

**LOCAL PARAMETRIC ANALYSIS
OF DERIVATIVES PRICING
AND HEDGING**

by

P. BOSSAERTS*
and
P. HILLION**

97/121/FIN

* California Institute of Technology, Pasadena, CA 91125, USA.

** Professor of Finance at INSEAD, Boulevard de Constance, 77305 Fontainebleau Cedex, France.

A working paper in the INSEAD Working Paper Series is intended as a means whereby a faculty researcher's thoughts and findings may be communicated to interested readers. The paper should be considered preliminary in nature and may require revision.

Printed at INSEAD, Fontainebleau, France.

Local Parametric Analysis Of Derivatives Pricing And Hedging*

Peter Bossaerts

Pierre Hillion

December 1997 (first draft: January 1997)

*First author's affiliation: California Institute of Technology; Mailing address: HSS 228-77, California Institute of Technology, Pasadena, CA 91125; Phone (626) 395-4028; e-mail: pbs@rioja.caltech.edu. Second author's affiliation: INSEAD; Mailing address: INSEAD, F-77305 Fontainebleau; Phone +33.1.6072-4214; e-mail: hillion@quartz.insead.fr. This version benefited from comments during seminars at UFSIA (Antwerp, Belgium) and ECARE (Brussels, Belgium). We are grateful to the *Deutsche Terminbörse* for sending us the transactions data, and to Wolfgang Bühler for the interbank deposit rates. The present version reports results from the 1992-3 dataset only; as diligent researchers, we have not touched the 1994 data yet. These will be analyzed once this version has had enough exposure to the criticisms of our colleagues.

Abstract

We present a novel methodology for the analysis of derivatives pricing in incomplete markets. Unlike existing procedure, we do not commit to a particular model for the prices of the underlying asset(s). Instead, we locally fit the hedge ratios from a parsimonious complete-market model to changes in the prices of options and underlying securities. We then compute prices from the locally estimated parameters. Thus, we extract information from the *local co-movement of prices of several securities*, as opposed to the standard procedure, which obtains parameter values either from time series estimation or by implying them from option prices directly. We illustrate the methodology on a dataset of DAX index options and futures transactions from the computerized German Futures Exchange, over the period 1992-94. We form our estimates entirely on the basis of lagged information. While they predict final payoffs with a precision that matches that of market prices, our prices often deviate substantially from synchronous market prices. Moreover, self-financed delta hedges based on our estimates of the hedge ratios perform worse than those with hedge ratios computed from market-implied volatilities. All this is caused by anomalous behavior of the locally fitted hedge ratios. For instance, the hedge ratios of in-the-money options are sometimes estimated to be so low that they imply a volatility above 200% (per year). Since a typical estimation is based on 5% of the "training" sample, with a minimum of 50 observations, sampling error cannot explain the anomalous behavior.

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

In a major departure from existing methodology, we do not commit to a particular model for the prices of the underlying asset. Instead, we fit the hedge ratios of a parsimonious complete-market model *locally*. This means that hedge ratios are estimated from changes in the prices of options with neighboring values of moneyness and maturity. We also limit the estimation to observations for which the recorded interest rate level was similar.

Under certain assumptions about actual pricing in the options market and the composition of the hedge portfolio, the estimate of the size of the hedge portfolio (i.e., dollar amount to be invested per option) is a consistent estimate of the correct value for the option under a quadratic loss function. The analysis is a variation on the ideas expounded in Bossaerts and Werker [1997].

Hence, we determine the optimal hedge portfolio and the right pricing directly from the pattern of correlation between derivatives prices and the values of the underlying assets. We do not have to specify a continuous-time stochastic process for the latter; we do not have to solve analytically intractable partial differential equations; we can avoid the well-known difficulties in estimating (volatility) parameters of continuous-time processes from discrete-time data.

As a matter of fact, since theoretical derivatives analysis always starts with a *correlation analysis* of the payoffs on several assets, we find our approach much more natural than the standard empirical procedure of obtaining the necessary hedge ratios and hedge portfolio values from a pure *time series* analysis of the values of the underlying assets.

The estimation technique that our methodology is based on, namely, *local parametric estimation* (LPE), has only recently been suggested in the statistics literature as an attractive alternative to the more agnostic nonparametric (kernel) estimation procedure. See Hjort and Glad [1995]. It was introduced in the analysis of hedging in Bossaerts and Hillion [1997].

When measuring net gains, we advocate the use of, e.g., the clock defined by the number of transactions of the underlying security. This change in time scale cleans the payoffs of the underlying security for nonnormalities (this is also documented in Tauchen and Pitts [1983] and Harris [1986], [1987]). More importantly, it simultaneously reduces the randomness in volatility, which implies that fewer securities are needed when hedging the risk of a derivative. In this sense, the paper applies the ideas in Bossaerts, Ghysels and Gouriéroux [1996], where the impact of time changes on the hedging and pricing of derivatives is studied from a theoretical point of view.

We illustrate the methodology using a dataset of *all* DAX index options and futures transactions on the computerized German Futures Exchange (*Deutsche Terminbörse*, DTB), over the period 1992-94. We use the (liquid) futures contract, as well as overnight positions in the money market (*Tagessgeld*) to hedge and price index options. Hedge portfolios are rebalanced every 1 1/2 hour on the transactions clock.

Our estimates will be based entirely on *past information*. This version reports results for the 1992-1993 part of our dataset only. We have left the 1994 data untouched. We keep it for future robustness checks. We will include the results from the 1994 data in the next version of the paper, as soon as this version has received sufficient comments.

We find that our methodology generates prices that predict payoffs at maturity as accurately as market prices. Our prices often deviate substantially from market prices, however. And a self-financed delta hedge based on our estimates does not perform as well as one based on volatilities implied from transaction prices and Black's formula. These results appear to be caused by anomalous co-movement of the prices of options and underlying futures contracts.

The paper is organized as follows. The first section presents the methodology. Section 2 discusses the DAX index options in the DTB dataset. Section 3 studies the pricing of these options using our methodology and compares the results with the transaction prices in the market. Section 4 analyzes the hedge performance of our procedure. Section 5 relates our findings to anomalous co-movement between options and futures prices. Section 6 concludes.

1 The Methodology

We proceed by first presenting the theoretical background, assuming that one has an appropriate statistical procedure to estimate any unknown parameters. Subsequently, we elaborate on the estimation of the latter and introduce the idea of local parametric estimation. Finally, we discuss the estimation interval, the desirability of implementing a time change before estimation, and a few other technicalities.

1.1 The Theory

Since Black and Scholes' seminal contribution, derivatives analysis has followed a standard pattern. Consider, for instance, the pricing of a call option written on a futures contract whose price process is binomial. Assume that interest rates, and, hence, returns on riskfree bonds, are constant.

The first step is to construct the portfolio of the underlying futures contract and bond investment which will perfectly hedge the risk of the option. Let ΔF_u and ΔF_d denote the possible payoffs (gains) on a (long) futures position over the next time interval. After a gain ΔF_u , the future futures quote is F_u ; after ΔF_d , it is F_d . Let ΔB denote the payoff on the bond investment. Posit momentarily that the call price path lies on the same binomial tree as the futures quote, and let ΔC_u and ΔC_d denote the possible payoffs on holding one call over the next time interval. To obtain ΔC_u and ΔC_d , assume that the future value of the call in either state is known (e.g., because the call expires, at which point these values equal $\max(0, F_u - K)$ and $\max(0, F_d - K)$, where K denotes the strike price), and assign an arbitrary initial value, C .

The hedging problem is to find the number of bonds x_B and the futures position (number of futures contracts) x_F needed to hedge the payoff on the option. Mathematically, one solves the following system of equations:

$$\begin{cases} \Delta C_u &= x_B \Delta B + x_F \Delta F_u, \\ \Delta C_d &= x_B \Delta B + x_F \Delta F_d. \end{cases} \quad (1)$$

The value of the hedge portfolio (number of dollars to invest), to be denoted C^* , equals $x_B B$, where B is the initial bond price. (The futures position does not require investment.) Since the hedge portfolio insures the payoff on the call perfectly, its value is also the *no-arbitrage value of the call*.

The second part of the analysis is a fixed-point (equilibrium) argument: the no-arbitrage value of the call, C^* , must also be the basis value for ΔC_u and ΔC_d , which we denoted C . Recomputing ΔC_u and ΔC_d based on C^* , re-solving (1), an updated value for C^* is obtained. One iterates the algorithm until a fixed point is reached, i.e., x_B and x_F , and, hence, C^* , settle.¹

The fixed-point solution also verifies the assumption that we temporarily made that option prices move along the same binomial tree as the futures quote (which *a priori* could only be justified on the terminal nodes, i.e., at expiration). In other words, we obtain a fully consistent theory of option pricing.

In the above case, changes in option prices could be hedged perfectly. In the language of derivatives analysis, markets are said to be complete. What about the *incomplete-markets case*? Earlier theoretical continuous-time analysis appealed to intertemporal asset pricing models to obtain equilibrium values for derivative assets. See, e.g., the popular pricing model under stochastic volatility in Hull and White [1987]. We will follow an alternative route, however, inspired by the more recent theoretical continuous-time analysis of incomplete markets, and closer to the above hedging arguments. Consult Föllmer and Sondermann [1986], Duffie and Richardson [1991], Schweizer [1994], Gouriéroux, Laurent and Pham [1995].

Position oneself at a point in time. One has information about asset prices in the marketplace which can be summarized by a (random) vector z . This vector includes items such as the moneyness and maturity of the derivative to be analyzed, as well as the level of the interest rate, which is no longer assumed constant. Consider

¹In the binomial case, the fixed point could actually be obtained in *one step*, because the argument could be based on hedging the *level* of the option price, instead of its payoff or net gain. Instead, our exposition follows the standard approach in continuous time, where optimal hedge portfolios are computed for (instantaneous) *changes* in option prices, i.e., for payoffs or net gains. The fixed point is obtained by means of Itô's Lemma.

the following (conditional) projection of net gains on the option (ΔC) onto net gains from holding one bond (ΔB) and one futures contract (ΔF):

$$\Delta C = a(z) + x_B(z)\Delta B + x_F(z)\Delta F + \eta, \quad (2)$$

where

$$\begin{aligned} E[\eta|z] &= 0, \\ E[\eta\Delta B|z] &= 0, \\ E[\eta\Delta F|z] &= 0. \end{aligned}$$

x_B and x_F are interpreted as hedge ratios (number of units of the bond and futures contract to be taken on). The *hedge error*, i.e., the difference between the payoff on the option and that on the hedge portfolio, equals $a(z) + \eta$, and, by construction, is uncorrelated with the payoff on the hedge portfolio.

