel estimators in a broad non-parametric neighborhood around the parametric family, while losing little or not at all when the true density is far from the parametric family. Second, they are useful in the higher dimension case when standard non-parametric methods, including the kernel method, are quite imprecise. Third, they capture local curvature much better than other alternatives such as polynomial functions. This last property is quite attractive for the problem under investigation, i.e., hedging and pricing options.

Bossaerts and Hillion (1997) show that LPE can be used to estimate the hedge ratios of a parsimonious complete-market model locally. By imposing very little structure, hedge ratios are estimated from lagged information contained in the changes of option prices with neighboring values of moneyness and maturity. The approach fully exploits the local curvature of parametric option pricing models. Option prices can in turn be derived from the LPE hedge ratios. Under certain assumptions about the pricing of the options and the composition of the hedge portfolio, Bossaerts and Werker (1997) show that the estimate of the size of the hedge ratio, i.e., the amount to be invested per option, is a consistent estimate of the correct value of the option under a quadratic loss function.

Bossaerts and Hillion (1997) investigate the performance of LPE hedging in an environment where a continuous-time option pricing model holds exactly, but where hedging is constrained to take place in discrete time. They find that LPE delta hedging significantly outperforms Black-Scholes delta hedging in a Black-Scholes simulated world. In the real world, most of the Black-Scholes assumptions, such as constant volatility or the absence of price jumps, are unlikely to hold. Significant pricing biases have been documented in the empirical option pricing literature and there is no guarantee that the superior performance of LPE in a simulated economy would still be observed in practice.

This paper has two goals. The first is to analyze the pricing and hedging performance of LPE using market rather than simulated data. LPE is tested using a very comprehensive dataset that contains all the DAX index options and futures transactions on the computerized German Futures Exchange, the Deutsche Terminbörse (DTB), over the period 1992–1994. The second goal is to analyze the pricing and hedging performance of LPE using a small set of underlying securities. This is quite challenging especially when the underlying securities display stochastic volatility.

In a Black-Scholes environment that accommodates stochastic volatility, Bossaerts et al. (1996) show that the number of securities in the hedge portfolio must increase without bounds to replicate perfectly the payoff on derivatives. They suggest the use of a transaction clock to reduce the randomness in volatility. This alleviates the constraint on the number of securities required for optimal hedging. The time transformation is adopted here as a means to annihilate stochastic volatility. The performance of the LPE approach for pricing and hedging the DAX stock index options is tested with two securities only, the DAX futures contract and the overnight money market rates.

The empirical tests are carried out on two sets of prices and hedge ratios. First, the LPE option prices and hedge ratios are obtained by fitting the Black (1976) model locally. Second, synchronous market prices and hedge ratios are estimated using the Black (1976) implied volatility. The performance of the LPE pricing technique is gauged by studying its ability to predict the DAX stock index options at expiration, and by comparing the LPE prediction to its market counterpart.

The deviations between the LPE and the synchronous market prices are also examined. A hedging analysis is performed by investigating the hedge error, defined as the difference between the value of the LPE hedge portfolio at maturity and the expiration value of the option. Finally, the performance of a self-financed delta hedging strategy based on the estimated parameters is compared to the standard strategy based on synchronous market prices.

Three results emerge from the empirical analysis. First, the LPE approach predicts payoffs at maturity as accurately as market prices. This is a remarkable finding given the fact that the LPE option prices are just obtained from the correlation between movements in the option and the underlying asset prices and do not include information in current option prices. This supports the claim of an efficient options market. Second, LPE option prices often deviate from market prices. Third, in contrast with both the first result and the findings reported in Bossaerts and Hillion (1997), self-financed delta hedging strategies based on the estimated parameters perform worse than hedging strategies based on market prices. These contradictory results appear to be caused by anomalous co-movements in the options and futures price changes, in particular hedge ratios above one and below zero, an anomaly observed by Bakshi et al. (2000) in the S&P500 index option market.

Stochastic volatility or price jumps could explain these results. The major benefit of hedging in transaction time rather than calendar time is to annihilate random volatility. In addition, omitted risk factors are taken care of by the estimation procedure. This suggests that stochastic volatility is unlikely to explain the results reported in the paper. A similar conclusion is reached by Bakshi et al. (2000) in their investigation of the violations of the predictions of one-dimensional diffusion class models. They find that the stochastic volatility model of Heston (1993) improves the fit of intraday option price changes marginally and conclude that future research should go beyond the stochastic volatility model. This leaves price jumps as a possible explanation for the empirical anomalies reported in the paper. As shown by Bergman et al. (1996), non-Markovian and discontinuous processes may generate option (call) prices that are a decreasing and a concave function of the underlying asset price over some range because of the violation of the "no-crossing" property.

Local Parametric Estimation differs significantly from other numerical procedures such as, for example, Monte Carlo simulation. Boyle (1977) proposed the use of Monte Carlo simulation for estimating derivative prices. Simulation is useful for estimating derivative prices when there is no analytical expression for the terminal distribution of the security price, when there are multiple state variables and when there are path dependent payoffs. The underlying asset is assumed to follow a particular stochastic process and a sampling procedure is used in a risk-neutral world to calculate the expected option payoff.

Until recently, simulation could not easily handle the early exercise feature of American options. A recent paper by Longstaff and Schwartz (2001) presents a new approach for approximating the value of American options by simulation. The cross-sectional information in the simulation is used to estimate the conditional expectation of the payoff from continuing to keep the option alive. This is done by regressing the ex-post realized payoffs from continuation on a set of basis functions of the values of the relevant state variables. Different approaches have also been suggested to value American options. They are based on stratification or parameterization techniques to approximate the transitional density function or the early exercise boundary.

Significant differences emerge from the comparison of the LPE approach to other estimation techniques that have been suggested in the literature. Theoretical derivatives pricing starts with a correlation analysis of the payoffs on several assets. Local parametric analysis relies entirely on the return correlation pattern between the option prices and the underlying security price. It exploits the correlation between the derivatives and the underlying asset prices to generate both the optimal hedge portfolio and the derivatives prices. The correlation-based LPE approach is a simple and robust model-free approach.

LPE is more natural than the standard time-series approach where hedge ratios and option values are estimated from a time-series analysis of the underlying assets values. The LPE estimates are not based on price levels. The relationship between the option and the underlying securities prices is not exploited, unlike the approach of Hutchinson et al. (1994). The LPE estimates do not exploit either information available in current option markets. The term structure of implied volatilities backed out from option prices is not inferred, unlike the approach of Dumas et al. (1998).

Unlike standard estimation procedures, correlation is the key parameter, not volatility.

The approach suggested in this paper is novel in several respects. First, LPE is not a numerical procedure to compute prices given known parameter values but an estimation procedure treating the underlying parameters as unknown. Second, compared to Monte Carlo simulation and other techniques, the local parametric approach for pricing and hedging derivatives imposes very little structure in the estimation. Unlike most of the literature, there is no need (i) to specify a continuous-time stochastic process, (ii) to simulate sample paths, (iii) to exploit the information contained in the time-series of prices or to estimate the parameters of continuous-time processes from discrete-time data, (iv) to exploit the information in current option prices, (v) to specify the market price of risk or finally (vi) to solve analytically intractable partial differential equations.

The paper is organized as follows. Section 2 presents the methodology. Section 3 discusses the dataset, the estimation and the testing strategies. Section 4 presents the empirical results and Section 5 concludes.

2. The methodology

The goal of this section is to show how to obtain an estimate of an option value and hedge ratio in an incomplete market for investors with quadratic utility. The underlying statistical model is initially assumed to be known. Local parametric estimation is introduced next, to provide a robust and parsimonious approximation of the theoretical model. Technical issues pertaining to the choice of the bandwidth and the estimation interval are then addressed.

2.1. The theory

Early theoretical continuous-time analysis appealed to inter-temporal asset pricing models to obtain equilibrium values for derivative assets. A popular example is the stochastic volatility option pricing model of Hull and White (1987). An alternative approach is followed in this paper, inspired by the more recent theoretical continuous-time analysis of incomplete markets, developed by Föllmer and Sondermann (1986), Duffie and Richardson (1991), Schweizer (1995) and Gouriéroux et al. (1995).

Consider, for example, the pricing of a call option written on a futures contract. Let z be a random vector that contains information available in the market place, such as the moneyness, the maturity of the derivative to be analyzed and possibly the interest rate level.

Consider the conditional projection of the net gains on the option ΔC onto the net gains from holding one bond ΔB and one futures contract ΔF

$$\Delta C = a(z) + x_{\mathrm{B}}(z)\Delta B + x_{\mathrm{F}}(z)\Delta F + \eta, \tag{1}$$

where

$$\mathrm{E}[\eta|z] = 0,$$
$$\mathrm{E}[\eta\Delta B|z] = 0,$$
$$\mathrm{E}[\eta\Delta F|z] = 0.$$

Eq. (1) represents a conditional projection because the projection error is orthogonal to the regressors conditional on z. For the projection to be possible, the projection coefficients $x_{\mathrm{B}}(z)$ and $x_{\mathrm{F}}(z)$, the number of bonds and futures contracts to be held in the hedge portfolio, respectively, must vary with z. In general, the functional relationship between the projection coefficients $x_{\mathrm{B}}(z), x_{\mathrm{F}}(z)$ and z is unknown and must be estimated. The hedge error, i.e., the difference between the option and the hedge portfolio payoffs, is equal to $a(z) + \eta$. By construction, the hedge error is uncorrelated with the payoff on the hedge portfolio.

Since the hedge ratios are obtained as the coefficients of a conditional projection, the resulting hedge portfolio is optimal under a quadratic hedge error loss function. The projection coefficients $x_{\mathrm{B}}(z)$ and $x_{\mathrm{F}}(z)$ provide the optimal number of bonds and futures contracts for an arbitrageur who attempts to replicate the option payoff under a quadratic loss function. Given that the second security in the hedge is a futures contract and does not require an investment outlay, the value of the hedge portfolio is given by the number of bonds times their price. This implies that, conditional on z, the value of the optimal hedge portfolio, C^*, is equal to,

$$C^* = x_{\mathrm{B}}(z)B. \tag{2}$$

Additional results pertaining to the pricing of the option can be obtained under appropriate assumptions. Assume that investors have quadratic utility and that in the absence of the option, their optimal portfolios consist entirely of the bonds and futures contracts used in the hedging strategy. If in Eq. (1), $a(z)$ is equal to zero, then,[1]

$$\frac{\Delta C}{C^*} = \left(x_{\mathrm{B}}(z)\,\frac{B}{C^*}\right)\frac{\Delta B}{B} + \left(x_{\mathrm{F}}(z)\,\frac{1}{C^*}\right)\Delta F + \varepsilon, \tag{3}$$

with

$$\mathrm{E}[\varepsilon|z] = 0,$$
$$\mathrm{E}\left[\varepsilon\,\frac{\Delta B}{B}|z\right] = 0,$$
$$\mathrm{E}[\varepsilon\Delta F|z] = 0.$$

Suppose that the price of the option is equal to C^*. Then, Eq. (3) implies that the market offers an average return on the option equal to that of a portfolio invested in bonds and futures but with a higher variance. At a price of C^*, the option would not improve the performance of the optimal portfolio composed of the bonds and futures contracts in the hedge portfolio for an investor with quadratic utility. There

[1] This implication is true only if C^* can be obtained from the elements in z. This is the case since $C^* = x_{\mathrm{B}}(z)B$ and B can be computed from the interest rate assumed to be in z.

would be zero demand for the option.[2] If the option is truly a derivative instrument in the sense that its net supply is equal to zero, the options market is in equilibrium and C^* is the equilibrium value of the option.

