Tilburg University

Local Parametric Analysis of Hedging in Discrete Time

Bossaerts, P.L.M.; Hillion, P.

Publication date:
1995

[Link to publication in Tilburg University Research Portal](#)

Citation for published version (APA):

Bossaerts, P. L. M., & Hillion, P. (1995). *Local Parametric Analysis of Hedging in Discrete Time*. (CentER Discussion Paper; Vol. 1995-23). CentER.

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Local Parametric Analysis of Hedging in Discrete Time

Peter Bossaerts

Pierre Hillion

October 1994 (first draft); August 1995 (this draft)*

*First author's affiliation: California Institute of Technology and CentER, Tilburg University; Second author's affiliation: INSEAD.

Abstract

When continuous-time portfolio weights are applied to a discrete-time hedging problem, errors are likely to occur. This paper evaluates the overall importance of the discretization-induced tracking error. It does so by comparing the performance of Black-Scholes hedge ratios against those obtained from a novel estimation procedure, namely *local parametric estimation*. In the latter, the weights of the duplicating portfolio are estimated by fitting parametric models (in this paper, Black-Scholes) in the neighborhood of the derivative's moneyness and maturity. Local parametric estimation directly incorporates the error from hedging in discrete time. Results are shown where the root mean square tracking error is reduced up to 41% for short-maturity options. The performance can still be improved by combining locally estimated hedge portfolio weights with standard analysis based on historically estimated parameters. The root mean square tracking error is thereby reduced by about 18% for long-maturity options. Plots of the locally estimated volatility parameter against moneyness and maturity reveal the biases of the Black-Scholes model when hedging in discrete time. In particular, there is a sharp “smile” effect in the relation between estimated volatility and moneyness for short-maturity options, as well as a significant “wave” effect in the relation with maturity for deep out-of-the-money options.

Send Proofs To: Peter Bossaerts, HSS 228-77, California Institute of Technology, Pasadena, CA 91125.

Acknowledgement: The authors are grateful for comments from participants at seminars at CentER, Erasmus University, Universitat Pompeu Fabra and Wharton, at the Conference on Financial Modelling and Econometric Analysis (Paris, December 1994), the Conference on Nonparametric and Dynamic Modelling (Berlin, December 1994), the Programme on Financial Mathematics at the Isaac Newton Institute (Numerical Methods Emphasis, April 1995), and the Econometric Society World Congress in Tokyo (August 1995), as well as for the suggestions of the Editor (H. Lütkepohl) and two anonymous referees.

1 Introduction

Since the seminal paper of Black and Scholes [1973], derivatives analysis has been following a standard pattern. First, continuous-time or binomial processes are posited for the underlying securities. Second, the hedging strategy is determined, leading to a dynamic portfolio that replicates the payoff on the derivative perfectly. Third, the pricing restrictions imposed by absence of arbitrage are derived. These take the form of a partial differential equation (in the continuous-time case) or a difference equation (in the binomial case), most often to be solved numerically. Fourth, the parameters of the posited processes are estimated, with the aim of evaluating prices offered in the market and of hedging open positions.

There is nothing logically wrong with the traditional approaches. Nevertheless, they do have disadvantages. Focusing on the continuous-time modeling, one can immediately mention the impossibility of hedging in continuous time. This causes misspecification. The ultimate goal of the present paper is to better understand the nature of the error from blindly applying continuous-time modeling to a discrete-time setting. In addition, there is the difficulty of estimating the parameters of the continuous-time value processes. Even if the relationship between the available time series and underlying processes is straightforward, actual estimation has revealed plenty of problems, such as near-unit-root behavior of continuous-time interest rates (e.g., Gibbons and Ramaswamy [1993]), of stochastic volatility (e.g., Bossaerts and Hillion [1993]) and lack of precision in the estimation of the mean return (which is not irrelevant when the data come in discrete time; see Lo and Wang [1995]).

Let us start with the latter. One of the striking constants that emerges from a re-reading of empirical work in finance is the ease with which *correlation* across securities can be estimated. Not only are the estimates precise, they are reasonably stable over time. For common stock returns, this can best be exemplified by the almost universal practice of using five-year windows to estimate “betas” (already present in the early work of Black, Jensen and Scholes [1972] and Fama and MacBeth [1973]).

In fact, derivatives analysis is a prime example of an exercise that is founded on correlation: it

exploits (usually time-varying) correlation between securities in order to construct a perfect hedge and derive its pricing implications. Contrast this with the standard implementation of derivatives analysis, whereby univariate time series properties, mainly volatility, are focused on. Variables such as volatility are not of immediate interest. Nevertheless, they form the crucial input with which to compute hedge ratios (and their integral, prices).

It seems to be a roundabout way to compute auxiliary variables such as volatility in order to obtain estimates of quantities that are fundamentally related to correlation, namely hedge ratios. This impression grows when one realizes that the auxiliary variables are not always easily estimated, as pointed out before. Therefore, one ought to try an alternative approach. One that keeps the essence of derivatives analysis (hedging), but that is more amenable to implementation. The new approach should primarily be based on the one aspect of return data that has been proven to be reasonably stable and easily estimable, namely correlation.¹

Such an approach is suggested here. The idea is to formulate the derivatives problem in terms of hedge equations, to be estimated directly with a flexible, yet robust procedure. The estimation technique has become known as *local parametric estimation*: one fits parametric models (such as Black-Scholes) locally, thereby exploiting at a maximum the insights of standard derivatives theory as far as curvature (convexities) is concerned. It has recently been suggested by Hjort [1995].

One of the first issues that our approach can shed light on is the nature of the tracking error from hedging in discrete time. With local parametric estimation, one can directly estimate appropriate portfolio weights with which to duplicate the return on a derivative over noninfinitesimal time intervals. The present paper evaluates the performance of locally estimated hedges and

¹One often overlooks that standard time series estimators of the parameters of the processes of the underlying values, such as maximum likelihood estimators, are not necessarily optimal for somebody whose criterion is defined in terms of hedging (e.g., minimization of the squared tracking error). Decision theory teaches us that the optimal estimators are context-dependent. It would be an interesting theoretical exercise to derive the actual optimal estimator of the hedge ratios in a well-specified case, such as the model of Cox, Ingersoll and Ross [1985], and compare its performance to the maximum likelihood estimator.

compares it to that of Black-Scholes hedges. It follows Hutchinson, Lo and Poggio [1994] in measuring the tracking error. Not surprisingly, the paper reports that the errors from Black-Scholes hedging in discrete time are most important for those options whose return is nonlinearly related to that of the underlying security, namely, short-maturity, out-of-the-money options. Best out-of-sample tracking records are obtained by combining locally estimated hedge portfolio weights with Black-Scholes weights obtained from historical volatility estimates. The optimality of this combination indicates that our local parametric approach to options analysis and traditional derivatives analysis exploit complementary information from historical samples.

When fitting Black-Scholes portfolio weight functions locally, the estimated parameter can be plotted against the two factors, namely moneyness and maturity. In the Black-Scholes world, this parameter equals the volatility of the underlying stock price. It obviously loses this meaning in local parametric estimation. There, it is a parameter that absorbs the misspecification of using Black-Scholes hedge portfolio weights in discrete time, and, hence, will generally depend on moneyness and maturity. Its meaning is better compared to that of the notion of local volatility (used, for instance, in Rubinstein [1994]). Plots of the estimated volatility parameter against moneyness exhibit strong “smile” patterns. Less pronounced is a “wave” in the relationship between locally estimated volatility and maturity. This finding indicates that there is an additional explanation for the empirically observed smile and wave effects in local volatility (see, e.g., Derman and Kani [1994]), namely, the impossibility to hedge in continuous time. This does not mean that discrete-time hedging would displace traditional explanations of the smile effect. One of them, namely, stochastic volatility, is so relevant empirically that it is likely to continue to explain the bulk of the mispricing effects (see Renault and Touzi [1992]).