Equation (2) represents a conditional projection, because the projection error is orthogonal to the regressors *conditional on* z . To make such a projection possible, the projection coefficients must be allowed to vary with z . In (2), we made this dependence explicit. In general, the functional relationship between the coefficients and z is unknown, and must be estimated. For the time being, we will assume that there is a convenient way to do so. In other words, we assume temporarily that we know $a(z)$, $x_B(z)$ and $x_F(z)$.

Because the hedge ratios are obtained as the coefficients in a (conditional) projection, the resulting hedge portfolio is *optimal* under a quadratic hedge error loss function. In other words, $x_B(z)$ and $x_F(z)$ provide the optimal number of bonds and futures contracts as a function of z for an arbitrageur who attempts to replicate the payoff of the option when she faces a quadratic loss function.

Since the second security in the hedge portfolio is a futures contract, which does not require an investment outlay, the value of the hedge portfolio is determined by the number of bonds times their price. Conditional on z , this means that the value of the optimal hedge portfolio, C^* , equals:

$$C^* = x_B(z)B. \quad (3)$$

Under the right assumptions, one can make stronger statements. Assume that investors have quadratic utility, and that, in the absence of the option, their optimal portfolios consists entirely of the bonds and futures contracts used in the hedging exercise. If $a(z)$ in (2) equals zero, then²

$$\frac{\Delta C}{C^*} = (x_B(z)\frac{B}{C^*})\frac{\Delta B}{B} + (x_F(z)\frac{1}{C^*})\Delta F + \epsilon, \quad (4)$$

with

$$\begin{aligned} E[\epsilon|z] &= 0, \\ E[\epsilon\frac{\Delta B}{B}|z] &= 0, \\ E[\epsilon\Delta F|z] &= 0. \end{aligned}$$

If the market price of the option equals C^* , (4) means the following: the market offers an average return on the option which is the same as that of a portfolio of bonds and futures contracts, yet its variance is higher. Hence, at a price C^* , the option would not improve the performance of the optimal portfolio of an investor with quadratic utility, which we assumed to be composed of the bonds and futures contracts in the hedge portfolio. There would be zero demand for the option.³ If the option is truly derivative, in the sense that its net supply equals zero, the options market is in equilibrium. Hence, C^* is the equilibrium value of the option.

²The implication is true only if C^* can be obtained from the elements in z . But this is the case: $C^* = x_B(z)B$, and B can be computed from the interest rate, which is assumed to be in z .

³For a proof, see Ross [1978].

We conclude that if, (i) in the absence of the option, mean-variance optimal portfolios consist exclusively of the bonds and futures contracts used in the hedge portfolio, and (ii) $a(z)$ in (2) equals zero, then $C^* = x_B(z)B$ not only provides the value of the hedge portfolio, it also equals the value of the option under quadratic loss. In other words, our strategy produces an estimate of the value of an option in an incomplete market for investors with quadratic utility.

The assumption that mean-variance optimal portfolios invest entirely in the bonds and futures contracts of the hedge portfolio is strong, but could be justified when studying the pricing of *index options* (options on some broad market index). Indeed, it is known under what conditions optimal portfolios can be constructed from a few securities such as a bond portfolio and an index futures contract. See the portfolio separation results in Ross [1978]. In other cases, the range of securities in the hedge portfolio would have to be extended in order to cover all those that generate mean-variance optimal portfolios.

1.2 Local Parametric Estimation

In what follows, we will first discuss the case where $a(z)$ in (2) equals zero. As just mentioned, if this restriction is satisfied, and mean-variance optimal portfolios are composed of only the bonds and futures contracts in the hedge portfolio, $x_B(z)B$ immediately provides an estimate of the value of the option under quadratic loss.

To estimate $x_B(z)$ and $x_F(z)$, we proceed as follows. Assume that the true hedge ratios are smooth functions of z . Let z consist of the options' moneyness, maturity, as well as the interest rate level. Now take a parsimonious, parametric, complete-markets model of the hedging and pricing of options, and fit its hedge ratios *locally*. For instance, take the Black model of the pricing of options on futures contracts (see Black [1976]). The number of units of bonds to purchase and the number of futures contracts to acquire in order to hedge the payoff on the option are both functions of z and the volatility of the futures quote, σ . We reflect this by writing the Black hedge ratios as follows:

$$x_B^B(z; \sigma), \quad x_F^B(z; \sigma). \quad (5)$$

To estimate x_B and x_F at z , we take N observations with values for moneyness, maturity and interest rates equal to z_1, \dots, z_N , and corresponding payoffs $(\Delta C_1, \Delta B_1, \Delta F_1), \dots, (\Delta C_N, \Delta B_N, \Delta F_N)$, and implement weighted nonlinear least squares to estimate σ in:

$$\Delta C_i = x_B^B(z_i; \sigma) \Delta B_i + x_F^B(z_i; \sigma) \Delta F_i + \eta_i, \quad (6)$$

$i = 1, \dots, N$. The weight to be given to observation i in the nonlinear least squares procedure depends on the distance between z_i and z . The weights for observations with z_i sufficiently far away from z are set equal to zero.

This provides an estimate of x_B at z as

$$x_B^B(z; \hat{\sigma}_z),$$

where $\hat{\sigma}_z$ is the weighted nonlinear least squares estimate of σ . The latter is subscripted with z , to stress that it may change with z . Therefore, an appropriate name for $\hat{\sigma}_z$ is *local volatility*.

The weights in the nonlinear least squares exercise are determined by means of a *kernel function*. Loosely speaking, the neighborhood around z for which this kernel function generates nonzero weights is referred to as *bandwidth*. The larger the bandwidth, the more accurate the estimation (because of the increase in observations), but, if the hedge ratios deviate substantially from the ones provided by the Black model, the higher the bias. As the number of observations increases, the bandwidth can be reduced without loss of accuracy, and any bias will be reduced.

Bossaerts and Hillion [1997] provide a detailed discussion of the estimation technique. The theoretical properties are discussed in Hjort and Glad [1995]. The above technique, whereby a *parametric* model is fit *locally*, differs from the usual nonparametric kernel estimation procedure (see, e.g., Härdle [1990]). In the latter, polynomials (most often of zero order)⁴ are fit locally. The advantage of fitting a parametric models is that they may capture much better the local curvature in the hedge ratios as a function of z than would generic functions such as polynomials. This is best seen with an extreme example. Suppose the hedge ratios of the parametric model can be fit globally.

⁴For an application of *higher-order polynomial* kernel estimation in finance, see Bossaerts, Hafner and Härdle [1996].

In particular, let us assume that, for some choice of σ , $x_B(z)$ and $x_F(z)$ in (2) equal $x_B^B(z; \sigma)$ and $x_F^B(z; \sigma)$, respectively, for all z .⁵ In that case, fitting the Black model locally would obviously be far more efficient than fitting any polynomial locally.

In a study of hedging in discrete time, Bossaerts and Hillion [1997] document that hedge portfolio weights estimated by locally fitting a parametric model outperform those obtained from the Black-Scholes formula *even* in an environment where *the Black-Scholes model holds exactly*. The present paper is an attempt to verify whether LPE performs equally well in a *real-world setting*.

While LPE may seem to be an automatic and robust procedure, it does require one to choose a bandwidth. This choice is at once important and delicate. Too large a bandwidth produces estimates of option prices and hedge ratios which may be too close to the (usually invalid) parametric model that one is fitting locally. Too small a bandwidth may reduce the bias, yet introduces huge sampling variance. In the exercise we report below, we chose the bandwidth on the basis of the hedging performance of LPE on a subset of the data (all options expiring 3/92).

Let us briefly return to estimation of the intercept $a(z)$ in (2). Theory has little to say about what exactly $a(z)$ could be. Usually, it is interpreted as a systematic “pricing error,” because its presence would indicate a discrepancy between the average net gain on the option and the optimal hedge portfolio. But the latter may not be optimal under a nonquadratic loss function. If option prices reflect different preferences, it should be no surprise that the average net gain on the option differs from that on the quadratic-loss hedge portfolio.

Absent clear indications about its functional form, $a(z)$ should be estimated by locally fitting polynomials. In its extreme form, when a zero-order polynomial is estimated, this would coincide with the familiar nonparametric kernel estimation.

1.3 Estimation Interval

Before illustrating our technique, we need to be more specific about the time interval over which ΔC , ΔB and ΔF ought to be recorded.

Of course, this interval should be small, but not to the extent that transaction costs would dominate the hedging performance. The issue is more one of finding the appropriate time interval such that the payoff on the option can be hedged successfully with a very limited set of securities (in the present case, bonds and futures contracts).

Start from the observation that the prices of many securities exhibit *stochastic volatility*. In a simple extension of the Black-Scholes environment that accomodates stochastic volatility, Bossaerts, Ghysels and Gouriéroux [1996] illustrate how the payoff on derivatives can only be replicated perfectly if the number of securities in the hedge portfolio is allowed to increase without bound. A straightforward time transformation, however, generates a perfect hedge with only a few securities. The result obtains because the time change annihilates the random volatility. (It does introduce random expiration times, however!)

The result in Bossaerts, Ghysels and Gouriéroux [1996] calls for rebalancing according to a clock under which the price of the underlying security shows little or no stochastic volatility. In the application to be discussed later, a simple switch to the transactions clock (here, the clock defined by the number of transactions in the futures market) reduced the randomness of the volatility to a minimum. The performance of hedging based on the transactions clock can be expected to be improved correspondingly. The incremental value from adding securities that insure remaining volatility changes is minor.

In our application, the switch to the transactions clock happened to simultaneously purge nonnormalities from the data.⁶ In particular, when observed at a pace in transactions time that neutralizes the impact of overnight intervals, futures price changes appear close to normal. This effect, together with the reduction in the randomness of volatility, imply that the behavior of futures prices is close to that assumed in the Black model.

⁵Bossaerts and Werker [1997] outline an incomplete market where this would obtain. See their Section 5.

⁶This effect has been observed elsewhere. See, e.g., Tauchen and Pitts [1983], Harris [1986], [1987].

1.4 Further Remarks

1. We should emphasize once more that our estimates are *nowhere based on price levels*. We neither exploit the relationship between the level of option prices and that of the prices of underlying securities (as in, e.g., Hutchinson, Lo and Poggio [1995]); nor do we infer a term structure of implied volatilities from option price levels, to be used to predict future prices (as in, e.g., Dumas, Fleming and Whaley [1996]). Instead, our estimates are generated entirely from the *correlation pattern* between options gains, on the one hand, and (i) futures gains, (ii) gains from overnight interbank investments, on the other hand.