This suggests that if (i) mean-variance optimal portfolios consist exclusively of the bonds and the futures contracts used in the hedge portfolio and (ii) the intercept $a(z)$ in Eq. (1) is equal to 0, then $C^* = x_B(z)B$ provides the value of the hedge portfolio and gives the value of the option under quadratic loss. This strategy produces an estimate of the value of an option in an incomplete market for investors with quadratic utility.

Theory has little to say about what a non-zero $a(z)$ is. It is usually interpreted as a systematic "pricing error". Its presence indicates a discrepancy between the average net gain on the option and the optimal hedge portfolio. The latter need not be optimal under a non-quadratic loss function. A difference between the average net gain on an option and the quadratic-loss hedge portfolio may simply reflect different preferences.[3]

2.2. Local parametric estimation

The goal of this section is to discuss the estimation of the regression coefficients. The issue is to estimate the two coefficients $x_B(z)$ and $x_F(z)$, the optimal number of bonds and futures in the hedge portfolio, respectively. They depend on z but the functional relationship between these coefficients and z is unknown.

They are estimated using a technique that fits a parametric model locally. Local parametric estimation (LPE) has been suggested as an efficient alternative to non-parametric kernel estimation procedures that fit polynomials locally. Bossaerts and Hillion (1997) provide a detailed discussion of the estimation technique.

The Black (1976) model is used here as a parsimonious parametric complete-market model to capture the curvature and to estimate the hedge ratios locally. Under the Black model, the hedge ratios depend on z and σ the volatility of the futures price changes. They are denoted by

$$x_B^B(z;\sigma), x_F^B(z;\sigma), \tag{4}$$

where the superscript B stand for "Black". In this formulation, the estimates are based on local fits of the single volatility parameter.[4] The hedge ratios $x_B(z)$ and $x_F(z)$ are estimated locally at z by running a weighted non-linear least-squares

[2] See Ross (1978).

[3] The estimation strategy ensures that the error in the projection of option returns onto bond and index futures returns is conditionally uncorrelated with bond and index futures returns. If we were to extend the estimation strategy to ensure that this error become mean-independent of the bond and index futures returns, then our pricing result would obtain for any risk averse preference type. That is, quadratic utility is no longer needed. However, such an extension would require an unexplored and more sophisticated approach than proposed in the paper. Whence, we leave the issue to future research.

[4] In another formulation, the interest rate was treated as an additional parameter and was fitted locally. Convergence problems due to small sample sizes were met. Since hedge ratios are more sensitive to volatility than interest rates, the paper focuses on the first formulation.

regression,

$$\Delta C_i = x_{\mathrm{B}}^{\mathrm{B}}(z_i;\sigma)\Delta B_i + x_{\mathrm{F}}^{\mathrm{B}}(z_i;\sigma)\Delta F_i + \eta_i, \tag{5}$$

with $i = 1, \ldots, N$, using a sample of N observations with values of moneyness, maturity and interest rates equal to $z_1, \ldots, z_N$ and corresponding payoffs $(\Delta C_1, \Delta B_1, \Delta F_1), \ldots, (\Delta C_N, \Delta B_N, \Delta F_N)$.[5] The weights given to observation i in the non-linear weighted least-squares regression depends on the distance between z_i and z. The weights of observations that have a z_i far away from z are set equal to zero. As is discussed below, LPE involves the choice of a kernel function and a bandwidth. The weights are determined by means of a kernel. The bandwidth is the neighborhood around z for which the kernel function generates nonzero weights.

The regression specified in Eq. (5) provides estimates of $x_{\mathrm{B}}(z)$ and $x_{\mathrm{F}}(z)$ at z. They are denoted,

$$x_{\mathrm{B}}^{\mathrm{B}}(z;\hat{\sigma}_z), x_{\mathrm{F}}^{\mathrm{B}}(z;\hat{\sigma}_z),$$

where $\hat{\sigma}_z$ is the weighted non-linear least-squares estimate of σ. The subscript z captures the fact that volatility may change with z. It is referred to as the "local volatility" and remains constant throughout the sample for a given moneyness, maturity and interest rate level. Consequently, two options that have the same characteristics on two different dates have the same local volatility and hedge ratios.[6]

2.3. Hedging with the black model and a limited set of underlying securities

The functional form of the hedge ratios is obtained from the Black (1976) model. This does not mean that the Black (1976) model is assumed to hold. Rather, the modeling of the hedge ratios as a function of the moneyness, the maturity and the interest rate level is assumed to hold well as a local approximation.

Still, the choice of the Black model could be criticized. There exists substantial empirical evidence suggesting that asset prices exhibit stochastic volatility. The ability of LPE to hedge risk with only two securities, a futures and a bond, could be questioned. This is an important issue in light of the recent findings of Bossaerts et al. (1996). In a simple extension of the Black-Scholes environment that accommodates stochastic volatility, they illustrate how the payoff on derivatives can only be

[5] To alleviate estimation problems caused by non-stationarity of futures and options prices, the payoffs are normalized by the strike price.

The gains from the overnight positions do not vary much in the sample. Since they determine the LPE estimate only when overnight positions receive a positive weight, they act as a dummy variable. ΔB_i takes two values, either zero for intra-day observations or a positive number for overnight observations. This implies that the estimated option values depend on the co-movement between options and futures price changes and the difference between intra-day and overnight gains on options positions.

The impact of the gains on overnight inter-bank deposits is likely to be small for two reasons. First, their magnitude is smaller than the typical gain or loss on the underlying futures position. Second, there are relatively few overnight observations in the sample, reducing their potential impact even further.

[6] In the discussion above, the intercept $a(z)$ is assumed to be equal to zero. Without any precise information about its functional form, $a(z)$ should be estimated by locally fitting polynomials. In its extreme form, a zero-order polynomial coincides with the familiar non-parametric kernel estimation.

replicated perfectly if the number of securities in the hedge portfolio increases without bounds.

They show that a straightforward time transformation generates a perfect hedge with a few securities. Their result calls for rebalancing according to a clock under which the price of the underlying security shows little or no stochastic volatility. A simple switch from a calendar clock to a transaction clock annihilates the random volatility. The change of time scale is based on the early work of Tauchen and Pitts (1983) and Harris (1986,1987) who suggested mixture of normal distributions to model stock price changes. They show that the joint distribution of daily price changes and volume can be modeled as a mixture of bivariate normal distributions. One implication of the so-called Mixture of Distributions Hypothesis (MDH) is that price heteroskedasticity is caused by subordination of the return variance to a mixing variable usually interpreted as the rate of flow of information. The price variance changes through time when the probability distribution of the mixing variable changes through time.

Under the assumption that the number of transactions is proportional to the number of information events, the transaction count is a useful instrumental variable for estimating unobserved realizations of stochastic price variances. Under the additional assumption that transactions occur at a uniform rate in event time, Harris (1986) shows that (i) daily price changes divided by the square root of the daily number of transactions is less kurtotic than the distribution of the unadjusted price changes and (ii) price changes measured over fixed transaction intervals are normally distributed and more so the longer the transaction interval. The empirical results reported by Harris (1986) confirm that the asymptotic conditions of the Central Limit Theorem are approached as the measurement interval length increases.[7]

The time transformation is adopted here to reduce the randomness in volatility and to be able to hedge with a limited set of securities. The clock is defined as the number of transactions in the futures markets. The optimal number of transactions is empirically determined in such a way to make the behavior of the futures as close as possible to the Black assumptions.[8]

2.4. The kernel and the bandwidth choice

The kernel and the bandwidth are two parameters to be determined in the context of any local estimation procedure. In the weighted non-linear regression, the weight

[7] Harris (1986) reports for a sample of 50 NYSE stocks, the cross-sectional median of the excess kurtosis of daily price changes. The unadjusted and transaction-adjusted daily data estimates are equal to 2.92 and 0.65, respectively. The transaction data estimates are equal to 1.86, 1.33, 1.00 and 0.45 and 0.45 using 1, 5, 50 and 100 transactions, respectively.

[8] Option traders use both a calendar and a transaction clock to hedge their positions. They are usually delta neutral at the close of trading. This corresponds to a calendar clock. In addition, traders may have a delta neutral position at several points in time during the trading day. Intraday, traders do not hedge at fixed intervals of time. They typically hedge after large price movements. The results of Bossaerts et al. (1996) suggest that hedging errors are likely to be smaller intraday when performed in transaction time than in calendar time because of stochastic volatility. Additional evidence can be found in Bakshi et al. (2000). They show that rebalancing a Black-Scholes hedge in calendar time as often as possible is not necessarily optimal.

given to observation i with z_i closest to the target z is determined using a kernel function. The Epanechnikov kernel is used here.[9] Formally, the weight $w(z, z_i)$ given to observation i is equal to,

$$w(z, z_i) = 0.75\left[1 - \min\left(1, \frac{\|z_i - z\|}{b}\right)\right]^2,$$

where $\|z_i - z\|$ denotes the Euclidian distance between z_i and z and b denotes the bandwidth. Normalizing each component of z_i by its range, the Euclidian distance is equal to,

$$\|z_i - z\| = \sqrt{\sum_{j=1}^{3}\left(\frac{z_{i,j}}{\max_k(z_{k,j}) - \min_k(z_{k,j})}\right)^2},$$

where the summation is calculated over the three variables contained in z. The interpretation of the bandwidth parameter b is straightforward with an Epanechnikov kernel. It provides a direct measure of the size of the sample that receives a positive weight in the LPE estimation. If, for example, the bandwidth is set equal to 0.15, all observations for which,

$$\|z_i - z\| \geqslant 0.15,$$

receive zero weight in the LPE estimation. If the components of z_i are independently and uniformly distributed, exactly 15% of the observations in the sample receive a positive weight with a 0.15 bandwidth. In practice these conditions are not met resulting in a lower sample size with on average 5% of the sample being used.

The choice of a bandwidth is important and delicate. The bandwidth controls the trade-off between the bias and the efficiency of the estimated parameters. A large bandwidth increases the sample size, which in turn raises the accuracy of the estimates. Biased estimates are obtained when the parametric model being fit locally is invalid. Conversely, a small bandwidth reduces the bias in the parameter estimates but increases its sampling variance.

Despite its popularity for the determination of the optimal bandwidth in non-parametric procedures, cross-validation would be prohibitively costly in terms of computation time in this particular study. The bandwidth is chosen on the basis of the hedging performance of LPE on a subset of the data.[10] The search for the optimal bandwidth is rudimentary. Several values within a range of 0.05–0.30 are tried. Within that range, little sensitivity in the hedging performance is noticed but the 0.15 bandwidth leads to the lowest average hedge error.[11] Even though the bandwidth choice of 0.15 may not be optimal for other contracts, the experiment is not repeated on other subsets to refrain from data mining.