We should point out here that there is independent work on discrete-time hedging. In particular, Gouriéroux and Laurent [1994] report the results from an alternative approach to estimating the optimal discrete-time hedge ratios. In the spirit of the first author’s work on indirect inference (see Gouriéroux, Monfort and Renault [1993]), an estimation strategy is developed whereby the correct discrete-time hedge formulae (or a linearized version) from a misspecified model are fit to the data. The parameters are not obtained from the minimization of a statistical criterion, but of

a hedging criterion.² While different, the estimation procedure of Gouriéroux and Laurent [1994] generates analogous results. Unfortunately, stochastic interest rates and volatility were introduced simultaneously with discrete-time hedging in their paper, so that the effect of the latter alone is difficult to discern (in particular, they do not report smile and wave effects in the estimates of the volatility parameter). There is also the paper of Liang [1994], which is an embryonic version of ours. He approximates the optimal hedge ratios by means of polynomials and estimates the parameters by least squares. For his procedure to recover the correct parameters, however, the polynomial degree has to increase with the sample size.

The remainder of this paper is organized as follows. Section 2 provides a more detailed discussion of our estimation approach. Section 3 discusses the performance of our approach when hedging in discrete time in a world where Black-Scholes holds. It compares the tracking error with that of using the (continuous-time) hedge portfolio weights suggested by Black-Scholes. Section 4 investigates the patterns in the relation between the locally estimated free parameter, the volatility parameter, on the one hand, and moneyness and maturity, on the other hand. Section 5 concludes by suggesting further research topics.

2 Local Parametric Estimation

2.1 The Derivatives Pricing Problem

To understand our approach, let us introduce a *canonical representation of the derivative securities pricing problem*. Let C denote the value of the derivative and ΔC the change in value. Let V_i denote the value of the i th risky asset used in the replicating portfolio ($i = 1, \dots, n$), and ΔV_i the corresponding change. The first step in derivative securities analysis is to determine coefficients $a()$ and $b_i()$ such that the error e in the following equation is minimized:

$$\frac{\Delta C}{C} = a() + \sum_{i=1}^n b_i() \frac{\Delta V_i}{V_i} + e. \quad (1)$$

²The idea of local estimation is actually used in the specification test in Gouriéroux and Laurent's paper. The theoretical analysis of this specification test can be found in Gouriéroux, Monfort and Renault [1994].

If the sum of squared errors is used as minimization criterion, we essentially ask for the least squares intercept and slope coefficient in the projection of changes in the derivative's value onto the payoffs of the assets constituting the replicating portfolio. The slope coefficients determine the *weights* accorded to each of the assets in the replicating portfolio, with the remainder, namely,

$$1 - \sum_{i=1}^n b_i(),$$

invested in a one-period riskfree asset.

The *power of derivatives securities analysis* lies in allowing the hedge portfolio weights (i.e., the regression coefficients) to vary over time as a function of the available information, such as moneyness, maturity, interest rates, etc. For this reason, we have been writing the coefficients $a()$ and $b_i()$ with parentheses, to indicate that they are functions of as yet unspecified information. But this is also the *difficulty* of derivative securities analysis: in general, it is not clear which hedge functions would lead the error to be negligible.

We have written Equation (1) in terms of *returns*. Usually, however, theoretical derivatives pricing problems are formulated in terms of *payoffs*. Once one realizes that we are going to estimate the hedge portfolio weights, it should be clear why the former is preferable. It is likely that the values of the underlying assets are nonstationary. For instance, they could follow geometric Brownian motions. In that case, returns are stationary, but the payoffs on the underlying assets and those of a derivative with a fixed moneyness will be nonstationary, rendering the theoretical-statistical analysis of estimation of (1) extremely difficult if it had been written in terms of payoffs.

In particular parametric cases, the replication error can be reduced to zero by a judicious choice of $a()$ and $b_i()$. One example is Black and Scholes [1973]' model. There, the derivative is a call option, written on common stock. One uses only one risky security in the hedge portfolio ($n = 1$), namely the stock itself. Let m denote the option's moneyness, i.e., its stock price divided by the exercise price, and let τ denote its maturity. Let r be the interest rate, assumed positive

and constant, and let σ be the instantaneous volatility of the stock price. Set:³

$$b_1(m, \tau; \sigma) = \frac{1}{1 - \frac{e^{-r\tau}}{m} \frac{N(d_2(m, \tau; \sigma))}{N(d_1(m, \tau; \sigma))}}, \quad (2)$$

and $a() = (1 - b_1(m, \tau))\gamma$, for some $\gamma > 0$. $N()$ denotes the standard normal distribution function, and:

$$\begin{aligned} d_1(m, \tau; \sigma) &= \frac{\log m + r\tau}{\sigma\sqrt{\tau}} + \frac{1}{2}\sigma\sqrt{\tau}, \\ d_2(m, \tau; \sigma) &= d_1(m, \tau; \sigma) - \sigma\sqrt{\tau}. \end{aligned}$$

This choice for the hedge portfolio weight eliminates the tracking error ($e = 0$) provided: (i) the stock price follows a geometric Brownian motion, (ii) the hedging interval is infinitesimal. If there are no arbitrage opportunities, $\gamma = r$.

2.2 Estimation Of Hedge Portfolio Weights

Instead of deriving the functional form of the coefficients $a()$ and $b_i()$ from a full specification of the stochastic properties of the underlying assets, we propose to *directly estimate* them in a robust way from the correlation properties in a dataset of call price changes and changes in the values of the underlying assets. A rough procedure would be to approximate the weight functions with polynomials and apply least squares (see Liang [1994]). The correctness of this approach would require that the degree of the polynomial increase with the sample size.

Instead, we propose *local parametric estimation*. This is a cross between least squares and kernel estimation. The idea is the following. To obtain estimates of $a()$ and $b_i()$ at a particular value for the factors on which they (the hedge portfolio weights) are supposed to depend, one implements weighted least squares, using observations with neighboring values for the factors, whereby the weights depend on how far away the latter are. To obtain best results, one does

³ $b_1(m, \tau; \sigma)$ is usually referred to as the hedge elasticity, which is the hedge ratio divided by the option premium. By Itô's Lemma, the hedge ratio is the derivative of the theoretical call price with respect to the price of the underlying asset. Eqn. (2) is then obtained after substituting the analytical call price formula for the call price in the definition of the hedge elasticity.

not fit a linear or polynomial function locally. Instead, a parametric model is used. For instance, the Black-Scholes hedge portfolio weights (Eqn. (2)) could be fitted locally. This exploits at a maximum the local curvature (convexities) of parametric option pricing models.

The resulting procedure, local parametric estimation, has recently been proposed by Hjort [1995] in the context of duration models. It is an extension of local polynomial estimation, when polynomials are fitted locally, initially proposed by Stone [1977] and Cleveland [1979].