2. In fact, gains from overnight interbank investments did not vary much in our sample. And, since they determine the LPE estimate only when overnight observations receive a positive weight, they act as a dummy variable. In terms of Equation (6), ΔB_t essentially takes only two values: zero (for observations over intraday intervals), or some positive number (for overnight observations). Therefore, our estimates of the option value will not only be based on the co-movement between options and futures price changes, but also on a comparison between intraday and overnight gains on options positions.

3. One should not exaggerate the influence of the gains on overnight interbank deposits, however. For one thing, they are an order of magnitude smaller than the typical overnight gain or loss on the underlying futures position. Also, there were relatively few overnight observations in our sample, reducing their potential impact even further.

4. Our estimates are based on local fits of one parameter only, namely, the volatility. We could have treated another variable in the Black model as parameter. For instance, we could have fitted the interest rate locally. We refrained from doing so, because our procedure is essentially a local version of nonlinear least squares. Convergence problems often emerge in nonlinear estimation procedures when applied to small samples. Since discrimination of the two parameters (volatility, interest rate) would have to be based on only a few overnight observations, we expected to encounter convergence problems, and, hence, we preferred to focus on the one parameter to which hedge ratios are most sensitive, namely, volatility.

2 The Data

2.1 Raw Data

The German Futures Exchange (*Deutsche Terminbörse*, DTB) sent us the entire record of futures and options transactions over the period 1992-94. This includes files with transactions in futures and options written on the German Stock Exchange Index (DAX). The options are *European-style*, and we considered only those that matured simultaneously with a futures contract. We used the futures contract to hedge the risk of the option, instead of the underlying index. This should reduce the actual transactions costs in the dynamic hedges. Also, it avoids our having to filter the DAX index series for dividend payments.⁷

Wolfgang Bühler, of the University of Mannheim, gave us a record of daily overnight money market (*Tagesgeld*) rates. These were used to compute the return on overnight positions in a bank account, which replaced the usual Treasury bill riskfree investment in the hedge portfolio. (Since option positions that are acquired and liquidated during the same trading day do not involve an investment outlay, the corresponding hedge portfolio need not take a riskfree position either; hence, we needed only the return on overnight riskfree positions.)

Both the DAX index futures and options markets are deep. Trade is totally computerized, and, hence, the records are unusually clean. The time stamp is accurate up to 1/100 of a second.⁸ We do not have quotes, however; the dataset consists only of transactions (volume and price). Trade in both futures and options takes place between 9:30am and 4:00pm.⁹ *Open interest* in the futures contract fluctuates between 60,000 and 160,000 contracts (each

⁷There 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 1:21pm and 1:30pm the same day.

⁸Even with a time stamp of 1/100 of a second, there are ties in the futures files.

⁹There is also a pre-trading period, from 7:30am till 9:30am, and post-trading, from 4pm till 5:15. Our data only cover the official trading period between 9:30am and 4:00pm.

contract is worth DEM 100 per index point¹⁰); that for the options 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, this means that values are quoted in DEM per 1/10 contract. Exercise prices for the options are set at increments of 25 index points. Five exercise prices are introduced for each expiration month.

After exploratory analysis which revealed extremely high kurtosis as well as persistence in volatility, we decided to implement a time change, and compute changes in futures and options prices over intervals in transaction time. The transaction frequency of the futures contract with the shortest time till maturity was chosen as a basis for this time transformation. Computing changes in futures prices over intervals of 300 transactions virtually eliminated the excess kurtosis, and dramatically reduced the persistence in volatility. Moreover, periods that cover a night in calendar time are not discernible at this frequency.¹¹ On average, a period of 300 transactions corresponded to roughly a 1 1/2 hour period in trading time. The latter corresponds to calendar time, except that one considers the close-to-open interval to be of zero length.

2.2 The “Training” Sample

The “training” sample refers to the part of the dataset that was used to compute LPE estimates of the local volatility. We constructed it as follows.

Take an option with a particular expiration date and strike price. We record a transaction for this option at a certain point in time. This will be used to generate one observation in the training sample, say, observation i . Let C_i denote the call transaction price. We search the futures transactions record for the next available futures transaction price. Let F_i denote this price. Then we advance our (transactions) clock with 300 futures transactions. We record the new futures price, F_i^* , and turn back to the options market, where we search for the next available options transactions price, C_i^* . We then compute changes:

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

$$\Delta F_i = F_i^* - F_i.$$

To generate observation $i + 1$, we proceed in the options transactions dataset and find the next transaction, record the price (C_{i+1}), find the next futures transaction price F_{i+1} , etc. We then repeat the process for another option.

Figure 1 provides a schematic view of our sampling technique. Notice that we could also have picked the options transaction price A to construct observation $i + 1$ (the corresponding options transaction price at least 300 futures transactions after the next futures transaction would be A^*). That way, however, there would be substantial overlap between observation i and $i + 1$. We originally constructed our dataset in this way. But we were concerned that some of our findings were caused by the substantial time overlap in our observations. Hence, we re-ran our analysis on non-overlapping data. Beyond a small drop in precision, the results hardly changed.

Between 1/92 and 3/92, we only include options with expiration 3/92; from the expiration in 3/92 till 6/92, we only include options with expiration 6/92. This matches our choice of the underlying futures contract: between 1/92 and 3/92, we use the 3/92 contract; after that, we use the 6/92 contract; etc.

The training sample extends all the way to the point in time where we want to price or hedge an option. We made sure that the future options transaction (the one with price C_i^* in the above example) always occurred before this reference point. The sample runs back at least two calendar months (one month for the 3/92 contracts). To avoid unmanageably large training samples, we dropped observations from the training sample that occurred before the third month preceding the expiration month. This way, training sample sizes run from a low of about 1,000 observations to a high of around 5,000 observations.¹²

Payoffs on interbank deposits were computed as follows. For observations that did not straddle a close-to-open interval, we set the beginning investment equal to zero. Hence, $\Delta B = 0$. For observations that did straddle a

¹⁰The DAX index at the start of our series was about 1625. Hence, the size of a futures contract would then be DEM 162,500, and open interest would roughly be between DEM 10 billion and DEM 26 billion.

¹¹At higher frequencies, the presence of intervals which straddle a close-to-open period reintroduces excess kurtosis.

¹²When using overlapping observations, sample sizes easily reached 17,000 to 45,000.

close-to-open interval, we invest $B = \$1$ in a bank account, which pays $1 + r$ the next day (r denotes the overnight interest rate on a per-day basis). Hence, $\Delta B = 1 + r$.

2.3 The “Testing” Sample

The pricing and hedging analysis is based on a “testing” sample, constructed as follows. We pick an option, say, the call option with strike price (K) equal to 1700 (index) points and expiration 3/92. We estimated its value and hedged its risk over intervals of 300 transactions of the futures contract with the same expiration month. We started the first trading day of the month preceding the expiration month, and followed the pricing and hedging until expiration.

As mentioned before, the “training” sample consists of observations pre-dating the testing sample. In the above example, this would mean that only 1/92 observations were used to compute the LPE local volatility for the first observation in the testing sample. As we progress in the testing sample, we add observations to the training sample, still assuring that the end-of-period option transactions in the training sample always occurred prior to the beginning of the next observation in the testing sample. Later on, we drop old observations from the training sample, as explained in the previous subsection.

2.4 Pricing And Hedging Study

For each option in our dataset, we construct a testing sample, and priced and hedged the option over its duration.

For each observation in the testing sample, we ran LPE on the corresponding training sample, thus obtaining a local volatility estimate. We implemented the procedure detailed in the previous section, except that we *normalized net gains by the exercise price*, in order to alleviate obvious estimation problems caused by nonstationarity of futures and options prices.

The local volatility estimate depends on the value of the option’s factors (z , which consists of: richness, maturity and interest rate). From this estimate $\hat{\sigma}_z$, we computed the value of the hedge portfolio as:

$$x_B^B(z, \hat{\sigma}_z).$$

We compared the performance of this *pricing* technique against that of synchronous market prices, by studying their ability to predict the option value at expiration. Synchronous market prices were obtained as follows. We scanned the training sample for the observation with the most recent (end-of-period) options transaction price, provided the option had a moneyness within 0.5%. We implied the volatility from this price, using Black’s formula. Black’s formula was then re-evaluated using this implied volatility, and the maturity and richness of the option that was priced with LPE. The resulting estimate is referred to as the synchronous market price.

We also investigated the deviations between LPE prices and synchronous market prices. LPE may generate prices that predict expiration values better than the market, but, if they often deviate substantially from market prices, LPE’s superior pricing could only be exploited by holding positions until maturity.

For each testing sample, we ran a *hedging analysis*. We constructed simple, self-financed delta hedges, based either on LPE or on Black implied volatilities. These hedges are re-balanced every 300 futures transactions, as follows.

Let’s assume that the hedge starts somewhere within the trading day. $x_F^B(z_1, \sigma)$ futures contracts are acquired (the vector z_1 holds the moneyness and maturity of the option, as well as the interest rate; σ will be determined either by LPE or as Black implied volatility). No position in interbank deposits is taken. At the beginning of the next observation, 300 futures transactions later, the futures position is liquidated, and the gains (losses) are deposited in a bank account. A new futures position is acquired, namely, $x_F^B(z_2, \sigma)$. Let’s assume that this position is not liquidated until the next day. In that case, $x_F^B(z_2, \sigma)$ Deutsche mark are deposited in the bank account. Since the position is taken overnight, it will carry interest. Note that interest will also be earned (charged) on the gains (losses) from the liquidation of the futures position in the first interval of 300 futures transactions.

Observation 3 starts the next morning, and let’s assume that it is liquidated the same day. This means that a new futures position, of $x_F^B(z_3, \sigma)$ contracts, will be acquired. Gains or losses from the previous futures position are

deposited in the bank account, adding to the principal plus interest gained overnight. The bank account, however, will not accumulate any interest over observation 3, since it ends within the same trading day. We continue this exercise until expiration of the option.

σ is computed either as the LPE local volatility or as the Black implied volatility from the most recent option transaction with similar richness (cf. the description of the computation of synchronous market prices). In the first case, the hedge is referred to as *LPE hedge*. In the latter case, it will be called *Black hedge*.

Notice that our hedging strategy is a self-financed version of the standard delta hedge. Its properties are studied in depth in Bossaerts and Werker [1997]. In particular, under certain conditions, the cumulative hedge error will be normally distributed, despite (obvious) non-normality of the option's payoff.