Finally, when the bandwidth is fixed in local parametric estimation, there is no guarantee that enough observations receive positive weight to mitigate the sampling

[9] The results are insensitive to the choice of other kernel functions, including the Normal kernel.

[10] The optimal bandwidth was calculated using the contracts that expired on 3/92.

[11] The 0.15 value is the optimal bandwidth in the analysis of Bossaerts and Hillion (1997). It is determined using cross-validation.

variation. The effective sample size may potentially drop to zero. To minimize the sampling variance, the bandwidth is successively increased by 10% until at least 50 observations receive positive weights.

3. The estimation and testing strategy

3.1. The data

The LPE approach is tested using a very comprehensive dataset that includes futures and options written on the German Stock Exchange Index (DAX) over the period 1992–1994. The dataset includes all the transactions, volume and price, in European options on the DAX futures.

The risk of the option is hedged with the futures contract rather than the stock index to reduce the actual transaction costs. From the entire record of futures and options transactions available in that period, only those options that mature simultaneously with a futures contract are selected.[12] For example, for the two time periods between 1/92-3/92 and 3/92-6/92, only the option contracts with expiration 3/92 and 6/92 are included, respectively. These options are matched with the 3/92 and 6/92 futures contracts, respectively. The daily overnight money market rates Tagesgeld, are used to compute the return on overnight positions.[13]

Trading in futures and options is totally computerized and takes place between 9:30 a.m. and 4:00 p.m.[14] Both the DAX index futures and options markets are liquid. Open interest in the futures contract fluctuates between 60,000 and 160,000 contracts, where each contract is worth DEM 100 per index point.[15] Open interest in the option contract moves between 400,000 and 1,500,000 contracts. The contract size is DEM 10 per DAX point. Values for both contracts are quoted in index points (with one decimal place). For the options, values are quoted in DEM per 1/10 contract. Exercise prices for the options are set in increments of 25 index points. Five exercise prices are introduced for each expiration month.

An exploratory analysis reveals extremely high kurtosis and persistence in volatility. To reduce the randomness in volatility, the time change discussed in the previous section is implemented as follows. The futures contract with the shortest time left to maturity is used as the benchmark. Its transaction frequency is calculated

[12] There is a slight discrepancy between the settlement value of futures and options contracts that expire simultaneously. For the futures contract, the settlement value is computed on the basis of opening prices for the stocks constituting the DAX index. For the option contract, final settlement is based on the average value of the DAX between 1:21 p.m. and 1:30 p.m. on the same day.

[13] Only those are relevant since option positions acquired and liquidated during the same trading day do not involve an investment outlay. This implies that the corresponding hedge portfolio does not involve a risk-free position either.

[14] There is a pre-trading period, from 7:30 a.m until 9:30 a.m. and post trading from 4:00 p.m. until 5:15 p.m. The data only cover the official trading period between 9:30 a.m. and 4:00 p.m. The time stamp is accurate up to 1/100 of a second. Even with such a time stamp, there are ties in the futures files.

[15] The DAX index level is about 1625 at the beginning of the period. The size of a futures contract was DEM 165,000 and the open interest between DEM 10 and DEM 26 billion.

Table 1
Descriptive statistics

Contract	Futures price changes				Median time gap		
	N	mean	s.d.	k	G1	G2	G3
3/92	218	0.026	0.328	5.0	11	7682	588
6/92	162	0.014	0.314	6.8	14	8929	579
9/92	354	−0.034	0.492	10.1	7	3492	394
12/92	290	−0.006	0.471	4.8	10	4793	533
3/93	296	0.026	0.367	4.6	9	4418	283
6/93	218	−0.000	0.307	3.8	12	6855	422
9/93	400	0.018	0.341	4.8	7	3360	390
12/93	420	0.026	0.328	5.4	6	2962	261
3/94	661	−0.006	0.395	12.4	4	1911	292
6/94	481	−0.010	0.336	4.9	5	2409	272
9/94	492	0.008	0.305	4.6	5	2672	269
12/94	489	0.007	0.350	5.9	6	2857	364

Remarks. Contracts are indicated with their expiration month (first column). N = number of futures price changes over intervals of 300 transactions. Mean and standard deviation (s.d.) are in percent, k = kurtosis, G1 = median time gap between options transaction and subsequent futures transaction, in seconds; G2 = median time gap between the first subsequent futures transaction and the 300th subsequent futures transaction, in seconds; G3 = median time gap between the 300th subsequent futures transaction and the first subsequent options transaction, in seconds.

and used as a basis for the time transformation. Several transaction frequencies are tested. The estimation of futures price changes over intervals of 300 transactions leads to the sharpest decrease in both the excess kurtosis and the volatility persistence. Periods that cover a night in calendar time are not discernible at this frequency any longer.[16] On average, a period of 300 transactions correspond to roughly a $1\frac{1}{2}$ h period in trading time, after considering the close-to-open interval to be of zero length.

Table 1 reports descriptive statistics, including estimates of kurtosis. Table 2 provides autocorrelation coefficients up to the fifth lag. The absence of negative first-order autocorrelation suggests that there is no discernible bid-ask bounce.

3.2. The training sample

The "training" sample refers to the part of the dataset used to compute the LPE estimates of local volatilities. The sample is constructed using a four-step procedure, described below for an option with a particular expiration date and strike price. This generates one observation in the training sample say, observation i.

- *Step* 1: From the options dataset, record the option transaction time and the option transaction price. Let C_i denote the call transaction price.

[16] At higher frequencies, the presence of time intervals that straddle a close-to-open period re-introduces excess kurtosis.

Table 2
Autocorrelations of futures price changes

Contract	N	Autocorrelation order				
		1	2	3	4	5
3/92	218	0.037	0.038	−0.001	−0.007	−0.086
6/92	162	0.093	−0.010	0.005	0.033	−0.049
9/92	354	−0.008	−0.023	0.029	0.067	0.048
12/92	290	0.079	−0.079	0.000	0.077	0.025
3/93	296	0.021	−0.077	0.020	0.073	−0.097
6/93	218	0.030	0.060	−0.017	−0.074	0.002
9/93	400	0.073	−0.041	0.034	−0.018	0.026
12/93	420	0.077	0.010	−0.008	0.046	−0.014
3/94	661	0.012	0.033	0.037	0.025	−0.041
6/94	481	0.040	−0.045	0.041	0.036	−0.041
9/94	492	−0.004	−0.007	0.015	0.021	−0.033
12/94	489	0.041	0.091*	0.004	0.002	−0.039

Remarks. Contracts are indicated with their expiration month (first column). N = number of futures price changes over intervals of 300 transactions. * indicates significance at the 5% level.

- *Step* 2: From the futures dataset, search for and record the next available futures transaction price. Let F_i denote this price.
- *Step* 3: From the futures dataset, advance the transaction clock 300 futures transactions forward and record the new futures price. Let F_i^* denote this price.
- *Step* 4: From the options dataset, search for and record the next available option transactions price. Let C_i^* denote this price.

This sampling technique generates a time-ordered sequence of four prices, (C_i, F_i, F_i^*, C_i^*), from which the following two price changes can be calculated,

$$\Delta C_i = C_i^* - C_i,$$
$$\Delta F_i = F_i^* - F_i,$$

the changes in the call and the futures prices, respectively. The next observation $i+1$ is obtained by turning back to the options dataset and by repeating the four-step procedure. Fig. 1 provides a schematic view of the sampling technique.

Table 1 documents the median calendar time interval in seconds between the first option transaction, C_i, and the next futures transaction, F_i, between the two futures transactions, F_i and F_i^* and between the second futures transaction, F_i^*, and the subsequent option transaction C_i^*.[17]

[17] Consider Fig. 1. An alternative scheme is to select the option transaction price A as the observation $i+1$ with A^* as the matching option. This alternative sampling strategy generates substantial overlap between observation i and $i+1$. Initially, the original dataset was constructed in this fashion. The analysis was re-run on non-overlapping data because concerns that the substantial time overlap could affect the results. Beyond a small drop in precision, the results hardly changed.8

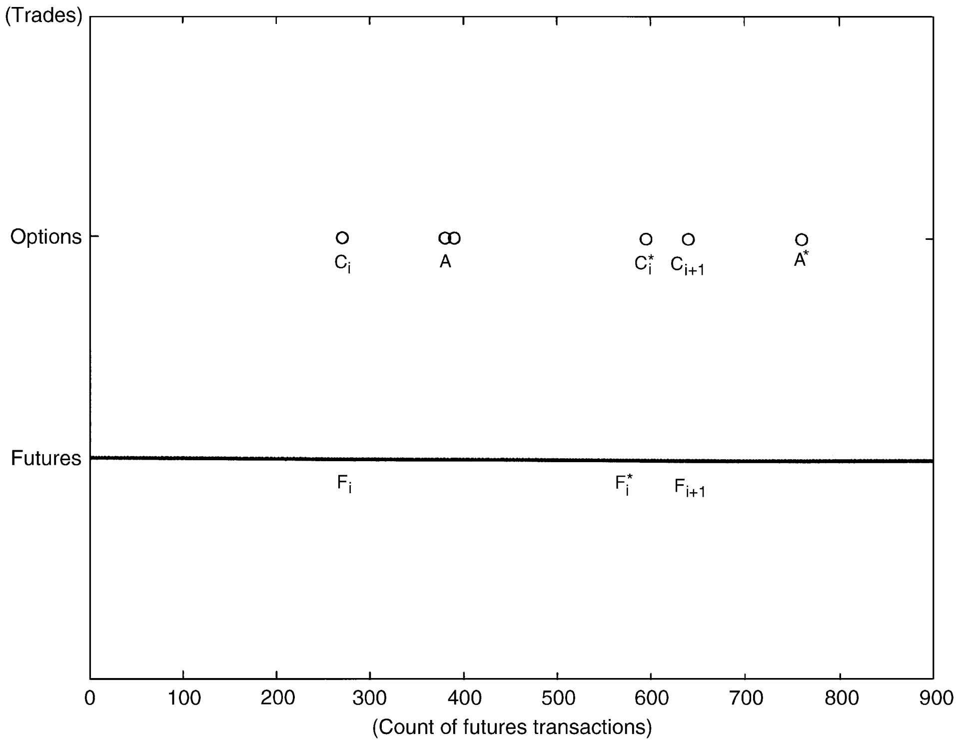

Fig. 1. Schematic representation of the construction of the training sample. The ith observation consists of: an options transactions price, C_i; the subsequent futures quote, F_i; the futures quote 300 futures transactions later, F_i^*; and the subsequent options transactions price, C_i^*. The $(i+1)$st observation starts with options transactions price C_{i+1} and futures quote F_{i+1}. In an alternative scheme (whose results are not reported in this paper), options transactions prices A and A^* were used to construct observation $i+1$. The substantial time overlap with observation i did not affect the estimation results qualitatively. In the text: $\Delta C_i = C_i^* - C_i$; $\Delta F_i = F_i^* - F_i$.