To illustrate the approach, consider hedging the payoff on a call (with price C) with the underlying asset and the riskfree security, assumed to have a constant return r . In Equation (1), $n = 1$. Let S denote the value of the underlying asset, i.e., $S = V_1$. We will refer to the underlying asset as the “stock.” ΔS denotes the change in the stock price. There is a dataset of M observations $((\Delta C/C)_j, (\Delta S/S)_j, m_j, \tau_j)$ ($j = 1, \dots, M$), where m_j and τ_j denote the option’s moneyness and maturity, respectively, for observation j .

Consider fitting Black-Scholes hedge portfolio weights (Eqn. (2)) locally to obtain $b_1(m, \tau)$. To estimate $a(m, \tau)$, one could fit one minus the Black-Scholes weight function times a constant, γ . This would mean that one minimizes the weighted sum of squared error in the following equation with respect to the two parameters, σ and γ :

$$\left(\frac{\Delta C}{C}\right)_j = (1 - b_1(m_j, \tau_j; \sigma))\gamma + b_1(m_j, \tau_j; \sigma) \left(\frac{\Delta S}{S}\right)_j + u_j. \quad (3)$$

If $\hat{\gamma}$ and $\hat{\sigma}$ denote the optimum, the estimates of $a(m, \tau)$ and $b_1(m, \tau)$ are defined as $(1 - b_1(m, \tau; \hat{\sigma}))\hat{\gamma}$ and $b_1(m, \tau; \hat{\sigma})$, respectively.

The weights in the optimization are determined by the values of *kernel functions* evaluated at the distance between (m_j, τ_j) and (m, τ) . An example is the Epanechnikov kernel, defined as follows:

$$K_h(u) = 0.75 \left(1 - \left(\frac{u}{h}\right)^2\right) 1_{\{|u/h| \leq 1\}},$$

where $1_{\{\cdot\}}$ denotes the indicator function. The argument of the function, u , could be the Euclidean distance between (m_j, τ_j) and (m, τ) .

The parameter of the weighting function, h , is referred to as the *bandwidth*. As the sample size increases, this parameter can be decreased in order to obtain consistency. A small bandwidth

generates low bias but may cause overfitting, and, hence, excessive variance. A large bandwidth tends to smooth fits, but increases the bias. The issue of bandwidth selection is important, and we will come back to it shortly.

2.3 The Nature Of Local Parametric Estimation

With the help of Figures 1 and 2, we can get a better idea of the nature of local parametric fitting of Black-Scholes portfolio weights. Figure 1 shows a plot of $\Delta C/C$ against $\Delta S/S$ and m , keeping τ constant at 0.1. Call price changes are computed from the stock price changes using the Black-Scholes formula. The following parameters were used: volatility = 0.4; stock price drift = 0.15; riskfree rate (r) = 0.05. Stock returns ($\Delta S/S$) and corresponding option returns ($\Delta C/C$) are measured over a time interval set equal to 0.01, the equivalent of approximately 2.5 trading days if one unit of time corresponds to one trading year.

The linearity of the relation between $\Delta C/C$ and $\Delta S/S$ for large m is apparent in Figure 1. This means that Eqn. (3) can be made to hold with almost zero error ($u_j = 0$) for an appropriate choice of volatility parameter σ . In fact, the best choice will be given by the true volatility of the underlying stock. In contrast, the relationship is clearly nonlinear for deep out-of-the-money options (small m). The choice of σ that minimizes the error u_j of (3) is now less obvious. Local parametric estimation provides a convenient way of determining the correct value. Keeping $\Delta S/S$ constant, the nonlinearity of the relation between $\Delta C/C$ and m stands out. This nonlinearity is captured by varying the free parameter σ in the hedge portfolio weight of Eqn. (3) as a function of m .

(Figure 1 about here)

Figure 2 displays a plot of $\Delta C/C$ against $\Delta S/S$ and τ , keeping m constant at 0.8. In all other respects, the parameter values are the same as in Figure 1. Figure 2 reveals that Eqn. (3) can be fit most easily for large values of τ . The appropriate choice of σ for low values of τ (short-maturity options), however, is less obvious. Again, local parametric estimation will determine the right

value. The nonlinearity in the direction of changing maturity will be captured by varying the free parameter in Eqn. (3), σ , as a function of maturity.

(Figure 2 about here)

In our application of local parametric estimation, the best choice of σ for each moneyness (m) and maturity (τ) will be determined in an empirical fashion: we use a dataset that is simulated under the Black-Scholes assumptions and find the values of σ that minimize the weighted sum of squared errors. This approach has the advantage of weighing the different outcomes (moneyness, maturity, stock price changes) with the frequencies that are consistent with the theoretical model. In the real-life implementation that we are presently carrying out (but not reporting here), the optimal choice of σ is obtained from the empirically observed frequency with which moneyness, maturity and stock price changes occur, rather than those obtained from simulating a theoretical model.

Traditional option pricing theory specifies assumptions such that equations like (3) hold with zero error. In Black-Scholes, for instance, it is argued that instantaneous large stock price changes happen with zero probability, and, hence, that a linear relationship like the one of Eqn. (3) obtains without error (a consequence of Itô's lemma). Such a theory is therefore essentially a linear approximation argument justified by the impossibility of large changes over short periods of time. Likewise, the binomial option pricing model is based on the view that stock price changes over a discrete time interval can take only two possible values. Eqn. (3) without error obtains trivially.

In discrete time, or when stock price changes are not binomial, tracking errors will be non-trivial. In order to understand those tracking errors when Black-Scholes portfolio weights are blindly applied to discrete-time hedging, Figures 3 and 4 replicate Figures 1 and 2, but, instead of displaying the actual option return against stock price changes, they show the *theoretical* option return that would have resulted *if* continuous-time relationships were extendable to discrete time. The theoretical continuous-time option return is given by Itô's lemma.

(Figure 3 about here)

A comparison of Figures 3 and 4 against 1 and 2 immediately reveals the inadequate nature of Black-Scholes hedge ratios when used to hedge deep out-of-the-money and short-maturity options. It is most apparent in the fact that the theoretical option return could be below -1 even for stock price changes that cannot be considered unreasonable over a time interval equal to 0.01 (approximately two trading days if one unit of time corresponds to a year)! Hence, the errors from blindly applying Black-Scholes in a discrete-time setting can be large. This paper studies how relevant these errors are *overall* if the theoretical discrete-time frequency of stock price changes from the Black-Scholes model is used as weighting norm.

(Figure 4 about here)

The discussion of Figures 1 through 4 is a good occasion to emphasize that our procedure implicitly takes into account the drift in stock prices when determining hedge portfolio weights. The optimal hedge ratio could indeed be a function of the underlying stock's drift, because the option's expected return will often be an explicit function of the latter (see, e.g., Rubinstein [1984]). Yet, the Black-Scholes hedge ratio ignores this drift, and, hence, would insure incorrectly if applied in a discrete-time setting.

Before we turn to a discussion of the bandwidth selection, we should briefly mention that we also estimated the hedge portfolio weights by means of *local polynomial estimation*. With this technique, polynomials are fit locally, instead of a parametric model as in Eqn. (3). Local polynomial estimation was originally suggested in Stone [1977] and Cleveland [1979], and later analyzed in Tsybakov [1986], Tibshirani and Hastie [1987], Cleveland, Devlin and Grosse [1988], Staniswalis [1989], Fan [1993], Fan and Gijbels [1992], Ruppert and Wand [1994], and applied in a financial context in Bossaerts, Härdle and Hafner [1995], Bossaerts, Hafner and Härdle [1995] and Gouriéroux and Scaillet [1994].