We are interested in the hedge error, which is defined to be the difference between the value of the hedge portfolio at maturity and the expiration value of the option. We will compare the properties of the hedge errors when the hedge portfolio is based on local volatilities from LPE, against those from the use of (Black) implied volatilities.

2.5 Kernel And Bandwidth Choice

There are two choices to be made in the context of any local estimation procedure: the kernel and the bandwidth. Choice of the kernel is usually straightforward. We chose the Epanechnikov kernel, which puts more weight on observations i with z_i close to the target z . Formally, the weight $w(z, z_i)$ used for observation i equals:

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

where $\|z_i - z\|$ denotes the distance between z_i and z , computed using a Euclidean measure, and b denotes the bandwidth (see below). We did experiment with other kernels (in particular, the normal kernel), but the results appear not to be sensitive.

The interpretation of the bandwidth becomes straightforward with an Epanechnikov kernel: it provides a direct measure of the size of the sample that receives positive weight in the LPE. In our application, we measured the distance between z_i and z in a standard Euclidean way, after normalizing each component of z_i by its range. More specifically,

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

If the bandwidth is set equal to 0.15, for instance, all observations with

$$\|z_i - z\| \geq 0.15$$

will receive zero weight in the LPE.

If the components of z_i are independently, uniformly distributed, exactly 15% of the sample would receive a positive weight when the bandwidth is set equal to 0.15. Of course, the components of z_i are neither independent nor uniform, so the effective size of the estimation sample may differ. Generally, about 5% of the sample is used. In Section 5, we will discuss an explicit example of the effective estimation sample sizes.

A small bandwidth reduces the bias in the estimation of the local volatility, but increases its sampling variation. We initially put a low constraint on the bandwidth: if less than 10 observations were included in the estimation sample with a bandwidth of 0.15, the bandwidth was successively increased by 10% until at least 10 observations received positive weight. As we encountered some anomalous behavior in the local volatility, we decided to tighten the constraint: we imposed a lower bound of 50 observations on the estimation sample. This did not eliminate the anomaly.

We used a bandwidth of 0.15. Experiments with several other bandwidths revealed that this choice led to the lowest average hedge error on a sub-sample of contracts that expired in 3/92. In fact, within a range of bandwidths

of 0.05-0.30, there was little sensitivity in the hedging performance. Of course, the bandwidth choice may not have been optimal for other contracts, but we wanted to refrain from mining the data. Also, our search for the optimal bandwidth was rudimentary: we tried several values, and picked the optimal one. A fully-fledged analysis (e.g., cross-validation) would have been prohibitively costly in terms of computation time. Finally, it should be pointed out that 0.15 was also the optimal bandwidth in the analysis of Bossaerts and Hillion [1997]. In that paper, we did use cross-validation (modified to account for time overlap) in order to find the optimal bandwidth.

3 Performance Of The Pricing Procedure

We studied the performance of LPE-based prices for all options in the dataset. (The present version only reports results for the 1992-3 subsample.) Prices were computed starting the first of the month preceding the expiration month, and every 300 futures transactions later.

Figure 2 provides an idea of the nature of the results. It plots the evolution of LPE-based prices (solid line) and synchronous market prices (dotted line), for the 6/93 call option with strike 1675. (Ignore the dashed line momentarily.) Notice how close LPE tracks the market price. This is remarkable, because we estimate prices on the basis of the (local) co-movement between options, futures, and overnight interest rates. Of course, option pricing theory predicts this outcome.

Notice, however, that there is a major aberration around the 80th 300-transaction interval. We actually chose to plot the results for the 6/93-1675 contract, because it does show a kind of anomaly that we will need to understand in order to explain the pricing and hedging performance. LPE prices accurately for most contracts, but there are a sufficient number of aberrations that the overall pricing and hedging results are affected.

Around the 80th time interval, the LPE-based valuation surges dramatically, returning, just a few periods later, back to the level of the market price. We will argue later that sampling error is not the cause. (When running any local estimation procedure, the bandwidth may be chosen so small that estimation is sometimes based on only a few observations.)

Before we attempt to explain the aberration, let us investigate the overall pricing performance of LPE. This is best done by comparing the ability of LPE-based call prices to predict the final payoff, against that of market-based prices.

Table 1 summarizes the prediction performance. Overall, LPE does at least as well as the market. We do want to stress that this is remarkable, because we have only been using information on co-movement of asset prices. Again, the fervent believers of option pricing theory would not find this surprising. It does demonstrate that option pricing theory not only generates price levels that are roughly in line with market prices, but also that it explains local correlation between option prices and prices of underlying assets. Without this local correlation, theoretical price levels cannot be derived in the first place.

We can directly compare LPE-based prices and market prices. This is done in Table 2. We observe some major deviations between the two. Of course, this is also what one would infer from a casual look at Figure 2. It appears that LPE often disagrees with the market on the correct price level. This does not affect the overall pricing performance, in the sense that both predict the final payoff equally well (see Table 1).

Which price is right: the one based on LPE, or the market's? We will be able to answer this question from a study of the hedging performance of delta hedges based on LPE or on (the volatility implied from) market prices. Hedging performance is studied in the next section.

At this point, we should look at a plot of the pricing deviations against the option's richness. See Figure 3.

Most (large) deviations occur for near-the-money options. This is not surprising. Synchronous market quotes are obtained from Black-implied volatilities. For way out-of-the money or in-the-money options, the Black price is insensitive to the volatility, and, hence, the volatility is implied to be the initial guess, 0.15. An analogous phenomenon occurs for local volatilities. Ignoring the overnight position, LPE finds the estimate of the volatility

at z by minimizing:

$$\sum_{i=1}^N w(z, z_i) (\Delta C_i - x_F^B(z_i; \sigma) \Delta F_i)^2$$

($w(z, z_i)$ is the weight given to observation i ; it is a function of the distance between z_i and z ; see above). The first-order condition is:

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

In the Black model,

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

where τ_i denotes the maturity of the option, r_i is the interest rate,

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

and F_i and K_i denote the (beginning-of-period) futures price and strike price of observation i , respectively. $n(\cdot)$ is the normal density function. The latter is very close to zero for way in-the money (high F_i/K_i) or way out-of-the-money (low F_i/K_i) options. Hence, the first-order condition would automatically be satisfied in those cases, and the optimization algorithm stops at the initial guess, 0.15. (An analogous phenomenon will occur when maturity decreases to zero.)

Because synchronous market prices and LPE-based prices are computed from the same initial guess (0.15), no discrepancies will be observed. This explains why price deviations decrease to zero away from a richness equal to 1.

Near-the-money, however, some of the pricing differences are substantial. An aberration was discovered in Figure 2. Figure 3 demonstrates that this aberration is not a rare event.

Figure 4 plots local volatilities against richness. It is the mirror image of Figure 3. Volatilities for moneyness away from 1 converge to 0.15. Close to the money, we see a serious anomaly: local volatilities are estimated to be as low as 0.02 (the lower bound in our estimation algorithm) or as high as 2.00 (the higher bound). Since maturity is expressed as a fraction of a year in the estimation, a volatility of 2.00 corresponds to a return volatility of 200% per annum!

It is not uncommon to estimate extremely high or low *implied* volatilities when these are based on single transactions. However, it is much more surprising that we obtain these extremes on the basis of local estimation of the co-movement between options and futures prices, certainly because the sample had to contain at least fifty observations. We will investigate the cause of this anomaly in Section 5. In the meanwhile, let us investigate the performance of LPE-based delta hedges.

4 Performance Of LPE-Based Delta Hedges

We explained earlier how we constructed the (self-financed) delta hedges. The dashed line in Figure 2 depicts the evolution of the value of (dollar investment in) our LPE-based portfolio, set up to hedge the risk of the 6/93-1675 DAX call option.

Notice that the first few periods, the value is close to zero. Initially, the hedge requires only futures trades, because the position is not carried overnight. The portfolio acquires nonzero value only because of gains/losses on the futures position. The value of the hedge portfolio changes dramatically, however, upon the first overnight position. Only at that point does one take the net (initial) investment in the hedge. Before and after that moment, no further cash inflows or outflows are necessary; the portfolio is self-financing. The large increase in the value of the hedge portfolio in Figure 2 coincides with the point at which the initial investment in the delta hedge is taken.

The hedge in Figure 2 does not perform well: while changes in the value of the hedge portfolio follow the pattern of the market price of the option, the level of the value is much too high. The eventual hedge error (difference between value of hedge portfolio and value of the option itself at expiration) is sizeable. This is in sharp contrast with LPE-based prices, which track the market price fairly well (see the solid line in Figure 2). Of course, hedging demands more than just accurate pricing. The latter only determines an appropriate *size* of the hedge portfolio (number of dollars to be invested). For the hedge to be successful, the portfolio's *composition* (number of futures contracts acquired) must be such that changes in its value reliably track changes in the option price.

The success of LPE as far as pricing is concerned does not seem to translate into accurate (delta) hedging. Our discussion indicates that this must be attributed to LPE's inability to form a portfolio that tracks changes in the option price. This is certainly surprising, because LPE is based on a statistical analysis of the co-movement between the option price and the prices of underlying securities!

The situation depicted in Figure 2 may not be representative. So, let us look at the hedging performance across all contracts in our dataset (156 in the 1992/3 part). We will compare the performance of LPE-based hedging with that of standard delta hedges that are computed from local volatilities.

Table 3 compares descriptive statistics of the values at maturity of the LPE hedge portfolios against those for delta hedges based on Black's model. The differences are minimal. 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 synchronous values of the option. Table 4 provides descriptive statistics, for hedge errors as of 180 and 90 transaction time units before maturity, as well as at the expiration date.

Overall, the performance of the LPE-based delta hedge is far worse than that of the standard hedge based on Black's delta and implied volatilities. Hedge errors can be as large as 26% of the strike price. The standard delta hedge leads to hedge errors that are at most 5% of the strike price.

Figure 5 plots the LPE hedge errors at maturity against the moneyness of the option. The biggest errors occur for moneyness values around 1. Of course, that is precisely where option prices are most nonlinear (have the highest second derivative with respect to the underlying security's price). It is difficult for delta hedges to capture this nonlinearity. Still, when implied volatilities are used, the eventual error is far less than with LPE. This suggests that the nonlinearity hardly contributed to the poor performance of LPE-based delta hedging.

The large hedging errors of LPE must be caused by the anomalous behavior of the local volatility. The often excessively high or low local volatility (see Figure 4) implies futures positions that are inadequate to hedge the risk of the DAX index options. Most of these excesses occur for near-the-money options, and, consequently, most of the poor hedging performance of LPE is concentrated in the same region. See Figure 5.