A major benefit of the sampling technique is to avoid estimation problems due to the lack of synchronicity between the options and the futures markets. The futures price F_i is close in time to the option price C_i because of the depth of the futures market. In contrast, C_i^* may sometimes be far away from F_i^*. Non-synchronicity generates a random pricing error, i.e., the difference between C_i^* and the unobserved option price synchronous with F_i^*.[18]

[18] The regression parameter estimates are not affected by the pricing error for the following reason. In the regression equation (5), the dependent and independent variables are the change in the option price and the change in the futures price, respectively. Unlike the dependent variable, the independent variable is measured without sampling error. As is well known from the standard econometric analysis of the errors-in-variables problem, the error in the dependent variable is absorbed in the error term, provided the independent variable is error-free. This last condition being satisfied, estimation problems due to non-synchronicity are avoided.

The training sample extends all the way to the point in time where an option is hedged or priced.[19] The sample runs back at least two calendar months except for the 3/92 contracts. To avoid unmanageably large training samples, observations that occur before the third month preceding the expiration month are dropped. The training sample sizes include a minimum of 1,000 observations and a maximum of around 5,000 observations.[20]

Finally, the payoffs on the interbank deposits are computed as follows. For observations that do not straddle a close-to-open interval, the initial investment is set equal to zero. This implies that $\Delta B = 0$. For the remaining observations, an investment of \$1 pays \$ $(1 + R)$ the following day, where R denotes the overnight interest rate. This implies that $\Delta B = 1 + R$.

3.3. *The "testing" sample*

The pricing and hedging analysis involves a "testing" sample constructed as follows.

Consider, for example, the call option with strike 1700 and expiration 3/92. LPE is used to price and hedge the option in transaction time defined as 300 transactions of the futures contract that expires on 3/92. The "training" sample consists of observations pre-dating the "testing" sample. In the example above, only the 1/92 observations in the "training" sample are used to estimate the LPE local volatility for the first observation in the testing sample. The first observation is then added to the "training" sample to estimate the LPE local volatility for the second observation in the "testing" sample. Moving over time towards maturity, observations are added to the "training" sample.[21]

The LPE pricing and hedging procedure starts on the first trading day of the month preceding the expiration month and continues until expiration. This process is repeated for every single option in the data set.

3.4. *Pricing and hedging analysis*

Several tests, carried out on all the options in the data set, are performed to assess the performance of LPE as a pricing and hedging technique. They involve two different estimates of volatility and consequently two different estimates of the option price. The first is the local volatility estimate obtained from LPE. The second is implied from the Black (1976) model. The corresponding option prices are referred to as the LPE price and the synchronous market price respectively. These estimates are obtained as follows.

[19] The data set is constructed in such a way that the time of the C_i^* transaction, the last transaction in the sequence, occurs before this reference point.

[20] In contrast, with overlapping observations, the samples would include between 17,000 and 45,000 observations.

[21] As discussed above, the end-of-period option transaction in the training sample always occurs prior to the beginning of the next observation in the testing sample. Also, there is a limit on how large the sample can be. Past that limit, old observations are dropped from the "training" sample.

For each observation in the testing sample, LPE is ran on the corresponding training sample. This generates an estimate of local volatility, denoted by $\hat{\sigma}_z$ where the vector z consists of the moneyness, M, the maturity, T of the option as well as the interest rate level R. From the local volatility estimate, the LPE option value is estimated as,

$$x_{\mathrm{B}}^{\mathrm{B}}(z, \hat{\sigma}_z)B,$$

i.e., the value of the LPE hedge portfolio.

The training sample is then scanned to find the most recent end-of-period option transaction price. Only those options with a moneyness within 0.5% of M and the same maturity T are considered in the process. The volatility is implied from the option price using the Black (1976) option pricing model. This estimate is denoted by σ_{bl}. The Black (1976) model is used a second time to generate the synchronous market price using M, T and σ_{bl} as inputs.

The performance of LPE is then assessed as follows.

- The performance of LPE *as a pricing technique* is assessed by comparing the ability of the LPE price and the synchronous market price to predict the option value at maturity. The deviations between LPE prices and synchronous market prices are also investigated. LPE may generate prices that predict expiration values better than the market. If large deviations are frequently observed, LPE superior pricing ability can only be exploited by holding the positions until maturity.
- The performance of LPE *as a hedging technique* is assessed by comparing hedging errors obtained with a LPE hedging strategy and a Black hedging strategy, based on local volatility and implied volatility, respectively. In both cases, the error is defined as the value of the hedge portfolio at maturity and the expiration value of the option. The hedging strategy is a self-financed delta hedge implemented in transaction time.[22]

The hedging strategy is implemented as follows. It is assumed to start within the trading day with no position in interbank deposits.

- Let $x_{\mathrm{F}}^{\mathrm{B}}(z_1, \sigma)$ be the number of futures contracts acquired, where the vector z_1 contains the moneyness, the maturity of the option and the interest rate level, with σ being either the local or the implied volatility.
- At the beginning of the next observation, 300 futures observations later, the futures position is liquidated and the gains/losses are deposited in a bank account. A new futures position, denoted by $x_{\mathrm{F}}^{\mathrm{B}}(z_2, \sigma)$ is acquired. Assume that this position is not liquidated until the next day. An amount of $x_{\mathrm{F}}^{\mathrm{B}}(z_2, \sigma)$ is deposited in the bank account. As an overnight position, it carries interest. Note that interest is also earned or charged on the gains/losses from the liquidation of the futures position in the first interval of 300 futures transactions.
- The third observation starts the next morning. Assume it is liquidated the same day. This implies that a new futures position of $x_{\mathrm{F}}^{\mathrm{B}}(z_3, \sigma)$ contracts is acquired. Gains or losses from the previous futures position are deposited in the bank

[22] Its properties are studied in Bossaerts and Werker (1997). They show that, under certain conditions, the cumulative hedge error is normally distributed despite the non-normality of the option's payoff.

account, adding to the principal plus interest gained overnight. The bank account does not accumulate any interest for the third observation since it ends within the same trading day.

- This hedging exercise is pursued until the expiration of the option. At maturity, the hedging error is obtained by comparing the value of the hedging portfolio to the option value.

4. Empirical results

4.1. An example

The pricing and hedging performance of the LPE approach is illustrated in Fig. 2. It displays the evolution of the LPE price, the synchronous market price and the mark value of the (self-financed) LPE delta-hedged portfolio for the 6/93 call option with strike 1675. They are represented by the solid line, the dotted line and the dashed line, respectively.

Three results emerge from Fig. 2.

- First, the LPE price tracks the market price very closely.

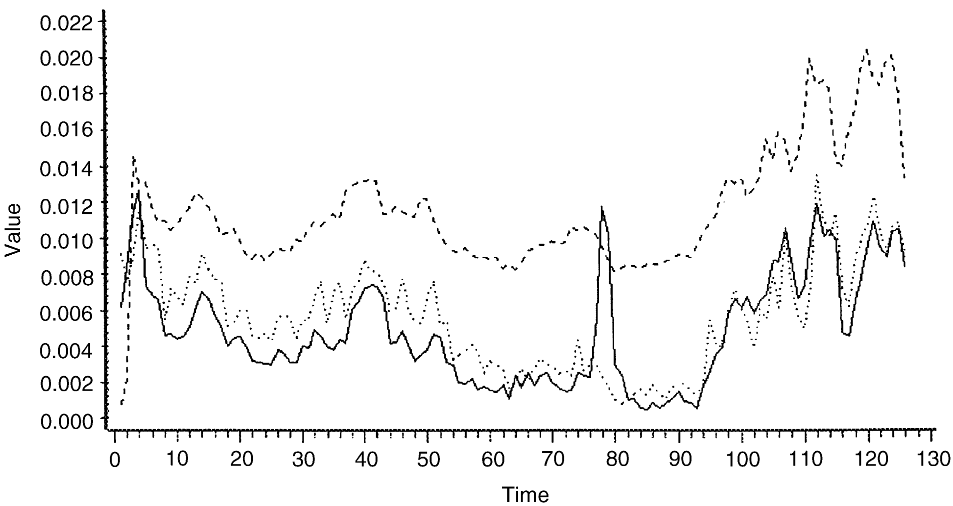

Fig. 2. Time series plots of values of the DAX call option with strike 1675 and maturity 6/93. The solid line depicts the price of the option, estimated using LPE. The dotted line depicts synchronous market prices. The latter are obtained by implying the volatility from the most recent transaction price of an option with similar richness and maturity, and using this implied volatility and the Black option pricing formula to derive a new quote for the maturity and richness of the estimated option price. The dashed line depicts values of a self-financed, modified delta hedge based on local volatilities from LPE. All prices and values are expressed as a percentage of the strike price (premia). Time is measured in number of intervals of 300 futures transactions (each interval corresponds to $1\frac{1}{2}$h calendar trading time, on average; this corresponds to slightly over $\frac{1}{4}$ trading day).

- Second, the LPE price deviates from the market price on rare occasions. For example, a major aberration is observed around the 80th 300-transaction interval. The LPE estimate surges dramatically and returns back to the level of the market price just after a few periods.
- Third, the LPE hedge portfolio does not perform well. While changes in the value of the hedge portfolio follow the pattern of the option market price, its level is much too high. The hedge error defined as the difference between the value of the hedge portfolio and the value of the option at maturity is sizeable.[23]

Fig. 2 suggests that the success of LPE at pricing options does not translate into accurate delta hedging. The situation in Fig. 2 may not be representative. A more systematic and rigorous analysis performed on the all the contracts in the dataset is needed to validate these preliminary results.

4.2. *Performance of the lpe pricing procedure*

The performance of LPE at pricing options is assessed by measuring the ability of the LPE estimates to predict the options final payoffs at maturity and by comparing the LPE prediction to the market prediction.

Table 3 reports summary statistics on two sets of pricing errors for option contracts with different maturities. They are computed as the difference between the option price, estimated either from LPE or from a synchronous market price, and the option's final payoff. Confirming the preliminary evidence of Fig. 2, the LPE pricing errors are found to be as small as their market counterparts. This suggests that LPE prices perform as well as current market prices as predictors of the final payoff of the option. This is a remarkable result for an estimation procedure that imposes so little structure and merely exploits the information in co-movements in asset prices, as opposed to all other prediction procedures, which always use at least information in price *levels*. Our results confirm option pricing theory, which predicts that the local correlation between the option and the underlying asset prices ought to be able to generate price levels that are in line with market prices.

The LPE prices are directly compared to the synchronous market prices in Table 4. As was inferred from Fig. 2, major deviations between the two are observed. Though the overall pricing performance is not affected, in the sense that both predict the option final payoff equally well, LPE often disagrees with the market on the correct price level. Table 4 provides only unconditional price deviations. Additional investigations must be pursued to understand their origin. Useful information is obtained by looking at the relationship between the pricing errors and the strike

[23] Notice that in the first few periods, the value is close to zero. Initially, the hedge portfolio requires only futures trades because the position is not carried overnight. The nonzero value of the portfolio follows from gains/losses in the futures position. The first overnight position has a dramatic impact on the value of the hedge portfolio. It is only at that point that a net investment in the hedge is taken. Before and after that, no further cash inflows or outflows are necessary since the portfolio is self-financing. The large increase in the value of the hedge portfolio in Fig. 2 coincides with the point at which the initial investment in the delta hedge is taken.