We fitted linear and quadratic functions locally. The performance of these estimates, however, were far inferior. This indicates that parametric models such as Black-Scholes provide useful information about the local curvature of the duplicating portfolio weights as a function of moneyness

and maturity. We decided not to report any results from the more agnostic local polynomial estimation here. Instead, we focus entirely on local parametric estimation.

2.4 Bandwidth Selection

Until now, we have presented local parametric estimation of hedge portfolio weights as a simple exercise of local nonlinear least squares fitting. It provides the estimates that minimize the weighted sums of hedge errors. This estimation criterion, however, can be justified from specific *statistical assumptions*. An exploration of these assumptions is not only illuminating about the kind of restrictions local parametric estimation exploits in the data. It also facilitates bandwidth selection.

Let us introduce a time index, t . In terms of our example of call options, to be hedged with the underlying stock, local parametric estimation then generates estimates of the functions $a()$ and $b_1()$ in the following equation:

$$\frac{\Delta C_{t+1}}{C_t} = a() + b_1() \frac{\Delta S_{t+1}}{S_t} + e_{t+1}. \quad (4)$$

Local parametric estimation provides consistent estimates of these functions if there indeed exist smooth $a()$ and $b_1()$ which are functions *solely* of moneyness m and maturity τ such that

$$E[e_{t+1} | \mathcal{F}_t] = 0, \quad (5)$$

and

$$E[e_{t+1} \frac{\Delta S_{t+1}}{S_t} | \mathcal{F}_t] = 0, \quad (6)$$

where $\mathcal{F}_t$ is information available at time t (which consists, at a minimum, of the option's moneyness and maturity). Hence, our procedure is justified if $a()$ and $b_1()$ are functions of only moneyness and maturity, and neither the hedge error nor the correlation between the hedge error and stock price change are predictable from past information.

The assumptions in (5) and (6) facilitate optimal bandwidth selection by means of *crossvalidation*. In this technique, an out-of-sample hedge error is computed for each observation. This

is accomplished by estimating $a()$ and $b_1()$ on the basis of information that is orthogonal to the true hedge error for the observation at hand. Subsequently, one sums the squared estimated out-of-sample hedge errors and minimizes the result with respect to the bandwidth parameter.

Assumptions (5) and (6) determines that prior information is orthogonal to the hedge error of a particular observation. Consequently, the out-of-sample hedge error could be computed on the basis of estimates of $a()$ and $b_1()$ that make use only of prior observations. Pairs of call and stock price changes that occur simultaneously with a given observation must *not* be used.

Notice that this crossvalidation technique differs from the standard one. In the latter, an out-of-sample hedge error would be computed on the basis of estimates of $a()$ and $b_1()$ that use all data points except the observation at hand. But Assumptions (5) and (6) allow hedge errors to be predictable from, among other things, contemporaneous hedge errors on calls with a different moneyness and/or maturity. If one were to implement the traditional crossvalidation technique, serious overfitting will result: the hedge errors are not entirely out-of-sample, because correlated information is used in the estimation.

In a previous version of this paper, we reported the out-of-sample performance of hedges based on local parametric analysis where we had employed the standard crossvalidation technique. We reported serious overfitting: the locally estimated volatility parameter changed erratically with moneyness and maturity. Altering the crossvalidation procedure in accordance with the assumption in (5) and (6) lead to a much smoother relationship between the locally estimated volatility parameter and moneyness or maturity. It also improved substantially the out-of-sample hedging performance. We will report those results in the next section.⁴

⁴By setting the weights in the out-of-sample weighted mean square prediction error equal to zero for some observations, a drastic reduction in computation time can be obtained without invalidating the bandwidth selection technique. We implemented this trick because the size of our samples made comprehensive cross-validation impracticable. In previous versions of this paper, we also increased the bandwidth whenever less than 40 observations were assigned positive weights. This bandwidth adjustment was essentially a simple version of the Gasser-Müller estimation technique (Gasser and Müller [1979]). As the sample size increases, however, the adjustment would have become unnecessary, which implies that the usual asymptotic properties continue to hold. With the new bandwidth selection technique, however, this adjustment appeared unnecessary.

Let us add here that assumptions like the ones in (5) and (6) could easily be tested. Since we are in a local estimation context, the techniques of Gouriéroux, Monfort and Tenreiro [1994] and Gozalo and Litton [1994] could be used. Such tests are an important way of assessing the validity of a crossvalidation technique in a particular context. Absent this, only an extensive out-of-sample performance analysis would indicate whether a bandwidth selection technique is appropriate.

3 The Tracking Error of Hedge Portfolios Obtained From Local Parametric Estimation

We implemented several variations of local parametric estimation and investigated their out-of-sample hedging performance against Black-Scholes hedges. In the results to be reported here, we fitted Black-Scholes portfolio weights locally while setting $\gamma = r$ in Eqn. (3).

We followed closely Hutchinson, Lo and Poggio [1994] in the setup of the hedging performance evaluation. We thus generated several “training samples”, to be used to (i) determine the optimal bandwidth through cross-validation as explained in the previous section, (ii) provide the necessary data for local fits with which to carry out an out-of-sample tracking analysis. The training samples were obtained from simulating stock price changes over two years, introducing call options in the way that the CBOT does, and computing the theoretical Black-Scholes prices for each of them. Since the interval between portfolio rebalancing was set at 0.01, 200 stock price changes were obtained, as well as price changes for all the calls that were trading concurrently. Typically, this generated a dataset of about 3,000 call price changes and corresponding stock price changes. The stock price drift and volatility were set equal to 0.15 and 0.40, respectively, and the short-term interest rate was fixed at 0.05.

Although the training sample included several thousand observations, we picked only 200 observations to determine the optimal bandwidth. All training sample observations were used to obtain both the in-sample cross-validation fits as well as the out-of-sample fits. We observed little improvement in the out-of-sample hedging performance when increasing the number of observa-

tions that were included in the cross-validation exercise.⁵

The out-of-sample hedging performance was evaluated on fifty “testing paths”. We generated independent stock price paths and observed the performance of a portfolio long in a call option with a certain moneyness and a maturity of 0.1 and short in the duplicating portfolio. The duplicating portfolio weights were determined by the local parametric estimation results or by Black-Scholes. In the latter case, we used the training-sample estimate of the volatility of the stock price (instead of the true volatility, as in Hutchinson, Lo and Poggio [1994]).⁶

For each variation of local parametric estimation, we repeated ten times the construction of training samples, the cross-validation and the out-of-sample performance evaluation on fifty independent testing samples. This produces 500 out-of-sample tests in total.

The out-of-sample hedging performance was evaluated as in Hutchinson, Lo and Poggio [1994] by computing several statistics generated from the model’s “tracking error”. The latter is essentially the *maturity-date dollar payoff on a portfolio long one call option and short the duplicating portfolio*. We looked at the average absolute tracking error, the average squared tracking error and the frequency with which the absolute tracking error of the locally estimated duplicating portfolios was lower than that of Black-Scholes.

We soon observed a bias in the out-of-sample tracking error of locally estimated duplicating portfolios. In particular, the average payoff of the portfolio long in the call and short in the duplicating portfolio was almost invariably positive when using locally estimated portfolio weights. With Black-Scholes hedge ratios, the average payoff was much closer to zero. To understand the impact of this bias, we also report the standard deviation of the tracking error, as well as the frequency with which the (signed) payoff of the portfolio with the hedged call was higher under

⁵We selected observations $[N/200]$, $[2N/200]$, ..., N , where N denotes the sample size and $[\cdot]$ denotes the rounding function. Because the data are ordered, first by maturity, then by exercise price and finally by trading time, a representative subsample was generated. The representativeness was confirmed when we observed only miniscule changes in the optimal bandwidth when increasing the crossvalidation sample size.

