Dynamic Nonparametric State Price Density Estimation
Using Constrained Least Squares and the Bootstrap

Wolfgang Härdle* and Adonis Yatchew**

The economic theory of option pricing imposes constraints on the structure of call functions and state price densities (SPDs). Except in a few polar cases, it does *not* prescribe functional forms. This paper proposes a nonparametric estimator of option pricing models which incorporates various restrictions within a single least squares procedure thus permitting investigation of a wide variety of model specifications and constraints. Among these we consider monotonicity and convexity of the call function and integration to one of the state price density. The procedure easily accommodates heteroskedasticity of the residuals. Static and dynamic properties can be tested using both asymptotic and bootstrap methods. Our monte carlo simulations suggest that bootstrap confidence intervals are far superior to aymptotic ones particularly when estimating derivatives of the call function. We apply the techniques to option pricing data on the DAX.

Keywords: option pricing, state price density estimation, nonparametric least squares, bootstrap inference, monotonicity, convexity

September 7, 2001.

1.   STATE PRICE DENSITY ESTIMATION

Parametric or Nonparametric?

Option price data have characteristics which are both nonparametric and parametric in nature. The economic theory of option pricing predicts that the price of a call option should be a monotone decreasing convex function of the strike price. It also predicts that the state price density (SPD) which is proportional to the second derivative of the call function, is a valid density function over future values of the underlying asset price, and hence must be non-negative and integrate to one. Except in a few polar cases, the theory does *not* prescribe specific functional forms. (Indeed the volatility smile is an example of a clear violation of the lognormal *parametric* specification implied by Black-Scholes.) All this points to a nonparametric approach to estimation of the call function and its derivatives.

On the other hand, multiple transactions are typically observed at a finite vector of strike prices. Thus, one could argue that the model for the option price as a function of the strike price, is intrinsically parametric. Indeed given sufficient data, one can obtain a good estimate of the call function by simply taking the mean transactions price at each strike price. Unfortunately, even with large data-sets, accurate estimation of the call function at a finite number of points does not assure good estimates of its first and second derivatives, should they exist. To incorporate smoothness and curvature properties, one can select a parametric family which is differentiable in the strike price, and impose constraints on coefficients. Such an approach, however, risks specification failures.

Fortunately, nonparametric regression provides a good reservoir of candidates for flexible estimation of the call function and its derivatives. (See e.g., Aït-Sahalia and Lo (1998, 2000) and Aït-Sahalia and Duarte (2000) who use such procedures.)   In addition, there is a growing literature on the imposition of curvature properties on nonparametric estimators. (See Wright, and Wegman, (1980), Robertson, Wright

and Dykstra (1988), Mammen (1991), Mukarjee and Stern (1994) and Ramsay (1998), among others.)

In earlier work, Yatchew and Bos (1997) showed how nonparametric least squares can easily incorporate a variety of constraints such as monotonicity, concavity, additive separability, homotheticity and other implications of economic theory. Their estimator uses least squares over sets of functions bounded in Sobolev norm. Such norms provide a simple means for imposing smoothness of derivatives of various order.

In the current paper, we apply these estimators to option pricing and show how various restrictions can be incorporated within a single least squares procedure. Constraints include smoothness of various order derivatives, monotonicity and convexity of the call function and integration to one of the SPD. We construct confidence intervals and tests of various static and dynamic properties using both asymptotic and bootstrap methods. In addition to providing simulation results we apply the procedures to option data on the DAX index for the period January 4-15, 1999.

As an initial illustration of the benefits of smooth constrained estimation, particularly when estimating derivatives, we have generated 20 independent transactions prices at each of 25 strike prices. Details of the data generating mechanism are contained in Section 3 below. The top panel of Figure 1A depicts all 500 observations and the 'true' call function. Note that the variance decreases as the option price declines. The second panel depicts the estimated call function obtained by taking the mean transactions price at each of the 25 strike prices. The bottom panel depicts our smooth constrained estimate. Both estimates lie close to the true function.