Table 3
LPE prediction errors versus market prediction errors, for different maturities

	Maturity					
	60		130		210	
	LPE	Market	LPE	Market	LPE	Market
N	251	251	201	201	162	162
Average	0.0053**	0.0058**	−0.0014	−0.0040	0.0147**	0.0092
s.d.	0.0201	0.0194	0.0441	0.0376	0.0456	0.0367
Skewness	0.38	−0.03	0.41	−0.84	1.38	1.10
Kurtosis	5.8	4.2	5.3	3.1	6.6	5.5
Minimum	−0.0518	−0.0518	−0.0999	−0.0999	−0.0607	−0.0607
5%	−0.0272	−0.0231	−0.0677	−0.0849	−0.0531	−0.0508
95%	0.0409	0.0415	0.0472	0.0445	0.1092	0.1084
Maximum	0.1019	0.0570	0.1743	0.0579	0.2424	0.1241

Remarks. The prediction error is defined to be the difference between the option price (estimated from LPE or a synchronous market price) and the final payoff on the option. The synchronous market price is obtained by implying the volatility from the most recent transaction price of an option with similar richness and maturity, and using this implied volatility and the Black option pricing formula to derive a new quote for the maturity and richness of the estimated option price. All prices are expressed as a percentage of the strike price (premia). N denotes sample size. Maturity is measured in number of intervals of 300 futures transactions (each interval corresponds to $1\frac{1}{2}$ h calendar trading time, on average; this corresponds to slightly over $\frac{1}{4}$ trading day). ** indicates: significant at the 1% level (two-tailed t-test).

Table 4
Deviations between LPE-based prices and market prices

	Maturity						
	60	90	120	150	180	210	240
N	251	251	234	201	181	162	144
Average	−0.0005	0.0049**	0.0014	0.0019	−0.0003	0.0055**	0.0094**
s.d.	0.0085	0.0320	0.0224	0.0209	0.0119	0.0276	0.0335
Skewness	13.09	5.56	7.36	7.03	1.18	4.88	3.43
Kurtosis	198.9	34.6	61.6	57.6	12.4	32.2	17.0
Minimum	−0.0173	−0.0165	−0.0202	−0.0257	−0.0505	−0.0221	−0.0381
5%	−0.0066	−0.0071	−0.0091	−0.0093	−0.0122	−0.0082	−0.0137
95%	0.0013	0.0091	0.0071	0.0090	0.0254	0.0521	0.0810
Maximum	0.1251	0.2298	0.2195	0.1882	0.0582	0.2249	0.2195

Remarks. Statistics are given for the deviations between the option price from LPE and a synchronous market price. The latter is obtained by implying the volatility from the most recent transaction price of an option with similar richness and maturity, and using this implied volatility and the Black option pricing formula to derive a new quote for the maturity and richness of the estimated option price. N denotes sample size. All prices are expressed as a percentage of the strike price (premia). Maturity is measured in number of intervals of 300 futures transactions (each interval corresponds to $1\frac{1}{2}$ h calendar trading time, on average; this corresponds to slightly over $\frac{1}{4}$ trading day). * indicates: significant at the 5% level (two-tailed t-test).

price of the options. Fig. 3 plots the deviations of LPE prices from synchronous market prices as a function of the "richness", defined as the underlying futures quote divided by the strike price.

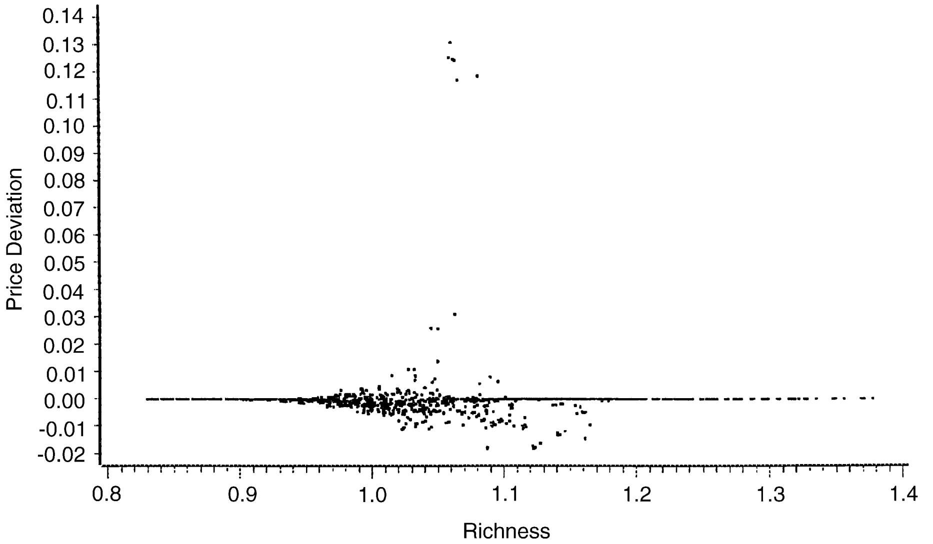

Fig. 3. Plot of deviations of LPE-based prices from synchronous market prices, against richness (underlying futures quote divided by strike price). The synchronous market price is obtained by implying the volatility from the most recent transaction price of an option with similar richness and maturity, and using this implied volatility and the Black option pricing formula to derive a new quote for the maturity and richness of the estimated option price. All prices are expressed as a percentage of the strike price (premia). Only observations with a maturity between 60 and 64 are included. Maturity is measured in number of intervals of 300 futures transactions (each interval corresponds to $1\frac{1}{2}$ h calendar trading time, on average; this corresponds to slightly over $\frac{1}{4}$ trading day).

Fig. 3 show that the largest deviations occur for near at-the-money options. This is not surprising. Synchronous market prices are obtained from the Black (1976) implied volatilities. For deep out-of-the money or deep in-the-money options, the Black (1976) price is insensitive to the volatility and hence the volatility is implied to be the initial guess set equal to 0.15. A similar phenomenon occurs for local volatilities. Ignoring the overnight positions, LPE finds the estimate of the local volatility at z by minimizing,

$$\sum_{i}^{N} w(z, z_i)(\Delta C_i - x_{\mathrm{F}}^{\mathrm{B}}(z_i; \sigma)\Delta F_i\ )^2,$$

where $w(z, z_i)$ is the weight given to observation i, a function of the distance between z_i and z. The first order condition is,

$$2\sum_{i=1}^{N} w(z, z_i)(\Delta C_i - x_{\mathrm{F}}^{\mathrm{B}}(z_i; \sigma)\Delta F_i)\,\frac{\partial x_{\mathrm{F}}^{\mathrm{B}}(z_i; \sigma)}{\partial \sigma}\,\Delta F_i = 0. \tag{6}$$

In the Black model,

$$\frac{\partial x_{\mathrm{F}}^{\mathrm{B}}(z_i; \sigma)}{\partial \sigma} = \sigma\sqrt{\tau_i}n(d_i)\mathrm{e}^{-r_i\tau_i},$$

where τ_i, K_i and F_i denote for observation i, the option maturity, the option strike price and the beginning-of-period futures price, respectively and r_i is the interest rate. The parameter $d_i$is equal to,

$$d_i = \frac{\ln(F_i/K_i)}{\sigma\sqrt{\tau_i}} + \frac{1}{2}\sigma\sqrt{\tau_i},$$

with $n(\cdot)$ being the normal density function. The latter is very close to zero for both deep-in-the-money, i.e., high F_i/K_i and deep-out-of-the-money i.e., low F_i/K_i options. The first order condition is automatically satisfied in those cases and the optimization algorithm stops at the initial value set equal to 0.15. For these options, no discrepancies are observed, i.e., the pricing errors decrease to zero, since synchronous market prices and LPE prices are both computed from the same initial value. An analogous phenomenon occurs when the option maturity decreases to zero. For near-the-money options, the pricing errors are substantial in certain cases.

These results suggest that the aberration noticed in Fig. 2 is not a rare event.

Fig. 4 plots local volatility against richness. It is the mirror image of Fig. 3. Local volatilities for moneyness away from 1.0 converge to 0.15. In contrast, for near-the-money options, local volatilities converge in certain cases to either the lower bound,

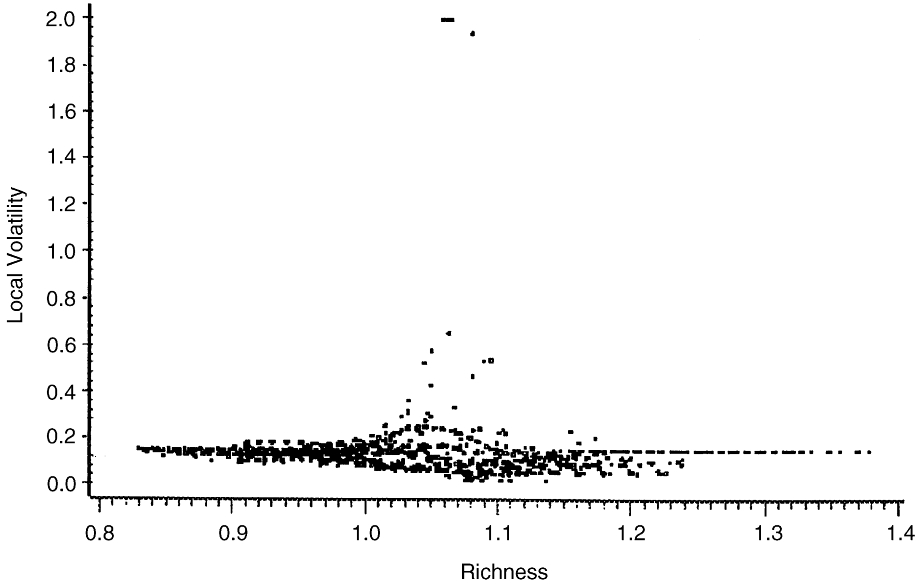

Fig. 4. Plot of local volatility against richness (underlying futures quote divided by strike price). The local volatility is defined to be the LPE parameter estimate of the volatility in the Black model. The Black model is used to approximate locally the covariance between DAX option prices, on the one hand, and DAX futures quotes and interest on overnight bunk deposits, on the other hand. Only observations with a maturity between 60 and 64 are included. Maturity is measured in number of intervals of 300 futures transactions (each interval corresponds to $1\frac{1}{2}$h calendar trading time, on average; this corresponds to slightly over $\frac{1}{4}$ trading day).

set equal to 0.02, or the higher bound, set equal to 2.00. Since maturity is expressed as a fraction of the year in the estimation, a volatility of 2.00 corresponds to an unrealistic volatility estimated of 200% per annum. It is not uncommon to estimate an extremely high or low implied volatility from a single observation. It is much surprising to obtain such extremes from the information contained in the co-movements between options and futures prices knowing that the sample is required to contain at least 50 observations.

The cause of this anomaly deserves additional investigations. The performance of the LPE as a hedging procedure is examined first.

4.3. Performance of the LPE hedging procedure

The performance of the LPE hedging procedure is assessed by comparing two self-financed delta hedging strategies. The first involves the local volatility and the LPE hedge ratios. The second involves implied volatility and the Black hedge ratios. They are referred to as the LPE hedge and the market hedge, respectively.

Table 5 reports descriptive statistics on the values of the hedge portfolios at maturity. The differences are small. LPE appears to perform as well as standard delta hedging. A much tougher test is to compare the hedge errors, i.e., the difference between the value of the hedge portfolio and the synchronous option value. Table 6 provides descriptive statistics for the hedge errors observed at maturity and at 180 and 90 transaction time units before maturity. As Table 6 suggests, the performance of the LPE hedge is far worse than the market hedge. Whereas the market hedge errors do not exceed 5%, the LPE hedge errors can be as large as 26% of the strike price.