⁶The relative hedging performance of Black-Scholes hedges were hardly affected by the use of estimated volatility in lieu of the true one. Of course, this reflects the fact that the volatility of a geometric Brownian motion can be estimated very accurately over only 200 observations.

local parametric estimation than under Black-Scholes.

One could attribute the bias in the out-of-sample tracking error of locally estimated hedges to the well-known biases of local estimation when an optimal bandwidth is selected. This conjecture was proven wrong in the following experiment. Biases in local estimation decrease as the bandwidth is lowered. Hence, we ought to observe decreases in the bias of the out-of-sample tracking error of locally estimated hedge portfolios as the bandwidth is reduced. Instead, we recorded *no* changes in the average return on the hedge portfolio, rejecting the conjecture. As we will argue later, the average positive return on the hedge portfolio may reflect mispricing of options in a world where one can only hedge in discrete time.

The average correlation between the tracking error of locally fitted hedge portfolio weights and Black-Scholes portfolio weights was found to be surprisingly low.⁷ Hence, we also investigated the tracking performance of a portfolio whose hedge ratio is obtained as an equally weighted average of the locally fitted ratio and the Black-Scholes ratio. As we will see, this improves on either way of obtaining the hedge ratio. In other words, it is preferable to combine both procedures. This finding essentially means that *hedge ratios estimated with local parametric estimation are based on different information from that used to compute the historical volatility*. The latter is used to compute the Black-Scholes hedge ratio. Consequently, the optimal strategy from a decision-theoretic point of view uses a combination of traditional option pricing and local parametric derivatives analysis.

One could also conjecture that the improved performance of the combined hedge strategy is due to the superiority of local parametric analysis for hedging of out-of-the-money options, while this superiority disappears as the derivative moves in-the-money. To evaluate this possibility, we also investigated the hedging performance of a strategy whereby we switched from locally estimated hedge portfolio weights to Black-Scholes weights as the option's moneyness increased above 1.05.⁸ It will be clear from the results, however, that this strategy is dominated even by the one where

⁷A similar phenomenon is present in the results reported in Gouriéroux and Laurent [1994].

⁸We thank an anonymous referee for suggesting this tactic.

locally estimated hedges are used throughout. Consequently, this alternative explanation of the impressive performance of the combined strategy could be proven false.

Finally, we checked the performance of local parametric estimation for options with extremely low moneyness and maturity. In such cases, the hedge portfolio weights are very large. When transformed to hedge ratios (defined to be the *units* of stock to be shorted in order to hedge the payoff on one option), however, the values should still be below one. Portfolio weights are translated to hedge ratios using the theoretical Black-Scholes call prices, which, unfortunately, do not necessarily provide the correct values when one is forced to hedge in discrete time. Hence, the corresponding hedge ratios may violate the bound without being incorrect. Hedge ratios were computed to be above one on average in less than 1% of the cases. Hence, the hedge portfolio weights translate well into hedge ratios even when using theoretical Black-Scholes option values. Also, we never observed hedge portfolio weights below one, thereby validating the lower bound.

Let us now turn to a discussion of the results. Table 1 displays the out-of-sample hedging performance of (i) hedges obtained from locally fitting Black-Scholes portfolio weights, (ii) Black-Scholes hedges based on historical volatility, (iii) hedges obtained as an equally weighted average of the former and Black-Scholes weights, (iv) hedges based on local parametric estimation, but switching to Black-Scholes as the moneyness of the option increases beyond 1.05. Results are reported for initial moneyness equal to 0.80, 0.85, ..., 1.05.⁹

(Table 1 about here)

Table 1 focuses on short-maturity options (maturity equals 0.1, or about 1.2 calendar months). The improvement of local estimation over Black-Scholes is pronounced for out-of-the-money options. The superiority is clearest in terms of squared tracking error or standard deviation of the (signed) tracking error, which means that someone with a quadratic loss function would be espe-

⁹In all the results to be reported, the bandwidth was set equal to the average optimal bandwidth from several trial runs, namely, $N^{-\frac{1}{5}}$, where N is the sample size. The distance measure used in the Epanechnikov kernel is the simple weighted Euclidean distance, with weights determined by the relative range of the two inputs (moneyness and maturity).

cially attracted by our technique. There are biases, as already pointed out before: the average return on the hedge portfolio based on local parametric estimation is *positive*.

The reduction in performance of local parametric estimation as a function of the initial moneyness of the option is not surprising. For in-the-money options, the relationship between call returns and stock returns is essentially linear (see Figures 1 and 2). Black-Scholes hedges *assume* linearity (see Figures 3 and 4). Local parametric estimation, however, is designed to capture *nonlinearities*, and, therefore, can be expected not to outperform when the true relationship is linear.

The tracking error is measured as the dollar payoff on a portfolio long in the option and short in the hedge. Hence, for an initial moneyness of 0.8, Table 1 indicates that the square root of the average squared payoff (root mean square tracking error) on hedges based on local parametric estimation is \$4.90, whereas that based on Black-Scholes equals \$8.34, a 41% improvement. This superiority reduces to less than 1% for at-the-money options ($m = 1.00$) (\$13.33 vs. \$13.39).

By far the most impressive performance is generated by the combination of hedges based on local parametric estimation and Black-Scholes: even for initially in-the-money options ($m = 1.05$), the reduction in root mean square tracking error is above 15% (above 9% for the absolute tracking error). As mentioned before, this is due to the less-than-perfect correlation between the tracking errors from local parametric estimation and Black-Scholes (0.469 across all the simulations used for Table 1).

The low correlation between the hedge error using local parametric estimation and that from Black-Scholes is a reflection of the low correlation between the sampling errors of the statistics behind each methodology. Black-Scholes hedges use the historical volatility as main statistical input; local parametric estimation exploits the correlation between call and stock returns. It appears that the error from estimating historical volatilities and correlations are not perfectly correlated. Hence, an improvement in the out-of-sample hedging performance is obtained by combining both procedures.

The promising track record of the combination of Black-Scholes hedges with locally estimated portfolio weights cannot be attributed to the former's enhanced performance for in-the-money

options. If that were the case, a policy whereby one switches from local parametric analysis to Black-Scholes from the moment the option’s moneyness reaches 1.05 would do much better. In fact, Table 1 documents that such a policy is inferior across the board.

The maturity of the options on which the performance analysis of Table 1 is based is short: only about 1.2 months. As we increase the maturity, the results will not alter, but there will be less differentiation across levels of initial moneyness. We did observe this. Table 2, therefore, reports only the *overall* results from a replication of the analysis behind Table 1 for a maturity of 0.6 (about 7.2 months). The numbers are based on simulations that are independent of those used to generate Table 1.

(Table 2 about here)

For this aggregate sample, local parametric estimation keeps its lead over Black-Scholes in terms of root mean square tracking error (\$8.72 vs. \$9.68, a 10% improvement). As mentioned before, this is important for hedgers with a quadratic loss function, i.e., those who would prefer to penalize outliers heavily. The outperformance of local parametric estimation disappears in terms of absolute tracking error.