Figure 1B contains estimates of the first derivative. The upper panel depicts first

order divided differences of the point means. By the mean value theorem these should provide a reasonable approximation to the true first derivative, but as can be seen, the estimate deteriorates rapidly as one moves to the left and the variance in transactions prices increases. The bottom panel depicts the first derivative of our smooth constrained estimate which by comparison is close to the first derivative of the true call function.
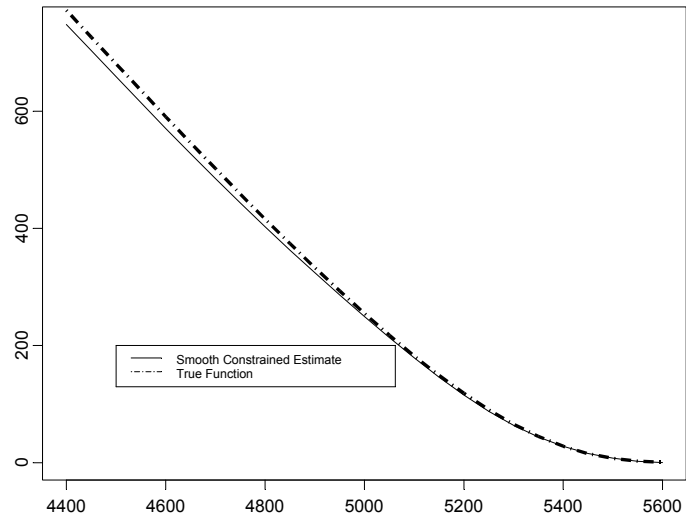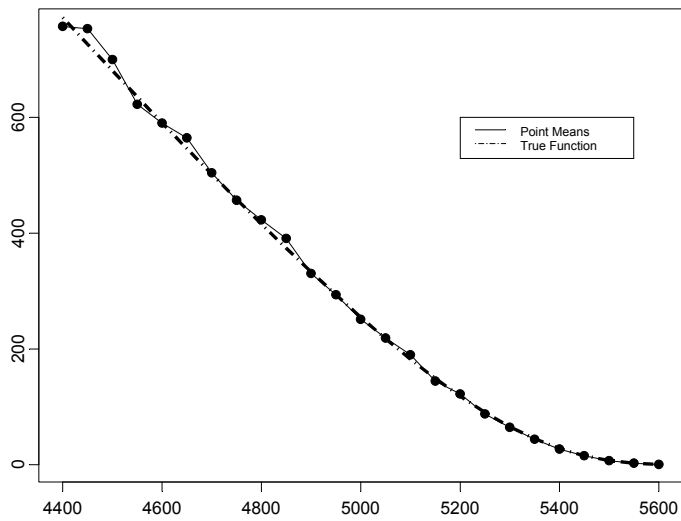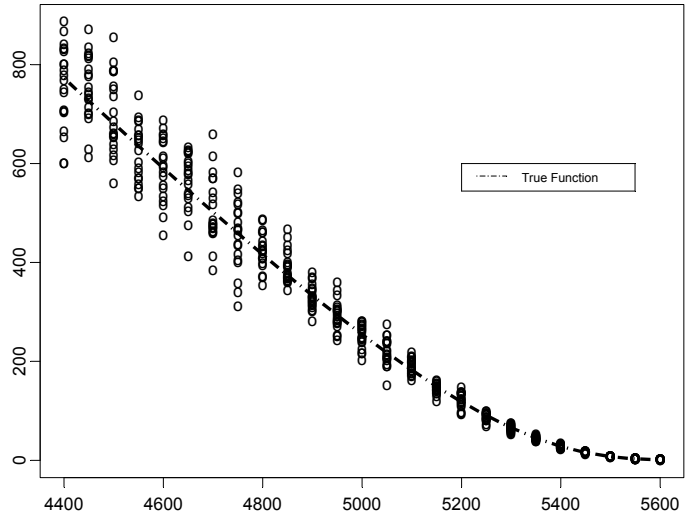
Figure 1C illustrates estimates of the second derivative of the call function. The upper panel depicts second order divided differences of the point means which gyrate wildly around the true second derivative. The lower panel depicts the second derivative of our smooth constrained estimate which tracks the true function reasonably well (note the change in scale of the vertical axis).