Fig. 5 plots the LPE hedge errors at maturity against the moneyness of the option. The largest errors occur for options that are near-the-money where non-linearity

Table 5
Values of hedge portfolios at maturity

	Basis for delta	
	LPE	Black model
N	252	252
Average	0.0848	0.0742
s.d.	0.0952	0.0901
Skewness	1.04	1.28
Kurtosis	3.0	3.9
Minimum	−0.0108	−0.0114
5%	−0.0048	−0.0014
95%	0.2768	0.2617
Maximum	0.3921	0.3921

Remarks. Hedge portfolios are constructed on the basis of a self-financed modified delta hedge. The delta is computed from either the local volatility estimated using LPE, or the volatility implied from the Black model and the last transaction price for an option with similar richness and maturity. N denotes sample size. All value are expressed as a percentage of the strike price (premia).

Table 6
Hedge errors

	Maturity					
	180		90		0	
	Basis for delta		Basis for delta		Basis for delta	
	LPE	Black model	LPE	Black model	LPE	Black model
N	181	181	251	251	252	252
Average	0.0140**	0.0024**	0.0145**	0.0036**	0.0155**	0.0049**
s.d.	0.0451	0.0086	0.0442	0.0100	0.0438	0.0102
Skewness	3.87	1.09	3.85	1.05	3.88	0.91
Kurtosis	19.1	7.1	19.1	6.1	19.5	6.4
Minimum	−0.0349	−0.0209	−0.0231	−0.0262	−0.0198	−0.0215
5%	−0.0125	−0.0101	−0.0107	−0.0114	−0.0111	−0.0101
95%	0.0942	0.0166	0.0962	0.0209	0.0980	0.0249
Maximum	0.2561	0.0412	0.2620	0.0503	0.2611	0.0533

Remarks. Hedge portfolios are constructed on the basis of a self-financed modified delta hedge. The delta is computed from either the local volatility estimated using LPE, or the volatility implied from the Black model and the last transaction price for an option with similar richness and maturity. The hedge error is computed as the difference between the value of the hedge portfolio and the market price of the option. All priced and values are expressed as a percentage of the strike price (premia). N denotes sample size. Maturity is measured in number of intervals of 300 futures transactions (each interval corresponds to $1\frac{1}{2}$ h calendar trading time, on average; this corresponds to slightly over $\frac{1}{4}$ trading day). ** indicates: significant at the 1% level (two-tailed t-test).

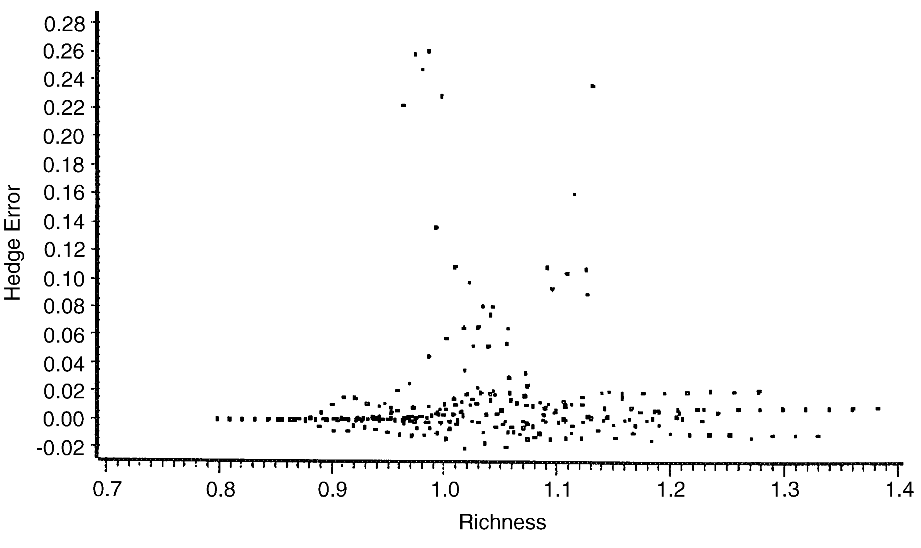

Fig. 5. Plot of hedge error (at maturity) against richness (underlying futures quote at maturity divided by strike price). The hedge error (at maturity) is defined to be the difference between the value of the hedge portfolio (at maturity) based on LPE and the expiration value of the option. All prices and values are expressed as a percentage of the strike price (premia).

matters most. This may question the ability of the LPE delta hedges to capture non-linearity. It is not clear that non-linearity per se contributes to the poor performance of LPE delta hedging. Rather, the large errors are likely to stem from the anomalous behavior of the local volatility. The often excessively high or low local volatility observed in Fig. 4 produce futures positions that are not optimal to hedge the risk of the DAX index options. The extreme values reached by local volatility are obtained for near-the-money options, explaining why the poor performance of the LPE hedges are mostly observed for these options.

4.4. The anomalous behavior of the local volatility estimate

Possible explanations for the poor performance of the LPE hedging procedure should focus on the abnormally low or high values reached by the local volatility estimates. Ignoring the return on interbank deposits, the local volatility is estimated by running the weighted non-linear least-squares regression,

$$\Delta C_i = x_{\mathrm{F}}^{\mathrm{B}}(z_i; \sigma)\Delta F_i + \eta_{i,}$$

with $i = 1, \ldots, N$. As is explained in Section 2.4, the weights are determined by the distance between z_i and z, a vector that contains the moneyness and the maturity of the option as well as the interest rate level.

4.4.1. The anomalous behavior of the local volatility estimate: an example

The weighted non-linear least-squares fit obtained by regressing the changes in the option prices onto the changes in the futures prices is illustrated in Fig. 6. The regression is performed on the training sub-sample used to compute the LPE option price of period 78 depicted in Fig. 2. The effective sub-sample includes 54 observations. Nothing unusual is observed. As option pricing predicts, the relationship between options and futures prices changes is almost linear due to the short observation interval. Still, with a moneyness of 0.80, the slope of the weighted least-squares regression line, equal to 0.38, is too high. The slope corresponds to a local volatility of 0.23, more than twice the local volatility implied by the estimated co-movement of options and futures prices around the same period.

Finally, Fig. 2 reveals that the spike in the local volatility estimates is short-lived. The local volatility estimate increases quickly from 0.10 at the transaction times $t = 71, \ldots 76$, to 0.15 at the time $t = 77$, increases to 0.23 at time $t = 78$, levels off at 0.22 until time $t = 79$ and rapidly drops to 0.14 at times $t = 80, 81$ and back to 010 at time $t = 82$.

Two types of arguments are put forward to explain the anomalous behavior of the local volatility estimate. The first line of arguments questions the econometric validity of the approach. The second questions the robustness of the option pricing model.

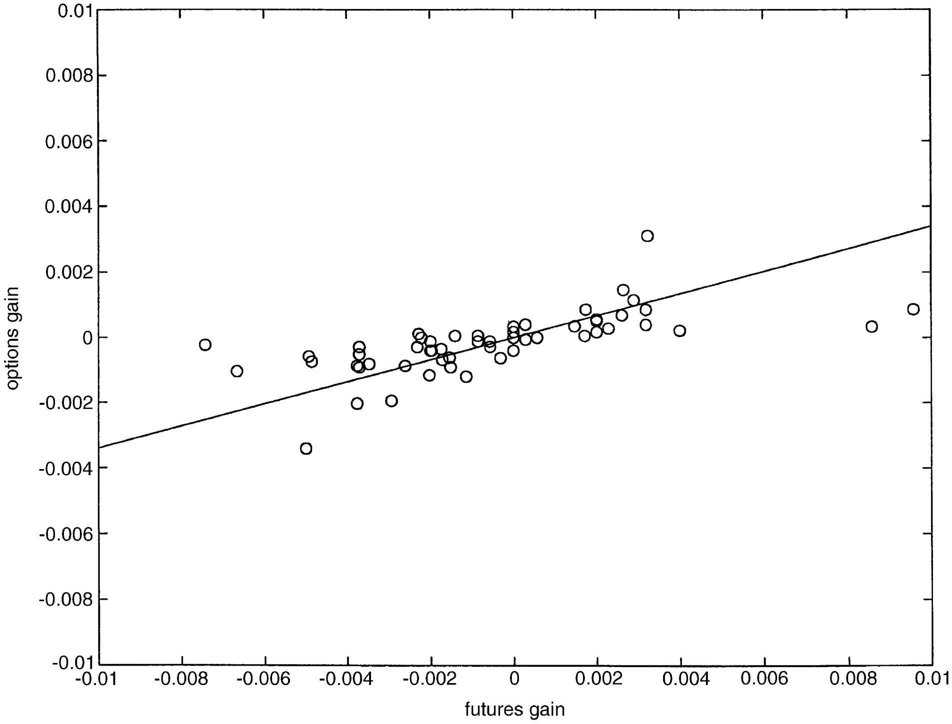

Fig. 6. Plot of options gains (changes in options prices) against futures gains (changes in futures prices) for the subsample of the training sample that is used to compute the LPE option price of period 78 depicted in Fig. 2. Gains are normalized by the strike price. The solid line is the (weighted) least squares fit.

4.4.2. Econometric explanations

Estimation biases, sampling error or omitted risk factors are three possible explanations for the aberrations in the behavior of the local volatility estimates. Each alternative is discussed below.

First, consider estimation biases. Non-synchronicity and/or the bid-ask bounce are known to generate an error-in-variables problem. As is discussed in Section 3.2, the sampling technique is designed to avoid estimation biases caused by the lack of synchronicity between the options and the futures markets. From this standpoint, the independent variable of the regression is measured without error. Additional evidence is reported in Table 2 that documents the lack of first-order autocorrelation in the futures price changes. This suggests that the bid-ask bounce does not generate any estimation biases. An error-in-variables problem is unlikely to generate the abnormal values observed for the local volatility estimates.