Most impressive is the uniform outperformance of the combined policy: even in terms of absolute tracking error, the improvement is of the order of 9% (\$3.40 vs. \$3.73). *The root mean square tracking error is reduced by 18%.*

Let us emphasize again that the data on which the performance analysis is based have been generated assuming Black-Scholes. The nice results from local parametric analysis can only be due to its ability to track the option return better in a discrete-time framework. Because of this, we expect the improvements of local parametric analysis over traditional Black-Scholes hedging to lower as we (i) move towards hedging in continuous time, (ii) reduce the volatility. A reduction in volatility, however, may improve the fit from local parametric estimation, because the training sample will include more observations with a moneyness and maturity similar to the ones used in the out-of-sample performance tests.¹⁰ Moreover, the number of observations in the training

¹⁰Remember that the training sample is constructed to reflect actual options data from the Chicago options

sample also increases, because we kept the original length of the sample as before (two years), but decreased the return interval from about 2.5 to one trading day. This lead to an increase in the number of observations in a typical training sample from about 3000 to 4500, generating a corresponding improvement in the precision.

Table 3 documents the trade-off between a reduction in outperformance of local parametric analysis over Black-Scholes as a result of more frequent rebalancing and lower volatility, on the one hand, and higher precision because of lower volatility and bigger training samples, on the other hand. It repeats the analysis of Table 1, but uses the parameters in Hutchinson, Lo and Poggio [1994], whereby rebalancing happens every trading day (based on a 253-day year) and the volatility is reduced to 0.20. This volatility value would be adequate for stock indices or foreign currency, whereas the value in Tables 1 and 2 (0.40) reflects the level for a typical common stock.¹¹ A closer look at Table 3 will reveal that local parametric analysis keeps its lead over Black-Scholes, despite the substantial increase in rebalancing frequency and concurrent decrease in volatility. The same pattern as in Table 1 appears: the outperformance of local parametric estimation and the combined strategy is most pronounced in terms of root mean square error. In terms of mean absolute tracking error, however, Black-Scholes beats hedging based on local parametric estimation alone, at least as far as at-the-money and in-the-money options are concerned. The switching strategy is dominated overall.

markets. When the volatility is high, short-maturity options with extreme moneyness are more likely than when the volatility is low.

¹¹The stock price drift is set equal to 0.1. This is perhaps a good occasion to point out that the purpose and approach of Hutchinson, Lo and Poggio [1994] differs markedly from ours. For one thing, they focus on nonparametric estimation of the *pricing function*. We estimate hedge portfolio weights. Also, they obtain hedge ratios as *first derivatives* of the estimated pricing function, which they convert to hedge portfolio weights, to be used in the out-of-sample tracking evaluation. We estimate the hedge portfolio weights directly. It should not come as a surprise, then, that in one of their analyses, where they assume Black-Scholes as we do, hedge portfolios on the basis of first derivatives of nonparametrically estimated pricing functions perform worse than Black-Scholes. Improvements over Black-Scholes can only be expected when Black-Scholes is known to be misspecified, as in the second application of Hutchinson, Lo and Poggio [1994], where real-market option data are investigated.

(Table 3 about here)

4 Patterns In The Locally Estimated Volatility Parameter

Local parametric estimation is based on locally fitting a parametric model using nonlinear least squares. If Black-Scholes is used as parametric model, one parameter only needs to be fit (provided one sets $\gamma = r$ in Eqn (3)), namely the one that plays the role of volatility in the original model. Because this parameter is fit locally, the estimates will generally differ across values of the factors that determine the optimal hedge ratio in a Black-Scholes world, namely moneyness and maturity. It is interesting to investigate what the differences are between the estimated volatilities across moneyness and maturity.

In the case of moneyness, the relation of the estimated volatility and moneyness is reminiscent of the well-known empirical “smile” effect. Here, it means that the optimal volatility parameter to be used to hedge deep out-of-the-money or in-the-money call options is higher than the one needed to hedge at-the-money options. Figure 5 illustrates this. It provides a scatter plot of volatilities estimated using local parametric estimation against moneyness for a randomly chosen training sample.

(Figure 5 about here)

Figure 5 may be confusing to some, because it displays locally estimated volatilities for several maturity levels. In order to enhance the interpretation, Figure 7 provides a three-dimensional plot of locally estimated volatility against moneyness and maturity. From it, one can deduce that the plot in Figure 5 should be interpreted as follows. The locally estimated volatilities for short-maturity options lie on a sharp V curve, extending as a straight line for moneyness beyond 1.25. As the maturity increases, the V curve becomes flatter, generating, for long-maturity options, a straight line through 0.4.

The effect of maturity on locally estimated volatility can also be seen from Figure 6. The uniform optimality of a volatility of 0.4 for long-maturity options is apparent. This is not surprising: we already know from Figure 2 that the relationship between call and stock returns is linear for long-maturity options. As the maturity decreases, the range of local volatilities increases. While not apparent in Figure 6, there is a significant “wave” effect in the relation between local volatility and maturity for deep out-of-the-money options. It is very apparent in Figure 7.

(Figure 6 about here)

(Figure 7 about here)

While the smile effect in volatilities implied from real-market call price data (using Black-Scholes) has been widely discussed, the wave effect seems to have gone unnoticed, but can nevertheless be discerned. See, e.g., Figure 7a in Derman and Kani [1994]. Is there any relation between these empirical effects and the patterns one observes in a local parametric analysis of hedging in discrete time on simulated data?

It turns out that there is. Take our smile effect (Figure 5), for instance. One can show that it implies that one needs to short more stock to hedge the risk of a deep out-of-the-money call than Black-Scholes prescribes. This can be translated in terms of a portfolio which would be used to create a riskfree asset. The portfolio is long one unit of stock and short in the call. It needs, however, a less extreme short position in the call than Black-Scholes advises. Provided the stock price and the riskfree rate are unaffected (i.e., provided markets are effectively complete), there will therefore be an upward pressure on the call price. The no-arbitrage price level will end up above the Black-Scholes level. Hence, the Black-Scholes implied volatility would be higher than the true volatility.

Notice that one could therefore explain effects such as the “smile” in terms of constraints on the frequency of hedge portfolio rebalancing. This contrasts with traditional explanations, which have focused on misspecification in the process of the underlying asset price. Renault and Touzi [1992], for instance, argue that the smile effect could be the consequence of misspecification of

Black-Scholes when the stock price process exhibits stochastic volatility. Of course, we do not want to push this point too far: stochastic volatility is a matter of fact in real markets, and, hence, a prime source for misspecification of Black-Scholes models.

Our alternative explanation of the smile effect, however, should be welcomed in view of the recent evidence that stochastic volatility alone cannot explain the observed smile. If the smile were to be driven solely by stochastic volatility, it would imply more persistence in volatility than is actually present in the data (see Dumas [1995]). We provide an explanation of the smile effect that could bridge the gap between mean-reverting volatility, on the one hand, and the observed smiles, on the other hand. Ours is based on factual constraints on rebalancing frequencies.

5 Conclusion

Local parametric estimation combines the power of analytical derivative securities analysis with the flexibility of local least squares estimation. The result is a technique that generates significant improvements in hedging performance in situations where the theoretical option pricing model is suspected to be inadequate, such as in discrete-time hedging.

The success of local parametric estimation for discrete-time hedging in a world where Black-Scholes holds inspired us to investigate further issues, such as out-of-sample hedging based on real-world data, outlier control, stochastic interest rates and volatility. We will report on those in separate papers.