We are left to explain the anomalous behavior of the local volatility. Let us do so now.

5 Explaining The Anomalous Behavior Of The Local Volatility

Ignoring the return on interbank deposits, the local volatility is estimated by a weighted nonlinear least squares fit of the following equation:

$$\Delta C_i = x_F^B(z_i; \sigma) \Delta F_i + \eta_i$$

($i = 1, \dots, N$). The weights are determined by the distance between z_i and z . z contains the values of the richness and maturity of the option to be priced or hedged, as well as the level of the interest rate.

$x_F^B(z_i; \sigma)$ is the hedge ratio in Black's model. As a function of the parameter, σ , this ratio is severely constrained. Figure 6 illustrates this. It plots $x_F^B(z_i; \sigma)$ as a function of richness (one of the elements of z_i), for a maturity equal to 0.05 year. σ is varied from a low of 0.02 to a high of 2.00 (these were also the bounds we imposed in the numerical nonlinear optimization within LPE).

Several observations can be made about Figure 6. First, the hedge ratio is always constrained to be between 0 and 1. The actual co-movement between options and futures prices may not agree. If the local co-movement is negative for out-of-the-money options, LPE will generate a corner solution and set the local volatility equal to 0.02. LPE will do the same for in-the-money options when the local co-movement is higher than 1.

There are a multitude of instances in Figure 4 when the local volatility is indeed estimated to be 0.02. In all cases, the richness is above one, indicating that the local co-movement between options and futures prices is far higher than the maximum predicted by the theory (1).

Second, there is an upper bound to the co-movement for out-of-the-money options. It is far below 1. It is clear from Figure 6 that this upper bound obtains for a volatility of 2.00. The corresponding hedge ratio is at most 0.6 (for a maturity of 0.05 year). If the actual co-movement is higher, LPE will generate an estimate of 2.00. We don't observe these estimates in Figure 4, but we did in the wider sample. (Figure 4 only displays estimates of local volatilities for a range of maturities between 31 and 60 transaction time periods till maturity.)

Third, there is a correspondingly tight lower bound on the co-movement for in-the-money options. As the moneyness increases, the volatility value for which this lower bound obtains increases as well. Hence, if the actual co-movement is below the lower bound, we expect excessively high local volatilities, not all necessarily equal to 2.00. This is borne out in Figure 4.

All this is related to a weakness of local parametric estimation: if the range of outcomes is restricted by the parametric model that one attempts to fit locally, and the data disagree, one cannot expect local parametric estimation to produce consistent estimates. This contrasts with the more traditional local nonparametric estimation technique, whereby (lower-order) polynomials are fitted locally: to obtain consistency, one merely needs smoothness conditions.

Does this mean that we should abandon local parametric estimation in favor of local polynomial estimation? Certainly not. The empirical results do point to a serious discrepancy between the data and option pricing theory which cannot be resolved by simply switching to local polynomial estimation. In particular, the co-movement between options and futures prices often contradicts theoretical restrictions. This is most blatant for cases where LPE generates a solution at the lower corner ($\sigma = 0.02$), suggesting that the actual hedge ratio should be below 0 (for out-of-the-money options) or above 1 (for in-the-money options).

The poor performance of LPE-based delta hedges reveals that the periods of anomalous co-movement are short-lived. If they were not, LPE-based hedge ratios would cover the risk of option price changes far better than they did.

The spike around period 80 in Figure 2 illustrates how short-lived anomalies are: local volatility increased quickly from 0.10 ($t = 71, \dots, 76$) to 0.15 ($t = 77$), 0.23 ($t = 78$) and 0.22 ($t = 79$), then rapidly dropped to 0.14 ($t = 80, 81$) and back to 0.10 ($t = 82, 83$).

Of course, there is the danger that the aberrations in the behavior of local volatilities are caused by sampling error. When the bandwidth is fixed in a local estimation procedure, one is never sure that enough observations receive positive weight to mitigate sampling variation. In fact, with our bandwidth of 0.15, the effective sample size was sometimes as low as one element (theoretically, it could even be zero). As mentioned before, however, we incrementally increased the bandwidth with 10% until the effective sample included at least 50 observations.

Figure 7 illustrates this variable bandwidth technique for the 6/93-1675 contract. The solid line depicts 5% of the total training sample size (it increases as one moves in time); the dotted line indicates the evolution of the effective sample size.

To get a better idea of the nature of the aberrating LPE options prices, we should investigate the part of the training sample that they are based on. Let us look at the worst case in Figure 2: the options price in period 78. The effective subsample included 54 observations, and they are plotted in Figure 8. Nothing seems to indicate that there is anything unusual. As option pricing theory predicts, the relationship between options and futures price changes is almost linear (due to the short observation interval). Still, with a moneyness of 0.98, the slope of the weighted least squares line through the datapoints, 0.34, is too high. This slope corresponds to a local volatility of 0.23. That is more than twice the local volatility implied by the estimated co-movement of options and futures prices around the same period.

In summary: the poor hedging performance is caused by aberrations in the local behavior of the co-movement between options and futures prices. Obviously, this also produced the often huge discrepancies between LPE-based option prices and synchronous market prices. Apparently, the co-movement between options and futures prices is at odds with the theory of option pricing.

6 Conclusion

We have presented a set of results on the pricing and hedging performance of a technique that relies entirely on the local correlation pattern between prices of the option and those of underlying securities. The local correlation is captured by means of a parametric option pricing model.

In theory, this technique should work well. Correlation between values of derivatives and underlying securities is at the heart of arbitrage-based option pricing models. We simply propose to take this idea to be the core of empirical analysis as well. This contrasts with the traditional approach, which extracts information from a pure time series analysis of the prices of underlying assets (or uses implied volatilities).

We applied the procedure, called local parametric analysis (LPE) to a dataset of DAX index options. On average, LPE-based prices predict final option payoffs as well as market prices. Yet, LPE-based prices often deviate substantially from market prices. Moreover, self-financed delta hedges based on LPE estimates of the hedge ratio underperform those based on implied volatilities.

At the root of the poor hedging performance (and the deviations between LPE and market prices) is frequent anomalous co-movement between DAX options and DAX futures prices. For instance, parametric option pricing models (like the Black model we have been employing here) restrict the coefficient from a local projection of options price changes onto futures price changes to be between zero and one. Yet, we find many violations of this constraint.

To close, we cannot put enough emphasis on the striking contrast between the result of this study and Bossaerts and Hillion [1997]. The latter investigated the performance of LPE-based hedging in an environment where a continuous-time option pricing model (Black-Scholes) holds *exactly*, but where hedging is constrained to take place in discrete time. LPE was found to significantly outperform Black-Scholes delta hedging. Contrast this with the *real-world* application of the present paper: LPE is far worse as far as delta hedging is concerned. This points to a major discrepancy between the theoretical and actual pattern of co-movement of the prices of derivatives and those of underlying securities.

References

- Black, F., 1976, "The Pricing of Commodity Contracts," *Journal of Financial Economics* 3, 167-179.
- Bossaerts, P., E. Ghysels and C. Gouriéroux, 1996, "Arbitrage-Based Pricing When Volatility Is Stochastic," Caltech working paper 977.
- Bossaerts, P., C. Hafner and W. Härdle, 1996, "Foreign Exchange Rates Have Surprising Volatility," in: P.M. Robinson and M. Rosenblatt, eds., *Athens Conference on Applied Probability and Time Series*, Vol II (Time Series Analysis). New York: Springer Verlag.
- Bossaerts, P. and P. Hillion, 1997, "Local Parametric Analysis of Hedging in Discrete Time," *Journal of Econometrics* 81, 243-272.
- Bossaerts, P. and B. Werker, 1997, "Martingale-Based Hedge Error Control," in: C. Rogers and D. Talay, eds., *Numerical Methods In Finance*. Cambridge (U.K.): Cambridge University Press.
- Duffie, D. and H.R. Richardson, 1991, "Mean-Variance Hedging in Continuous Time," *Annals of Applied Probability* 1, 1-15.
- Dumas, B., J. Fleming and R. E. Whaley, 1996, "Implied Volatility Smiles: Empirical Tests," Duke University working paper.
- Föllmer, H. and D. Sondermann, 1986, "Hedging of Contingent Claims Under Incomplete Information," in: A. Mas-Colell and W. Hildenbrand, *Contributions to Mathematical Economics*. Amsterdam: North-Holland.
- Gouriéroux, C., J.P. Laurent and H. Pham, 1995, "Quadratic Hedging and Numéraire," CREST working paper.

- Härdle, W., 1990, *Applied Nonparametric Estimation*. Cambridge (U.K.): Cambridge University Press.
- Harris, L., 1986, "Cross-Security Tests of the Mixture of Distribution Hypothesis," *Journal of Financial and Quantitative 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. and I.K. Glad (1995): "Nonparametric Density-Estimation With A Parametric Start." *Annals of Statistics* 23 (3, June), 882-904.
- Hull, J. and A. White, 1987, "The Pricing of Options on Assets with Stochastic Volatilities," *Journal of Finance* 3, 281-300.
- Hutchinson, J., A. Lo and T. Poggio, 1994, "A Nonparametric Approach to Pricing and Hedging Derivative Securities Via learning Networks," *Journal of Finance* 49, 851-890.
- Ross, S., 1978, "Mutual Fund Separation and Financial Theory – The Separating Distributions," *Journal of Economic Theory* 17, 254-86.
- Schweizer, M., 1995, "Variance Optimal Hedging in Discrete Time," *Mathematics of Operations Research* 20, 1-32.
- Tauchen, G. and M. Pitts, 1983, "The Price Variability-Volume Relationship on Speculative Markets," *Econometrica* 51, 485-505.

Table 1
LPE Prediction Errors Versus Market Prediction Errors, For Different Maturities

	Maturity					
	60		130		210	
	LPE	Market	LPE	Market	LPE	Market
<i>N</i>	155	155	105	105	66	66
Average	0.0048**	0.0051**	-0.0027	-0.0033	0.0012	-0.0057
St. Dev.	0.0225	0.0213	0.0395	0.0342	0.0496	0.0309
Skewness	0.27	-0.26	0.74	-0.29	1.93	-0.52
Kurtosis	5.6	4.0	5.6	2.3	10.7	1.9
Minimum	-0.0518	-0.0518	-0.0676	-0.0676	-0.0607	-0.0607
5%	-0.0453	-0.0429	-0.0623	-0.0623	-0.0576	-0.0576
95%	0.0395	0.0382	0.0468	0.0465	0.0976	0.0354
Maximum	0.1019	0.0570	0.1686	0.0579	0.2424	0.0374

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 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 1/2 hour calendar trading time, on average; this corresponds to slightly over 1/4 trading day). ** indicates: significant at the 1% level (two-tailed *t*-test).