A number of practical advantages ensue from the procedures we propose. First, various combinations of constraints can be incorporated in a natural way within a single least squares procedure. Second, our 'smoothing' parameter has an intuitive interpretation since it measures the smoothness of the class of functions over which estimation is taking place. The measure is actually a (Sobolev) norm. If one wants to impose smoothness on higher order derivatives, this can be done by a simple modification to the norm. Third, call functions and SPD's can be estimated on an hour-by-hour, day-by-day or 'moving window' basis, and changes in shape can be tracked and tested. Fourth, our procedures readily accommodate heteroskedasticity and time series structure in the residuals. Fifth, under certain assumptions which we outline below, the *asymptotic* theory of our estimator is elementary.
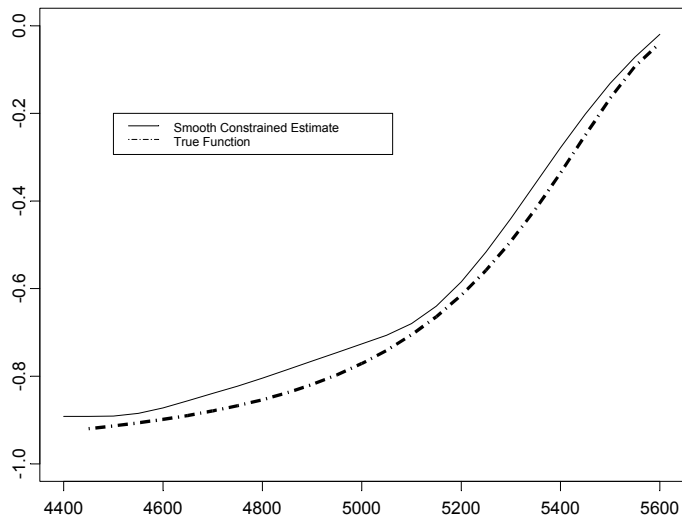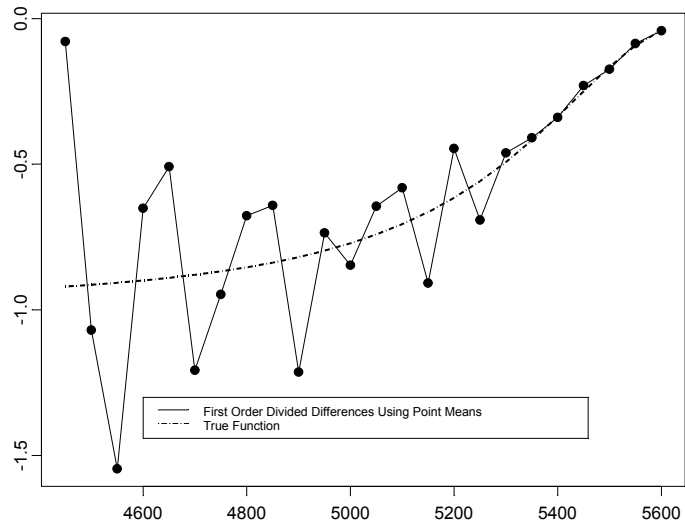
The paper is organized as follows. The remainder of this section outlines the relevant financial theory and establishes notation. Section 2 outlines the estimator as well as asymptotic and bootstrap inference procedures. Section 3 contains the results of a simple Monte Carlo study which permits us to convey some sense of the