Finally, another urgent topic for research is that of tracking error control. Our hedge portfolios blindly followed their prescribed path (as did those of Hutchinson, Lo and Poggio [1994]) even if it was clear along the way that they were off-track. No self-correction was built into the hedging procedure. Substantial improvements are likely if such correcting mechanisms are introduced into the hedging policy. Of course, one thereby deviates from the standard analysis of selffinancing portfolios: portfolios with feedback and control will need interim capital input or may generate excessive funds along the way. Again, this is an issue we are presently investigating.

References

- Ait-Sahalia, Y., 1994, "Testing Continuous-Time Models of the Spot Interest Rate," University of Chicago working paper.
- Black, F. and M. Scholes, 1973, "The Pricing of Options and Corporate Liabilities," *Journal of Political Economy* 81, 637-659.
- Bossaerts, P., C. Hafner and W. Härdle, 1995, "Foreign Exchange Rates Have Surprising Volatility," Institut für Statistik und Ökonometrie, Humboldt Universität zu Berlin SFB, working paper.
- Bossaerts, P., W. Härdle and C. Hafner, 1995, "A New Method for Volatility Estimation with Applications in Foreign Exchange Rate Series," *Physica Verlag*, forthcoming.
- Cleveland, W., 1979, "Robust Locally Weighted Regression and Smoothing Scatterplots," *Journal of the American Statistical Association* 74, 823-836.
- Cleveland W., S. Devlin and E. Grosse, 1988, "Regression By Local Fitting," *Journal of Econometrics* 37, 87-114.
- Cox, J., S. Ross and M. Rubinstein, 1979, "Option Pricing: A Simplified Approach," *Journal of Financial Economics* 7, 229-263.
- Derman, E. and I. Kani, 1994, "The Volatility Smile and Its Implied Tree," Goldman Sachs Quantitative Strategies Research Notes.
- Dumas, B., 1994, "Implied Volatility Trees and the Empirical Process for the Price of the Underlying Asset: The Case of SP500 Options," Duke University and HEC working paper.
- Fan, J., 1993, "Local Linear Regression and their Minimax Efficiencies," *Annals of Statistics* 21, 196-216.

- Fan J. and I. Gijbels, 1992, "Variable Bandwidth and Local Linear Regression Smoothers," *Annals of Statistics* 20, 2008-2036..
- Gasser, T. and H.-G. Müller, 1979, "Kernel Estimation of Regression Functions," in: *Smoothing Techniques for Curve Estimation*, Lecture Notes in Mathematics 757, Berlin: Springer verlag.
- Gozalo, P. and O. Linton, 1994, "Local Nonlinear Least Squares Estimation: Using Parametric Information Nonparametrically," Yale University, working paper.
- Gouriéroux, C. and J.P. Laurent, 1994, "Parameter Estimation for Discrete Time Dynamic Hedging," CREST working paper.
- Gouriéroux, C., A. Monfort and E. Renault, 1993, "Indirect Inference," *Journal of Applied Econometrics* 8, S85-S118.
- Gouriéroux, C., A. Monfort and C. Tenreiro, 1994, "Kernel M-Estimators: Non-Parametric Diagnostics for Structural Models," CEPREMAP working paper.
- Gouriéroux, C. and O. Scaillet, 1994, "Determination of the Term Structure From Bond Data: Correction for Cash-Flow Patterns," CREST working paper.
- Härdle, W., 1992, *Applied Nonparametric Estimation*, Econometric Society Monograph, Cambridge University Press.
- Hjort, N., 1995, "Dynamic Likelihood Hazard Rate Estimation," *Biometrika*, forthcoming.
- 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.
- Liang, J.X., 1994, "A Nonparametric Approach to Hedging Derivative Securities," Caltech working paper.

- Renault, E. and N. Touzi, 1992, "Option Hedging and Implicit Volatility," CREST/GREMAQ working paper.
- Rubinstein, M., 1984, "A Simple Formula for the Expected Rate of Return of an Option over a Finite Holding Period," *Journal of Finance* 39, 1503-1509.
- Rubinstein, M., 1994, "Implied Binomial Trees," *Journal of Finance* 49, 771-818.
- Ruppert, D. and M.P. Wand, 1994, "Multivariate Locally Weighted Least Squares Regression," *Annals of Statistics*, forthcoming.
- Staniswalis, J., 1989, "The Kernel Estimate of a Regression Function in Likelihood-Based Models," *Journal of the American Statistical Association* 84, 275-283.
- Stone, C., 1977, "Consistent Nonparametric Regression," *Annals of Statistics* 5, 595-645.
- Tibshirani, R. and T. Hastie, 1987, "Local Likelihood Estimation," *Journal of the American Statistical Association* 82, 559-567.
- Tsybakov, A., 1986, "Robust Reconstruction of Functions by the Local-Approximation Method," *Problems of Information Transmission* 22, 133-146.

Table 1
Performance Of Hedges Based On Local Parametric Estimation
Against Traditional Black-Scholes Hedges

		Absolute Tracking Error		Squared Tracking Error Root Mean	Tracking Error		
		Mean	%		Mean	Average S.Dev.	%
<u>$m = 0.80$</u>							
	LPE	1.55	40	4.90	0.36	4.87	50
	BS	2.20	NA	8.34	-0.84	8.27	NA
	50/50	1.80	47	5.78	-0.24	5.75	50
	SWITCH	1.84	40	6.41	0.30	6.30	50
<u>$m = 0.85$</u>							
	LPE	3.84	43	9.17	1.57	8.96	44
	BS	4.59	NA	11.33	0.74	11.21	NA
	50/50	4.04	49	9.24	1.15	9.11	44
	SWITCH	4.13	42	10.30	1.85	10.04	46
<u>$m = 0.90$</u>							
	LPE	6.62	41	11.84	1.30	11.67	42
	BS	7.30	NA	14.46	-1.40	14.27	NA
	50/50	6.47	44	11.18	-0.05	11.09	42
	SWITCH	6.95	41	12.63	2.69	12.27	49
<u>$m = 0.95$</u>							
	LPE	8.17	37	13.18	2.09	12.89	44
	BS	7.48	NA	13.94	-0.88	13.83	NA
	50/50	6.93	41	11.18	0.60	11.06	44
	SWITCH	9.54	33	15.08	3.71	14.54	59
<u>$m = 1.00$</u>							
	LPE	8.31	40	13.33	3.12	12.92	44
	BS	7.69	NA	13.39	-0.64	13.23	NA
	50/50	7.07	42	11.41	1.24	11.27	44
	SWITCH	9.41	32	14.14	3.36	13.58	54
<u>$m = 1.05$</u>							
	LPE	6.97	31	11.93	2.23	11.64	39
	BS	6.22	NA	12.08	-0.47	11.91	NA
	50/50	5.65	36	10.17	0.88	10.01	39
	SWITCH	8.37	21	13.64	0.74	13.52	38

Remarks: Out-of-sample hedging performance of duplicating portfolios determined by (i) locally fitting Black-Scholes hedge portfolio weights (“LPE”), (ii) Black-Scholes hedge portfolio weights based on historical volatility (“BS”), (iii) weights that are an equally weighted average of those under (i) and (ii) (“50/50”), (iv) weights that switch from those under (i) to those under (ii) when the moneyness is above 1.05 (“SWITCH”). Results are displayed for different values of the initial moneyness (m). The tracking error is defined as the payoff of a portfolio long in one call and short in the hedge portfolio. The call option has a maturity of 0.1, the underlying stock’s drift and volatility equal 0.15 and 0.40, respectively, and the short-term interest rate is set equal to 0.05. Rebalancing of the hedge portfolio occurs over intervals of length 0.01. The results are based on 500 out-of-sample stock price paths; the local parametric estimation and the volatility parameter for the Black-Scholes hedges are based on training samples which are redrawn every 50 testing samples. “%” denotes the frequency that the indicated procedure outperforms BS. “S.Dev.” denotes standard deviation.