Table 2
Deviations Between LPE-Based Prices And Market Prices

	<u>Maturity</u>						
	60	90	120	150	180	210	240
<i>N</i>	155	155	138	105	85	66	48
Average	-0.0003	0.0086*	0.0011	0.0024	-0.0002	0.0068	-0.0007
St. Dev.	0.0106	0.0403	0.0195	0.0230	0.0145	0.0334	0.0068
Skewness	10.67	4.29	10.38	6.40	0.70	4.97	0.67
Kurtosis	129.8	21.1	119.2	48.8	10.3	31.9	6.4
Minimum	-0.0173	-0.0165	-0.0202	-0.0257	-0.0505	-0.0221	-0.0174
5%	-0.0066	-0.0069	-0.0081	-0.0092	-0.0137	-0.0060	-0.0137
95%	0.0027	0.0892	0.0087	0.0090	0.0254	0.0735	0.0178
Maximum	0.1251	0.2298	0.2195	0.1882	0.0582	0.2249	0.0191

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 1/2 hour calendar trading time, on average; this corresponds to slightly over 1/4 trading day). * indicates: significant at the 5% level (two-tailed *t*-test).

Table 3
Values Of Hedge Portfolios At Maturity

	<u>Basis For Delta</u>	
	LPE	Black Model
<i>N</i>	156	156
Average	0.0909	0.0841
St. Dev.	0.0988	0.0956
Skewness	0.99	1.11
Kurtosis	3.2	3.4
Minimum	-0.0099	-0.0114
5%	-0.0010	-0.0009
95%	0.2883	0.2869
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 values are expressed as a percentage of the strike price (premia).

Table 4
Hedge Errors

	<u>Maturity</u>					
	180		90		0	
	<u>Basis For Delta</u>		<u>Basis For Delta</u>		<u>Basis For Delta</u>	
	LPE	Black Model	LPE	Black Model	LPE	Black Model
<i>N</i>	85	85	155	155	156	156
Average	0.0072**	0.0023**	0.0098**	0.0025**	0.0123**	0.0054**
St. Dev.	0.0202	0.0065	0.0292	0.0090	0.0284	0.0096
Skewness	2.66	0.37	4.18	1.80	4.39	1.61
Kurtosis	13.4	6.9	26.7	7.3	30.0	9.7
Minimum	-0.0349	-0.0209	-0.0218	-0.0146	-0.0198	-0.0213
5%	-0.0138	-0.0056	-0.0094	-0.0121	-0.0084	-0.0039
95%	0.0461	0.0120	0.0660	0.0172	0.0654	0.0228
Maximum	0.1064	0.0283	0.2308	0.0503	0.2373	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 prices 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 1/2 hour calendar trading time, on average; this corresponds to slightly over 1/4 trading day). ** indicates: significant at the 1% level (two-tailed *t*-test).

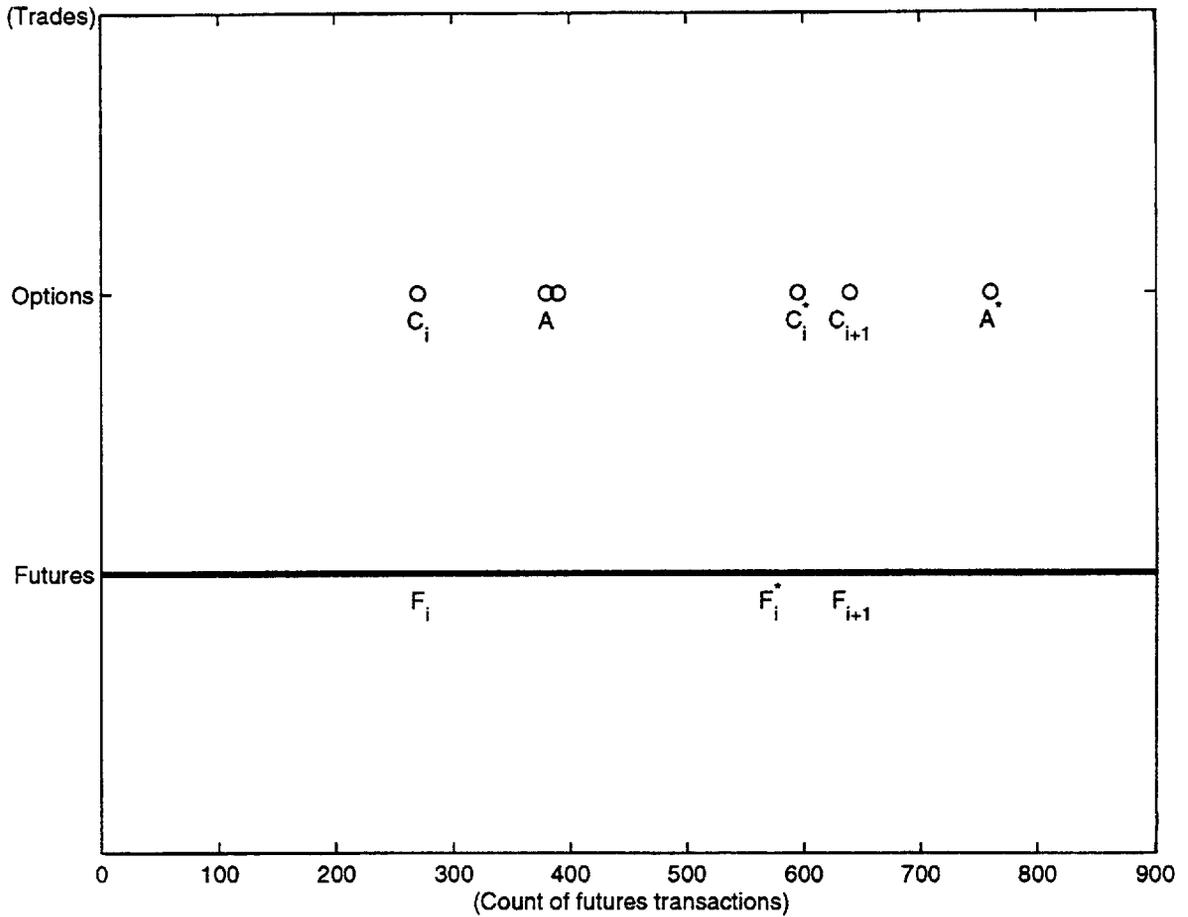


Figure 1: Schematic representation of the construction of the training sample. The i th observation consists of: an options transactions price, C_i ; the subsequent futures quote, F_i ; the futures quote 300 futures transactions later, F_i^* ; and the subsequent options transactions price, C_i^* . The $(i + 1)$ 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: $\Delta C_i = C_i^* - C_i$; $\Delta F_i = F_i^* - F_i$.

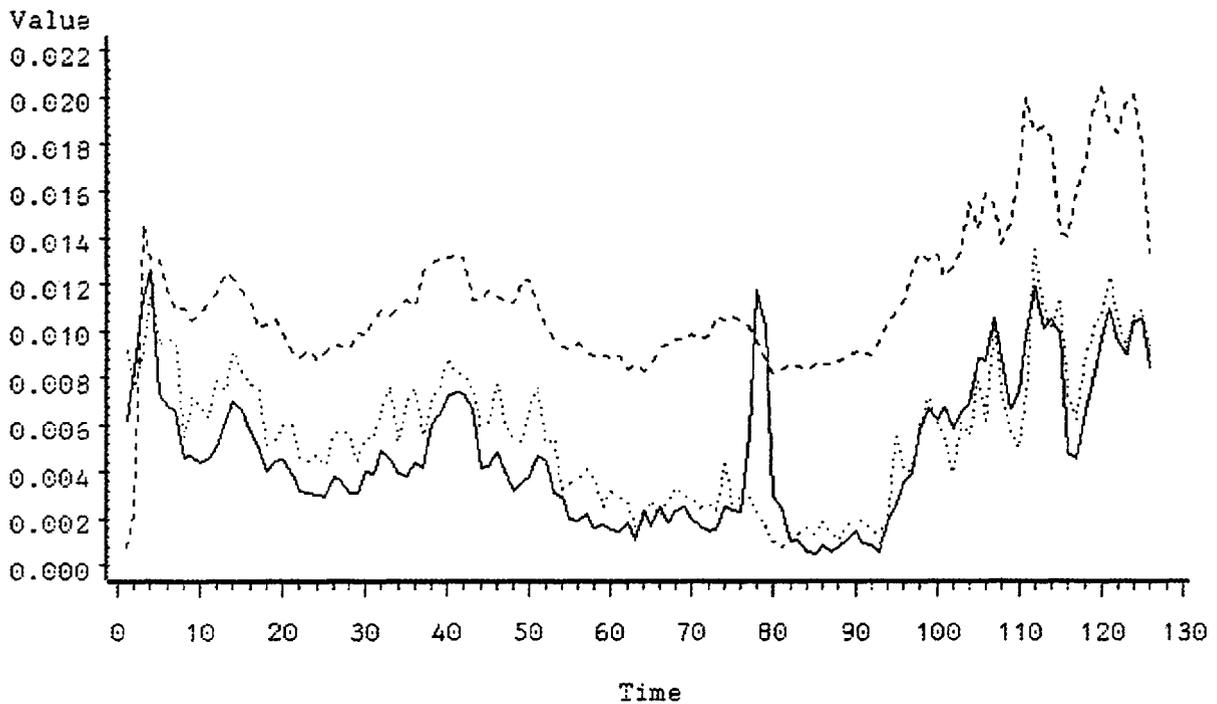


Figure 2: Time series plots of values of the DAX call option with strike 1675 and maturity 6/93. The solid line depicts the price of the option, estimated using LPE. The dotted line depicts synchronous market prices. The latter are obtained by implying the volatility from the most recent transaction 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 1/2 hour calendar trading time, on average; this corresponds to slightly over 1/4 trading day).