Second, consider sampling error. This error depends on the size of the training sample. There is no guarantee that enough observations receive positive weight in a local estimation procedure with a fixed bandwidth. The solution, discussed in Section 2.4, is to increase the bandwidth by 10% until the effective sample contains a

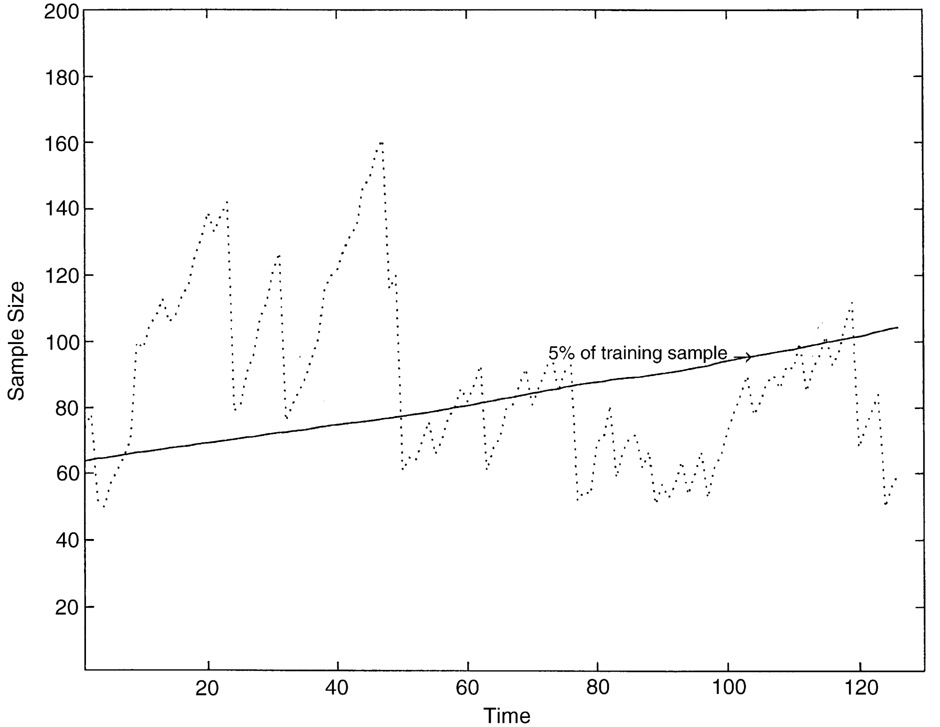

Fig. 7. Time series of the number of observations in the training sample that receive positive weight in LPE. Results for the 6/93 call option contract with strike 1675 are shown. The solid line indicates 5% of the total (training) sample size.

minimum of 50 observations. Fig. 7 illustrates this variable bandwidth technique for the contract expiring on 6/93 with a strike of 1675. The solid line depicts a value equal to 5% of the training sample size while the dotted line displays the evolution of the effective sample size. Sampling error is unlikely to generate the abnormal values observed for the local volatility estimates.

Third, consider omitted risk factors. Abnormally high or low local volatility estimates may reflect the presence of risk factors beyond futures price changes that are correlated with the latter, such as, for example, stochastic volatility. The estimation procedure implements a local projection of option price changes onto futures price changes and bond price changes. The contribution of the omitted risk factors to option price changes has two components, depending on whether their variability is correlated or not with the included risk factors. The contribution to the changes in the option prices of the changes in the omitted risk factors that are uncorrelated with the included risk factors is picked up by the error term. In contrast, the contribution to the changes in the option prices of the changes in the omitted risk factors that are correlated with the included risk factors is picked up by the projection coefficients, i.e., the hedge ratios.

The presence of omitted correlated risk factors could potentially explain the unusual hedge ratios. This explanation is not fully satisfactory. If this were the correct explanation, one would expect the hedging performance of the LPE approach to be far better. In particular, it would outperform the simple delta hedging strategy based on the Black implied volatility estimates, as the latter ignores the omitted risk factors entirely. In fact, the poor hedging performance of the LPE approach is an indication that omitted risk factors cannot explain the abnormal hedge ratios.

4.4.3. Theoretical explanations: the failure of option pricing models

The anomalous behavior of the local volatility estimates could be attributed to a possible inconsistency between the option pricing model and the observed option price dynamics.

The Black (1976) model is used in the regression to estimate locally the hedge ratio $x_F^B(z_i, \sigma)$. Though the functional form offered by the Black (1976) model is assumed to hold well locally, the model imposes severe constraints on the hedge ratio. This is illustrated in Fig. 8 that displays $x_F^B(z_i, \sigma)$ as a function of the richness, defined as the

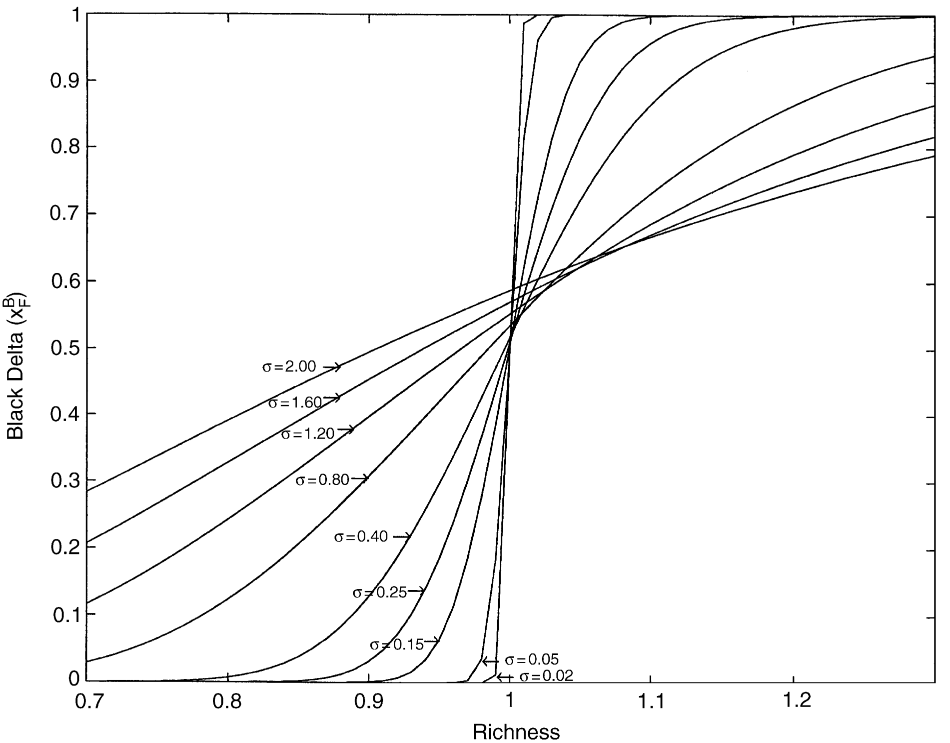

Fig. 8. Plot of the Black (futures) hedge ratio as a function of richness (underlying futures quote divided by strike price), for different volatilities. Volatilities are annualized. The maturity is set equal to 0.05 year, which is roughly equal to the maturities of the options for which local volatilities are plotted in Fig. 4.

ratio of the underlying futures quote divided by the strike price. The maturity is set equal to 0.05 year and the volatility, σ, is let to vary between a low of 0.02 and a high of 2.00, the two bounds imposed in the numerical non-linear optimization procedure. Three results emerge from Fig. 8.

- First, the Black hedge ratio is always constrained to be between zero and one. The actual co-movements between options and futures prices may not agree. If the local co-movement is negative for out-of-the-money options, LPE generates a corner solution and set the local volatility equal to the lower bound 0.02. A similar corner solution is observed when the local co-movement is higher than 1.0. Fig. 4 displays many cases of a local volatility estimate equal to 0.02. They occur for a richness above 1.0. This suggests that the local co-movement between options and futures prices is far higher than the maximum allowed by the Black (1976) model.
- Second, there exists an upper bound on the co-movement between options and futures for out-of-the-money options. It is far below 1.0. As is displayed in Fig. 8, the upper bound is obtained for a volatility of 2.0. The corresponding hedge ratio is at most 0.6 for a maturity of 0.05 years. If the actual co-movement is higher, LPE generates a corner solution and estimate of local volatility is equal to 2.0. Though these estimates are not observed in Fig. 4, they are obtained in the wider sample.[24]
- Third, there exists a correspondingly tight lower bound on the co-movement for in-the-money options. As the richness increases, the volatility value for which this lower bound obtains increases as well. If the actual co-movement is below the lower bound, excessively high volatilities, not all necessarily equal to 2.0, are obtained. They can be observed in Fig. 4.

The corner solutions obtained for the local volatility suggest that, unlike the Black hedge ratios, the actual hedge ratios do not always take values between 0.0 and 1.0. The co-movements between the observed options and futures price changes often contradict the theoretical restrictions. More complex option pricing models that let hedge ratios take values below 0 for out-of-the-money options and above 1.0 for in-the-money-options are necessary to capture the observed co-movements.

4.4.4. Local parametric versus local polynomial estimation

The previous discussion reveals that local parametric estimation cannot produce consistent estimates when the parametric model imposes restrictions that are not satisfied by the data. These constraints may call for a less parametric technique to estimate the hedge ratios. Local polynomial estimation, as an example of a local non-parametric technique, is a possible alternative. Local non-parametric estimation (LNPE) merely requires smoothness conditions to produce consistent estimates. While the simulations performed in Bossaerts and Hillion (1997) suggest that local

[24] Fig. 4 only displays estimates of local volatilities for a range of maturities between 60 and 64 transaction time periods before maturity.

polynomial estimation leads to an inferior hedging performance, it is worth asking how it would fare in the present case with market data.

The following version of local polynomial estimation is implemented. The hedge ratios are estimated by fitting locally the projection coefficients of the regression equation,

$$\Delta C_i = x_{\mathrm{B}}^{\mathrm{B}} \Delta B_i + x_{\mathrm{F}}^{\mathrm{B}} \Delta F_i + \eta_i, \tag{7}$$

with $i = 1, \ldots, N$. As before, the weights given to observation i in the linear least-squares procedure depends on the distance between the option strike, its maturity, the interest rate level and the option for which the hedge ratio is being estimated. The implementation of this procedure amounts to generalized linear least squares. The estimates can be estimated efficiently without any of the usual numerical problems encountered in non-linear estimation, such as LPE.[25]

LNPE generates even more aberrant hedge ratios. As an example, consider the call contract with expiration 3/94 and strike price 2050. Fig. 9 displays the evolution of the hedge ratios, estimated using LPE and LNPE, respectively. The hedge ratios obtained from the Black (1976) model are also displayed for comparison.[26] For lengthy periods, the LNPE hedge ratios differ substantially from their LPE or synchronous counterparts. In contrast, the LPE hedge ratios are close to their synchronous counterparts with a few exceptions.

The serious discrepancy between the data and option pricing theory cannot be resolved by simply switching to local polynomial estimation. In fact LNPE fares worse than LPE, confirming the previous findings of Bossaerts and Hillion (1997).

4.5. Discussion

The discrepancies between theoretical and actual patterns of co-movements of the derivatives and underlying securities prices imply an inconsistency between the data and the Black (1976) option pricing model. Bakshi et al. (2000) observe similar anomalous co-movements in the options and futures price changes in the S&P500 index option market. They find that when the underlying index goes up (down), quite frequently call prices go down (up) and put prices go up (down). This is referred to in their paper as "type 1" violations. In addition, they report evidence of hedge ratios above 1.0 for calls and below -1.0 for puts, referred to as "type 4" violations. While they attribute the "type 4" violations to market microstructure effects, they conclude that the "type 1" violations suggest that one-dimensional diffusion option models are inconsistent with the observed option price dynamics.

[25] The results obtained with LPE and LNPE are expected to differ for the following reason. In the LNPE approach, the estimates are obtained by setting to zero the weighted average cross-product of the equation error and the explanatory variables. In contrast, in the LPE approach, the estimates are obtained by setting to zero the weighted average cross-product of the equation error and the explanatory variables multiplied by the sensitivity of the coefficients with respect to the volatility parameter. See Eq. (6).

[26] The Black volatility is implied from the most recent option transaction price with a similar strike and maturity.