Table 2
Performance Of Hedges Based On Local Parametric Estimation
Against Traditional Black-Scholes Hedges

	Absolute Tracking <u>Error</u>		Squared Tracking <u>Error</u> Root Mean	<u>Tracking Error</u>		
	Mean	%		Mean	Average S.Dev.	%
LPE	3.86	35	8.72	1.28	8.58	47
BS	3.73	NA	9.68	0.01	9.58	NA
50/50	3.40	44	7.96	0.64	7.87	47
SWITCH	7.04	22	12.25	0.81	12.13	51

Remarks: Out-of-sample hedging performance of duplicating portfolios determined by (i) locally fitting Black-Scholes hedge portfolio weights (“LPE”), (ii) Black-Scholes hedge portfolio weights based on historical volatility (“BS”), (iii) weights that are an equally weighted average of those under (i) and (ii) (“50/50”), (iv) weights that switch from those under (i) to those under (ii) when the moneyness is above 1.05 (“SWITCH”). The tracking error is defined as the payoff of a portfolio long in one call and short in the hedge portfolio. The call option has a maturity of 0.6, the underlying stock’s drift and volatility equal 0.15 and 0.40, respectively, and the short-term interest rate is set equal to 0.05. Rebalancing of the hedge portfolio occurs over intervals of length 0.01. The results are averaged over 500 out-of-sample stock price paths per moneyness m , where m equals 0.80, 0.85, ..., 1.05; the local parametric estimation and the volatility parameter for the Black-Scholes hedges are based on training samples which are redrawn every 50 testing samples. “%” denotes the frequency that the indicated procedure outperforms BS. “S.Dev.” denotes standard deviation.

Table 3
Performance Of Hedges Based On Local Parametric Estimation
Against Traditional Black-Scholes Hedges

		Absolute Tracking Error		Squared Tracking Error Root Mean	Tracking Error		
		Mean	%		Mean	Average S.Dev.	%
<u>$m = 0.80$</u>							
	LPE	1.62	57	4.97	0.56	4.90	24
	BS	3.03	NA	9.70	-0.28	9.62	NA
	50/50	2.31	64	6.60	0.14	6.60	24
	SWITCH	1.65	57	5.09	0.64	5.00	25
<u>$m = 0.85$</u>							
	LPE	1.12	53	4.32	0.46	4.23	23
	BS	1.39	NA	5.40	0.26	5.31	NA
	50/50	1.15	60	4.26	0.36	4.16	23
	SWITCH	1.15	53	4.38	0.53	4.29	23
<u>$m = 0.90$</u>							
	LPE	1.00	41	4.09	0.28	4.07	57
	BS	1.13	NA	4.60	-0.17	4.54	NA
	50/50	1.05	50	3.98	0.06	3.94	57
	SWITCH	1.02	41	4.10	0.31	4.07	57
<u>$m = 0.95$</u>							
	LPE	3.87	37	8.97	0.50	8.83	40
	BS	3.75	NA	9.56	-0.03	9.39	NA
	50/50	3.59	43	8.14	0.23	7.99	40
	SWITCH	4.10	36	9.09	0.96	8.94	44
<u>$m = 1.00$</u>							
	LPE	5.33	38	10.16	1.77	9.97	36
	BS	5.21	NA	11.00	-0.35	10.91	NA
	50/50	4.62	42	8.90	0.71	8.83	36
	SWITCH	6.56	33	10.80	2.79	10.38	54
<u>$m = 1.05$</u>							
	LPE	3.01	33	7.31	1.07	7.12	32
	BS	2.80	NA	7.63	-0.13	7.57	NA
	50/50	2.48	40	6.35	0.47	6.27	32
	SWITCH	4.92	13	8.35	0.94	8.18	41

Remarks: Out-of-sample hedging performance of duplicating portfolios determined by (i) locally fitting Black-Scholes hedge portfolio weights (“LPE”), (ii) Black-Scholes hedge portfolio weights based on historical volatility (“BS”), (iii) weights that are an equally weighted average of those under (i) and (ii) (“50/50”), (iv) weights that switch from those under (i) to those under (ii) when the moneyness is above 1.05 (“SWITCH”). The tracking error is defined as the payoff of a portfolio long in one call and short in the hedge portfolio. The call option has a maturity of 20 trading days, the underlying stock’s drift and volatility equal 0.10 and 0.20, respectively, and the short-term interest rate is set equal to 0.05. Rebalancing of the hedge portfolio occurs over intervals of one trading day. The results are based on 500 out-of-sample stock price paths; the local parametric estimation and the volatility parameter for the Black-Scholes hedges are based on training samples which are redrawn every 50 testing samples. “%” denotes the frequency that the indicated procedure outperforms BS. “S.Dev.” denotes standard deviation.

Figure 1: Actual call return ($\times 0.1$!) over an interval of 0.01 year, as a function of the stock return (between -0.2 and 0.2) and moneyness (between 0.8 and 1.2), for maturity fixed at 0.10 year. The Black-Scholes model is assumed, i.e., the call return is obtained from the stock return using the Black-Scholes pricing formula.

Figure 2: Actual call return ($\times 0.01!$) over an interval of 0.01 year, as a function of the stock return (between -0.2 and 0.2) and maturity (between 0.02 and 0.20), for moneyness fixed at 0.80 year. The Black-Scholes model is assumed, i.e., the call return is obtained from the stock return using the Black-Scholes pricing formula.

Figure 3: Call return according to Itô's lemma ($\times 0.1!$), as a function of the stock return (between -0.2 and 0.2) and moneyness (between 0.8 and 1.2), for maturity fixed at 0.10 year. The call return is obtained from the stock return and Itô's lemma. The latter is strictly correct only in a continuous-time diffusion world, and, hence, inaccurately predicts call price changes for noninfinitesimal stock price changes.

Figure 4: Call return according to Itô's lemma ($\times 0.01!$), as a function of the stock return (between -0.2 and 0.2) and maturity (between 0.02 and 0.20), for moneyness fixed at 0.80 year. The call return is obtained from the stock return and Itô's lemma. The latter is strictly correct only in a continuous-time diffusion world, and, hence, inaccurately predicts call price changes for noninfinitesimal stock price changes.

Figure 5: Scatter plot of estimates of the volatility parameter (“local volatility”) against moneyness for a typical training sample obtained as described in the Remarks to Table 1. The volatility is estimated as the best nonlinear local least squares fit of Black-Scholes hedge portfolio weights at the indicated moneyness (“local parametric estimate”). Bandwidth size: $2829^{-\frac{1}{5}}$.

Figure 6: Scatter plot of estimates of the volatility parameter (“local volatility”) against maturity for a typical training sample obtained as described in the Remarks to Table 1. The volatility is estimated as the best nonlinear local least squares fit of Black-Scholes hedge portfolio weights at the indicated moneyness (“local parametric estimate”). Bandwidth size: $2829^{-\frac{1}{5}}$.

Figure 7: Scatter plot of estimates of the volatility parameter (“local volatility”) against moneyness and maturity for a typical training sample obtained as described in the Remarks to Table 1. The volatility is estimated as the best nonlinear local least squares fit of Black-Scholes hedge portfolio weights at the indicated moneyness (“local parametric estimate”). Bandwidth size: $2829^{-\frac{1}{5}}$.