kinds of results one might expect when the true data generating mechanism is known. Section 4 applies the techniques to data on DAX index options.
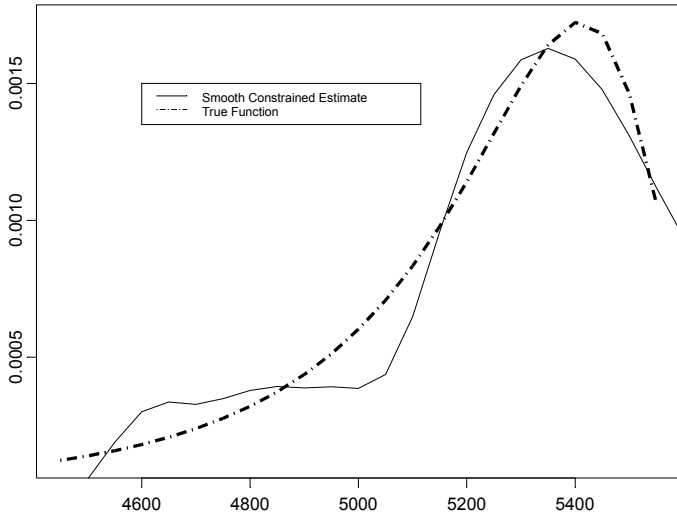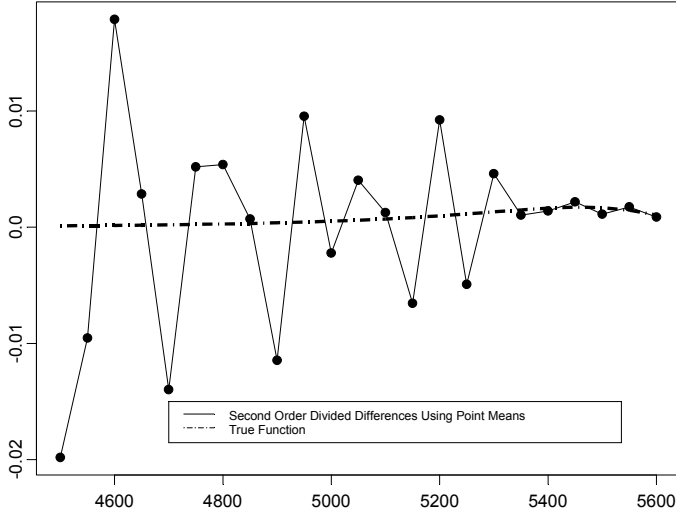
Figure 1A: Data and Estimated Call Function

Figure 1B: Estimated First Derivative

Figure 1C: Estimated SPDs

Second Order Divided Differences Using Point Means
True Function

Smooth Constrained Estimate
True Function

Before proceeding, we briefly review some of the relevant financial theory. Implicit in the prices of traded financial assets are Arrow-Debreu prices or in a continuous setting, the state price density. These are elementary building blocks for understanding markets under uncertainty. The existence and characterization of SPDs has been studied by Black and Scholes (1973), Merton (1973), Rubinstein (1976) and Lucas (1978) amongst many others. Under the assumption of no-arbitrage, the SPD is usually called the risk neutral density because if one assumes that all investors are risk neutral, then the return on all assets must equal the risk free rate of interest. Cox and Ross (1976) showed that under this assumption Black-Scholes equation follows immediately. Other approaches have been proposed by Derman and Kani (1994) and Barle and Cakici (1998).

Let X be the strike price for a call option which will expire at time T. Let t be the current time and τ the time to expiry. Let P_t and P_T denote prices of the underlying asset at times t and T respectively. Then the call pricing function at time t is given by:

$$C(X,\tau,P_t) \;=\; e^{-r\tau} \int_0^\infty \max(P_T - X, 0) f^* (P_T|P_t,\tau) dP_T \qquad (1.1)$$

where the function f^* is the state price density. It assigns probabilities to various values of the stock at time of expiration given the current stock price and the time to expiry. The density and hence the call function also depend on the current risk-free interest rate and the corresponding dividend yield of the asset. We have suppressed these in the notation. As stated earlier, the call function is monotone decreasing and convex in X.

Breeden and Litzenberger (1978) show that the second derivative of the call pricing function with respect to the strike price is related to the state price density by:

$$f^{*}\left(P_{T}|P_{t},\tau\right) \;\; = \;\; e^{r\tau}\frac{\partial^{2}C\left(X,\tau,P_{t}\right)}{\partial X^{2}} \tag{1.2}$$

Let us focus on data over a sufficiently brief time span so that we may take the time to maturity τ and the underlying stock price P_{t} as roughly constant. Our objective will be to estimate the call function $C(X)$ subject to monotonicity and convexity constraints and the constraint that the implied SPD is non-negative and integrates to one, (or at least does not exceed one over the range of observed strike prices).

We will use the following notational conventions. For an arbitrary vector z and matrices A, B we will use the usual notation z_{s}, A_{st}, $[Az]_{s}$ and $[AB]_{st}$ to denote elements. Occasionally, we will need to refer to sub-matrices of a matrix. In this case we will adopt the notation $A[a{:}b,c{:}d]$ to refer to those elements which are in rows a through b and columns c through d. Given a function f, we will denote derivatives using bracketed superscripts, e.g., $f^{(1)}$, $f^{(2)}$.