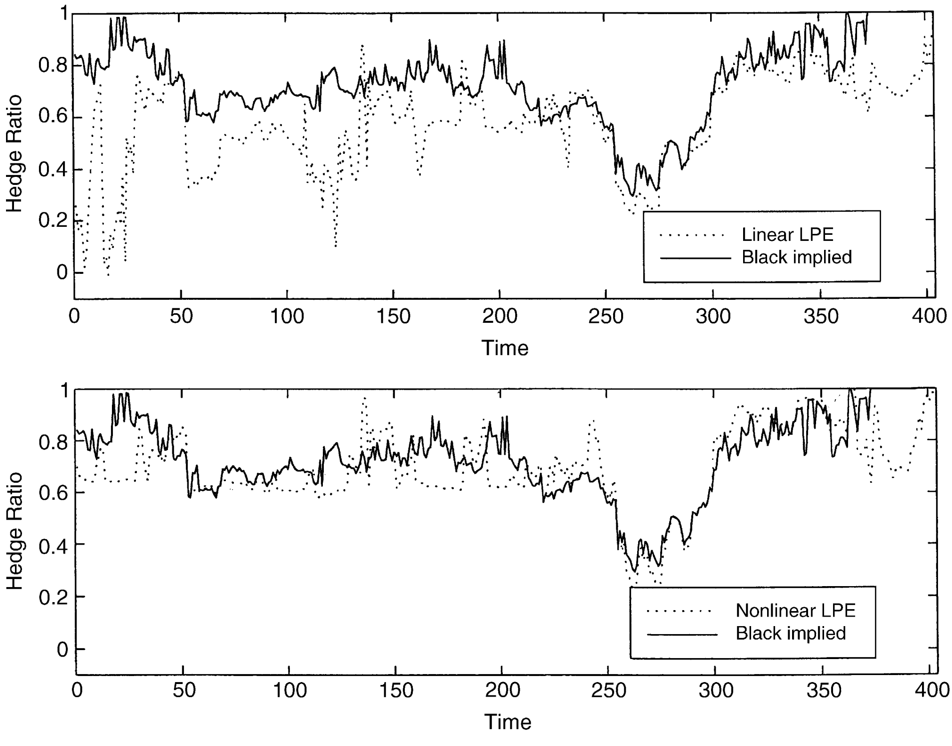

Fig. 9. Comparison of the evolution of the estimated hedge ratios over time, 3/94-2050 contract. Solid line depicts the hedge ratio computed on the basis of the volatility implied from the most recent transaction price of an option with similar richness and maturity, using Black's formula. The dashed line depicts the hedge ratio estimated using nonlinear LPE (top panel) and linear LPE (bottom panel). In linear LPE, the hedge ratio is directly estimated using (locally) weighted least squares. In nonlinear LPE, the hedge ratio is estimated nonlinearly using the Black model as local approximation. The latter procedure was also employed to generate the results reported in earlier tables and figures.

The empirical evidence reported in this paper and in Bakshi et al. (2000) paper can be better understood in light of the theoretical results derived by Bergman et al. (1996). They show that whenever the underlying stock price follows a diffusion process whose volatility depends only on time and the concurrent stock price, a call price is always increasing and convex in the stock prices, the so-called "monotonicity" property. When volatility is stochastic, or the stock price process is not a diffusion but is instead either discontinuous or non-Markovian, a call price can be a decreasing concave function of the stock price over some range. In the latter case, the call duplication may involve shorting the underlying stock and lending.

These results rest on the "no-crossing" property, the fact that for a given realization of the Brownian motion driving the risk-neutral stock price process, the realized value of the process on the option's maturity date is increasing in its starting value. A consequence of the "no-crossing" property is that in the one-dimensional diffusion setting, the option price inherits the monotonicity from the contractual payoff function. A stochastic process that is not a diffusion need not feature the no-crossing property and option prices need not exhibit the inherited monotonicity, i.e., the call price may be decreasing or concave over some underlying price range. Also, in contrast to the one-dimensional case, the no-crossing property need not hold in the two-dimensional case.

The theoretical results of Bergman et al. (1996) suggest that stochastic volatility and/or jumps in the price process may explain the anomalous co-movements in the options and futures prices changes. It has been widely documented that stock returns exhibit both stochastic volatility and jumps. There are two reasons why stochastic volatility is unlikely to explain the anomalous co-movements. First, as discussed in Section 4.4.2, omitted risk factors are taken care of by the estimation procedure. Second, as discussed in Section 2.3, stochastic volatility is controlled for by hedging in transaction time. This conclusion is consistent with the findings of Bakshi et al. (2000) who investigate whether the stochastic volatility model of Heston (1993) explains the anomalous co-movements reported in the S&P500 index option market. In Heston's (1993) model, volatility is not necessarily correlated with the underlying asset price. This implies that option prices can move independently of the underlying price. Bakshi et al. (2000) conclude that the stochastic volatility model falls short of explaining the anomalous price movements. About 47% of the "type 1" violations are consistent with the predictions of the stochastic volatility model.

This leaves price jumps as a possible explanation for the anomalous co-movements in the options and futures price changes. From a theoretical point of view, price jumps may cause violations of the "no-crossing" property as shown by Bergman et al. (1996). From an empirical point of view, there is increasing evidence that jump diffusion models fit option prices better than standard models. Bates (2000) find that stochastic volatility models require extreme parameters, in particular high volatility of volatility, that are implausible given the time-series properties of option prices. Bates (2000) reports evidence that there is a strong tendency for jump and non-jump risk to rise whenever the market falls and concludes that stochastic volatility/jump diffusion models are more compatible with observed option prices.[27]

5. Conclusion

The goal of the paper is to investigate the pricing and hedging performance of LPE using a small set of underlying assets. The LPE estimation technique combines both parametric and non-parametric methods. LPE is a kernel smoother for non-parametric regression that uses prior information on regression shapes in the form of a parametric model. By encompassing non-parametrically a parametric model, local

[27] See Section 4. A of Bergman et al. (1996) for an illustration of a mixed diffusion-jump process for which the "no-crossing" property is violated and for which the call price is not everywhere increasing and convex in the underlying asset price. The process is a non-proportional process, such that below a certain price level, it behaves like a mixed diffusion-jump process and, above it, it grows at the risk-free rate. This process is consistent with Bates' (2000) finding that the jump probability is a function of the price level.

parameters are estimated using information from a neighborhood of the point of interest.

The approach relies entirely on the local correlation pattern between the option and the underlying securities prices. The local correlation is captured by means of a parametric option pricing model. In contrast to the traditional approach, there is no need to exploit the information contained in the time-series of the underlying asset prices or to rely on implied volatility. In theory, the technique should work well, because correlation between the derivatives and the underlying securities prices is at the heart of arbitrage-based option pricing models.

The performance of local parametric analysis is tested empirically using a dataset of DAX index options. On average, LPE-based prices predict final option payoffs as well as market prices. This is a remarkable result given the fact that the LPE prices just exploit the correlation between the options and the underlying securities price changes, and do not include information in current option prices. Yet, they differ substantially from market prices. In addition, self-financed delta hedging strategies based on the LPE estimates of the hedge ratios under-perform those based on implied volatilities.

Frequent anomalous co-movements between DAX options and futures prices are at the root of the poor hedging performance and the deviations between the market and the LPE-based prices. Option pricing models, such as the Black (1976) model used in this paper, restrict the coefficient from a local projection of option and futures price changes to be between zero and one. Yet, many violations of this constraint are observed in the dataset. The issue of whether such anomalous co-movements disappear over longer time intervals deserves further investigation.

The results reported in this paper differ significantly from those reported in Bossaerts and Hillion (1997). They investigate the performance of LPE-based hedging in an environment where a continuous-time option pricing model holds exactly but where hedging is constrained to take place in discrete time. In a simulated Black-Scholes world, the LPE-based delta hedging strategy dominates significantly the Black-Scholes delta hedging strategy. In contrast, with market data, the LPE-based delta hedging strategy fares much worse.

The evidence of major discrepancies between theoretical and actual patterns of co-movements of the derivatives and underlying securities prices has important implications for theoretical developments of option risk management models. They imply that hedging risks can temporarily be far bigger than predicted because correlations increase or decrease far beyond theoretically feasible values.

References

Bakshi, G., Cao, C., Chen, Z., 2000. Do call prices and the underlying stock always move in the same direction? Review of Financial Studies 13, 549–584.

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

Bergman, Y.Z., Grundy, B.D., Wiener, Z., 1996. General properties of option prices. Journal of Finance 51, 1573–1610.

Black, F., 1976. The pricing of commodity contracts. Journal of Financial Economics 3, 167–179.

Bossaerts, P., Hillion, P., 1997. Local parametric analysis of hedging in discrete time. Journal of Econometrics 81, 243–272.

Bossaerts, P., Werker, B., 1997. Martingale-based hedge error control. In: Rogers, C., Talay, D. (Eds.), Numerical Methods In Finance. Cambridge University Press, Cambridge, UK.

Bossaerts, P., Ghysels, E., Gouriéroux, C., 1996. Arbitrage-based pricing when volatility is stochastic. Caltech working paper 977.

Bossaerts, P., Hafner, C., Härdle, W., 1996. Foreign exchange rates have surprising volatility. In: Robinson, P.M., Rosenblatt, M. (Eds.), Athens Conference Applied Probability and Time Series, Vol. II (Time Series analysis). Springer, New York.

Boyle, P., 1977. Options: a Monte Carlo approach. Journal of Financial Economics 4 (3), 323–3338.

Duffie, D., Richardson, H.R., 1991. Mean-variance hedging in continuous time. Annals of Applied Probability 1, 1–15.

Dumas, B., Fleming, J., Whaley, R.E., 1998. Implied volatility functions: empirical tests. Journal of Finance 53, 2059–2106.

Föllmer, H., Sondermann, D., 1986. Hedging of contingent claims under incomplete information. In: Mas-Colell, A., Hildenbrand, W. (Eds.), Contributions to Mathematical Economics. North-Holland, Amsterdam.

Gouriéroux, C., Laurent, J.P., Pham, H., 1995. Quadratic hedging and numèraire. CREST working paper.

Harris, L., 1986. Cross-security tests of the mixture of distribution hypothesis. Journal of Financial and Quantitive Analysis 21, 39–46.

Harris, L., 1987. Transaction data tests of the mixture of distribution hypothesis. Journal of Financial and Quantitative Analysis 22, 127–142.

Hjort, N.L., Glad, I.K., 1995. Nonparametric density-estimation with a parametric start. Annals of Statistics 23 (3), 882–904.

Hull, J., White, A., 1987. The pricing of options on assets with stochastic volatilities. Journal of Finance 3, 281–300.

Hutchinson, J., Lo, A., Poggio, T., 1994. A nonparametric approach to pricing and hedging derivative securities via learning networks. Journal of Finance 49, 851–890.

Longstaff, F., Schwatz, E., 2001. Valuing American options by simulation: a simple least-squares approach. The Review of Financial Studies 14 (1), 113–147.

Ross, S., 1978. Mutual fund separation and financial theory—the separating distributions. Journal of Economic Theory 17, 254–286.

Schweizer, M., 1995. Variance optimal hedging in discrete time. Mathematics of Operations Research 20, 1–32.

Tauchen, G., Pitts, M., 1983. The price variability-volume relationship on speculative markets. Econometrica 51, 485–505.