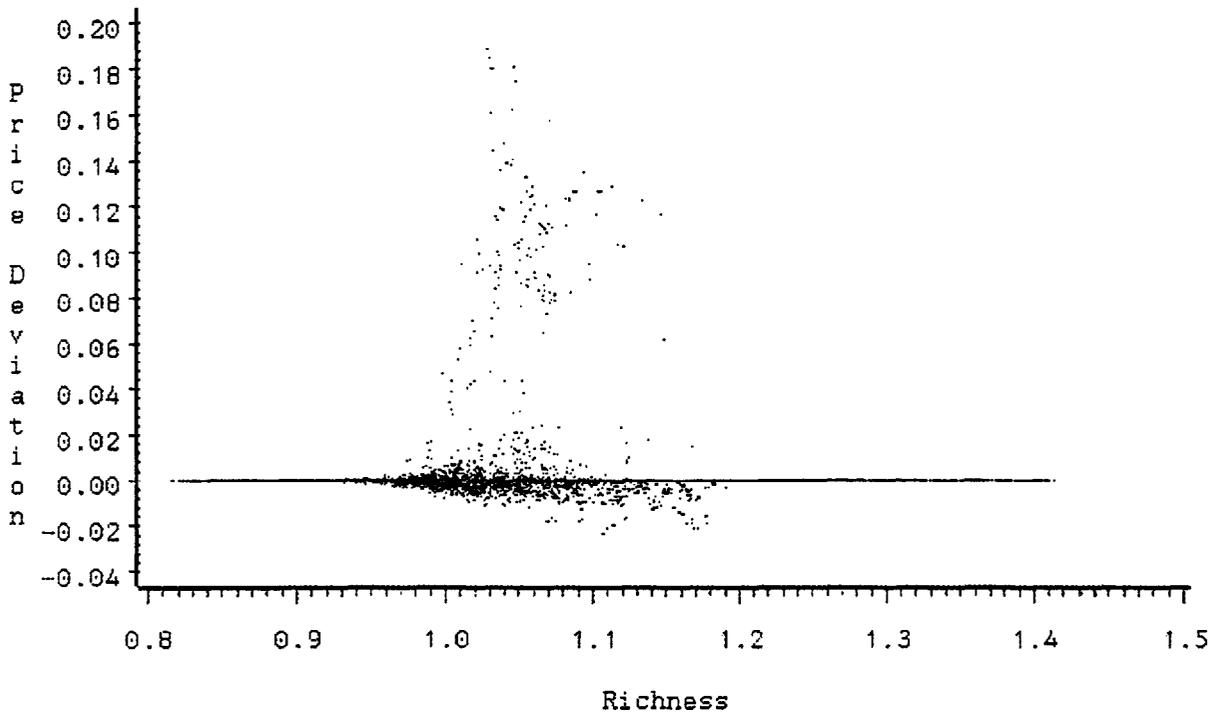


Figure 3: Plot of deviations of LPE-based prices from synchronous market prices, against against richness (underlying futures quote divided by strike price). The synchronous market price is obtained by implying the volatility from the most recent transaction price of an option with similar richness and 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 31 and 60 are included. Maturity is measured in number of intervals of 300 futures transactions (each interval corresponds to 1 1/2 hour calendar trading time, on average; this corresponds to slightly over 1/4 trading day).

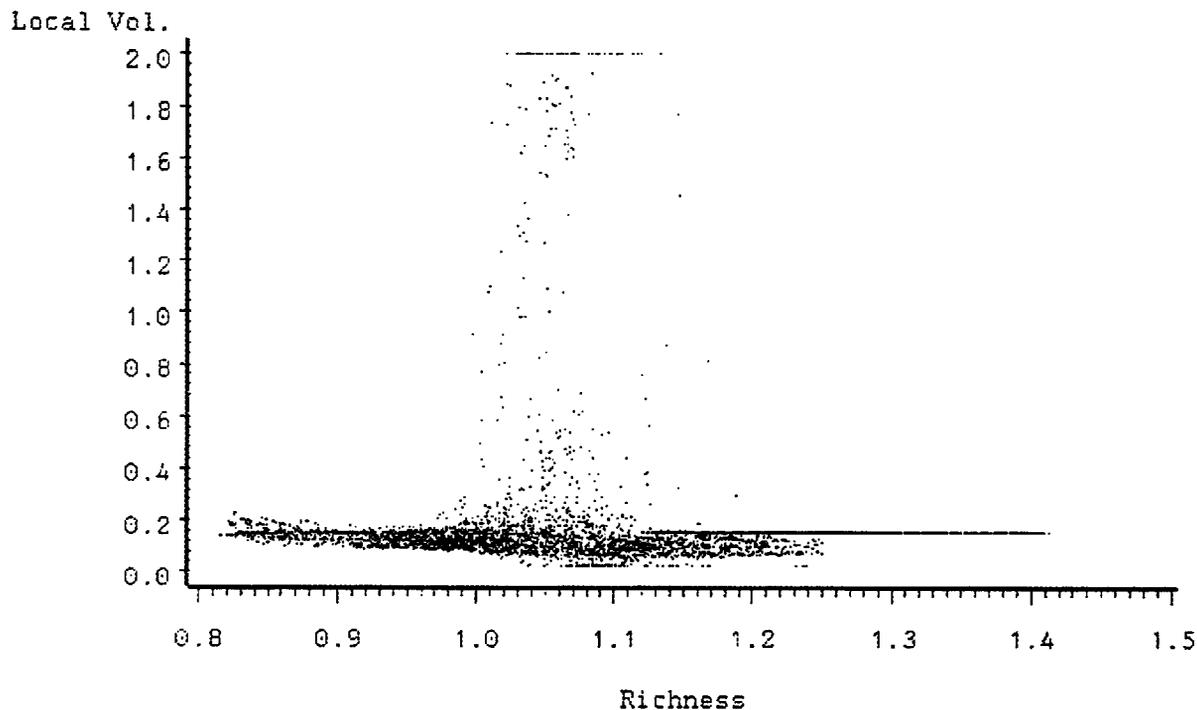


Figure 4: Plot of local volatility against richness (underlying futures quote divided by strike price). The local volatility is defined to be the LPE parameter estimate of the volatility in the Black model. The Black model is used to approximate locally the covariance between DAX option prices, on the one hand, and DAX futures quotes and interest on overnight bank deposits, on the other hand. Only observations with a maturity between 31 and 60 are included. Maturity is measured in number of intervals of 300 futures transactions (each interval corresponds to 1 1/2 hour calendar trading time, on average; this corresponds to slightly over 1/4 trading day).

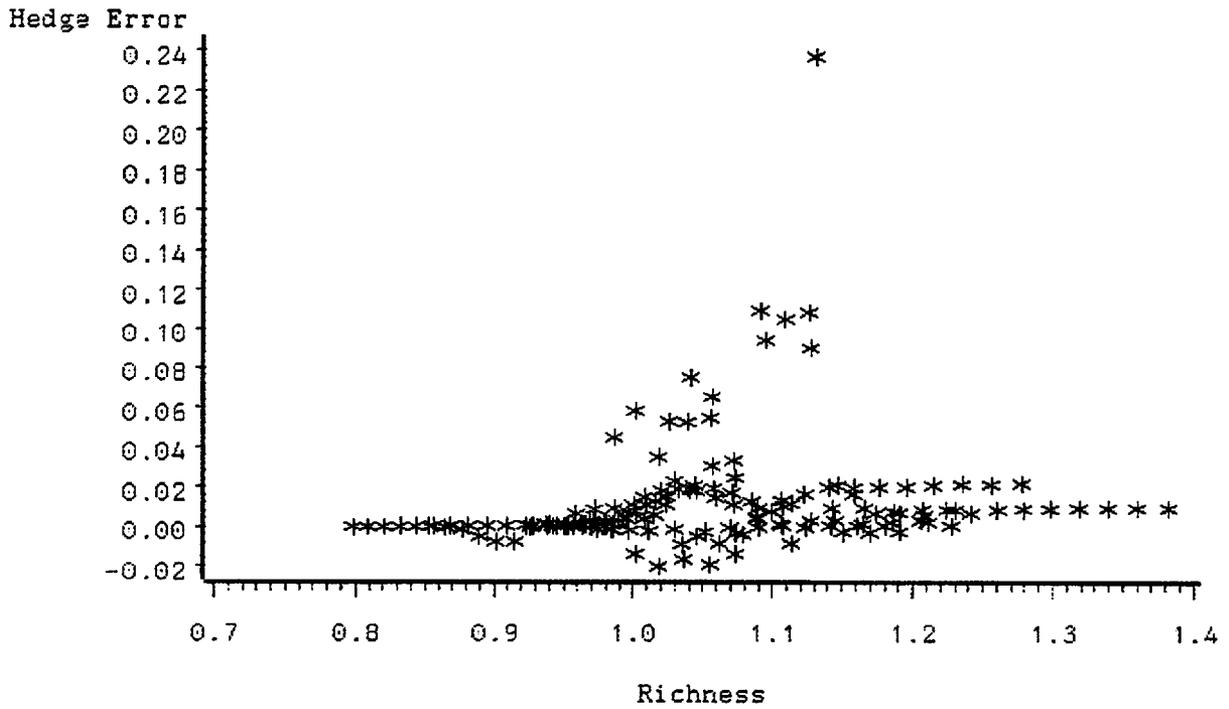


Figure 5: Plot of hedge error (at maturity) against richness (underlying futures quote divided by strike price). The hedge error (at maturity) is defined to be the difference between the value of the hedge portfolio (at maturity) based on LPE and the expiration value of the option. All prices and values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

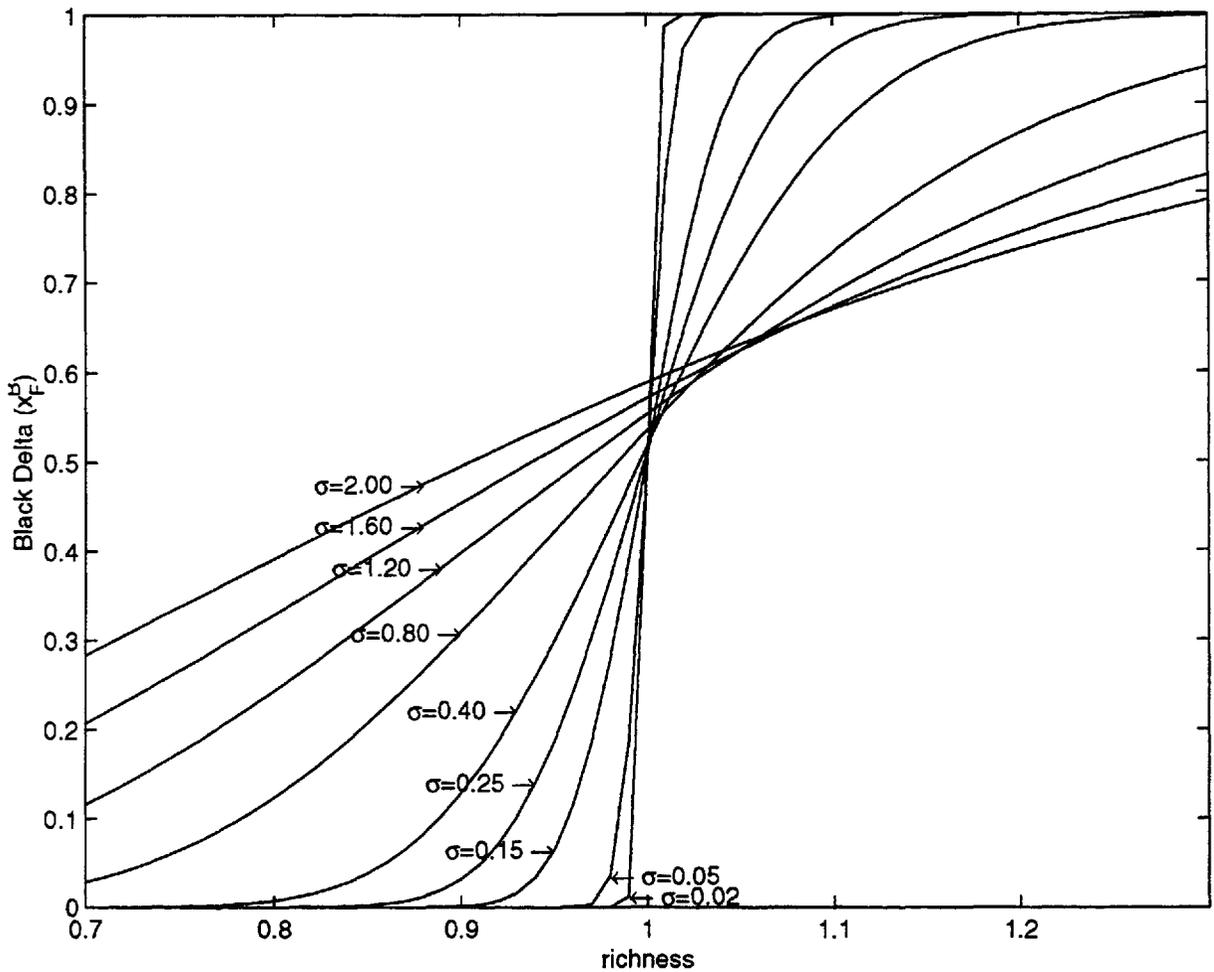


Figure 6: Plot of the Black (futures) hedge ratio as a function of richness (underlying futures quote divided by strike price), for different volatilities. Volatilities are annualized. The maturity is set equal to 0.05 year, which is roughly equal to the maturities of the options for which local volatilities are plotted in Figure 4.

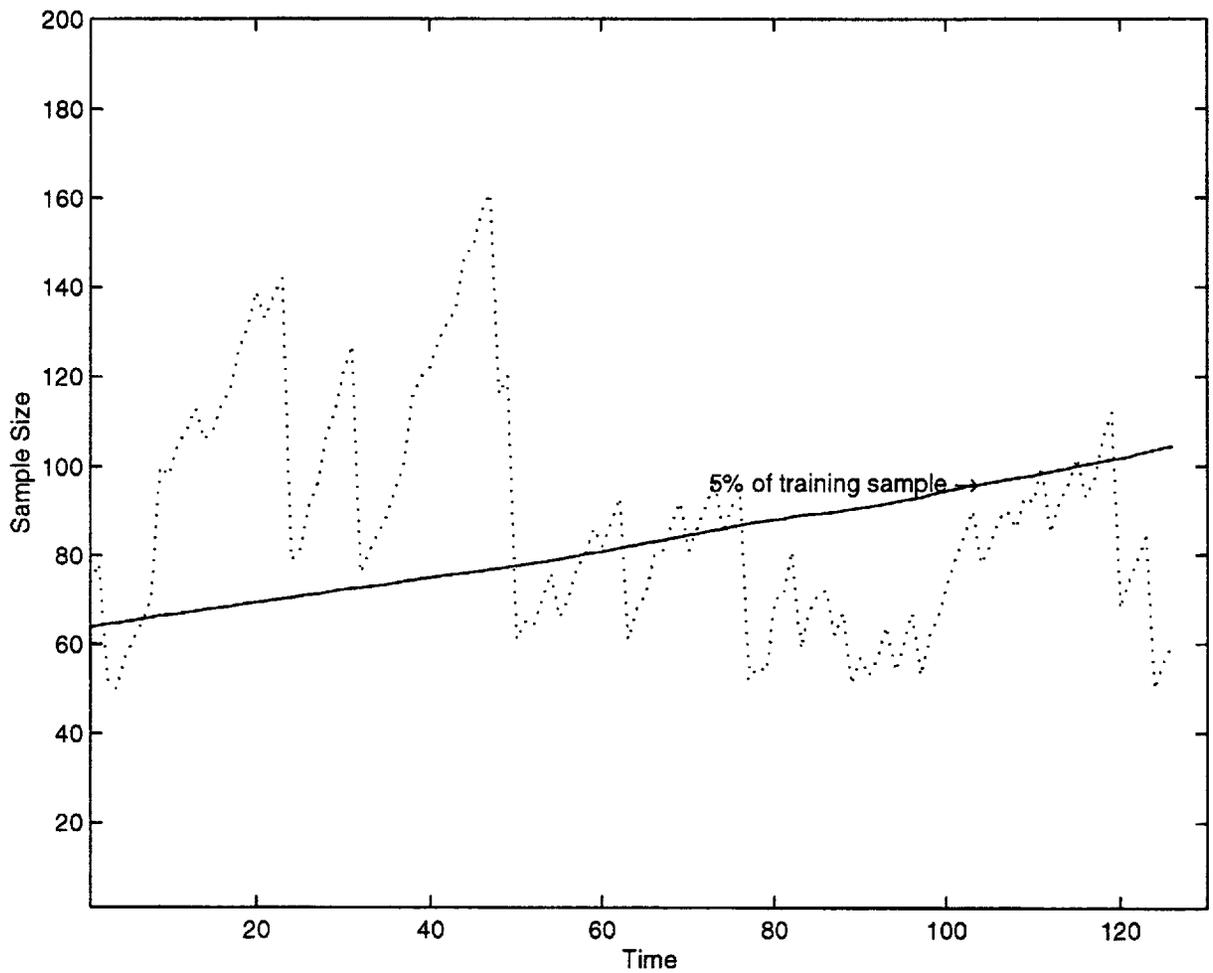


Figure 7: Time series of the number of observations in the training sample that receive positive weight in LPE. Results for the 6/93 call option contract with strike 1675 are shown. The solid line indicates 5% of the total (training) sample size.

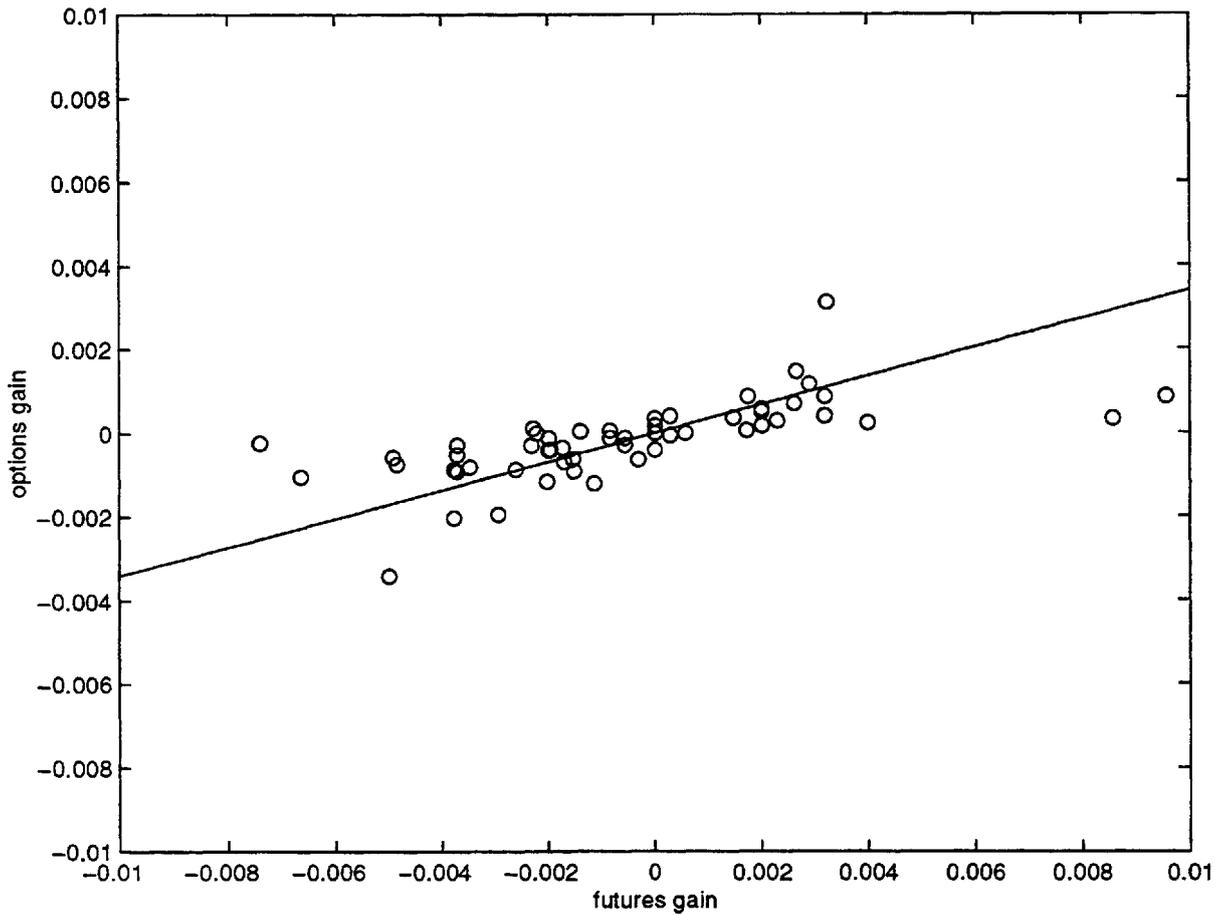


Figure 8: Plot of options gains (changes in options prices) against futures gains (changes in futures prices) for the subsample of the training sample that is used to compute the LPE option price of period 78 depicted in Figure 2. Gains are normalized by the strike price. The solid line is the (weighted) least squares fit.