2.   Constrained Nonparametric Procedures

Nonparametric Least Squares

We begin with constrained nonparametric least squares estimation of a function of one variable. Suppose we are given data $(x_i, y_1), ..., (x_n, y_n)$ where x_i is the strike price and y_i is the option price. Let $X = (X_1, ..., X_k)$ be the vector of k distinct strike prices. (In Figure 1A, there are $k=25$ distinct strike prices with 20 observations at each price so that n-500.)

We will assume that the vector X is in increasing order. With mild abuse of notation we will use x, y and X to denote both the variable in question and the vector of observations on that variable. Our model is given by:

$$y_i = m(x_i) + \varepsilon_i \qquad i = 1, ..., n \qquad (2.1)$$

We will assume that the regression function m is four times differentiable, which in a nonparametric setting, will ensure consistent and smooth estimates of its second derivative. (Other orders of differentiation can readily be accommodated using the framework below.) The vector of distinct strike prices X lies in the interval [a,b]. (For example, if the X variable was 'moneyness' rather than the absolute strike price, then [a,b] would typically be the interval [.8, 1.2].)

We will assume the residuals ε_i are independent but possibly heteroskedastic. Let $\sigma^2(X_1), ..., \sigma^2(X_k)$ be the residual variances at each of the distinct strike prices and $\underset{n \times n}{\Sigma}$ the diagonal matrix with diagonal values $\sigma^2(x_1), ..., \sigma^2(x_n)$.

Define the following Sobolev inner product:

$$<m_1, m_2>_{Sob} = \int m_1 m_2 + m_1^{(1)} m_2^{(1)} + m_1^{(2)} m_2^{(2)} + m_1^{(3)} m_2^{(3)} + m_1^{(4)} m_2^{(4)} \qquad (2.2a)$$

with corresponding norm:

$$\| m \|_{Sob} = \left[\int m^2 + \left(m^{(1)} \right)^2 + \left(m^{(2)} \right)^2 + \left(m^{(3)} \right)^2 + \left(m^{(4)} \right)^2 \right]^{\frac{1}{2}} \qquad (2.2b)$$

where bracketed superscripts denote derivatives. Consider the following optimization problem:

$$\min_{m} \ \frac{1}{n} \sum_{i} \left[y_i - m(x_i) \right]^2 \qquad \text{s.t.} \quad \| m \|_{Sob}^2 \leq L \qquad (2.3)$$

which imposes a smoothness condition with smoothing parameter L. By varying this parameter, we control the smoothness of the ball of functions over which estimation is taking place.

Before detailing the estimator, we rewrite (2.3) to reflect the fact that option pricing data are usually characterized by repeated observations at a fixed vector of strike prices. Let B be the $n{\times}k$ matrix such that:

$$\begin{aligned} B_{ij} &= 1 \quad \textit{if } x_i = X_j \\ &= 0 \quad \textit{otherwise} \end{aligned} \qquad (2.4)$$

11

We may now rewrite (2.3) as

$$\min_{\mathbf{m}} \ \frac{1}{n} \left[\underset{n\times 1}{y} - \underset{n\times k}{B} \ \underset{k\times 1}{m(X)} \right]^{\top} \underset{n\times n}{\Sigma^{-1}} \left[\underset{n\times 1}{y} - \underset{n\times k}{B} \ \underset{k\times 1}{m(X)} \right] \quad \text{s.t.} \quad \| m \|_{Sob}^{2} \leq L \qquad (2.5)$$

Using techniques well known in the spline function literature, it can be shown that the infinite dimensional optimization problem in (2.3) or (2.5) can be replaced by a finite dimensional optimization problem as we outline below, (see e.g., Wahba (1990) or Yatchew and Bos (1997)).

Let $\mathcal{H}^{4}$ be a Sobolev space of functions from $\mathbb{R}^{1}$ to $\mathbb{R}^{1}$. Since the evaluation of functions at a specific point is a linear operator, by the Riesz Representation Theorem, given a point X_o there exists a function r_{X_o} in $\mathcal{H}^{4}$ called a *representor* such that $<r_{X_o},h>_{Sob} = h(X_o)$ for any $h \in \mathcal{H}^{4}$. Let $r_{X_1}, ..., r_{X_k}$ be the representors of function evaluation at $X_1,...,X_k$ respectively. Let R be the $k\times k$ representor matrix whose columns (and rows) equal the representors evaluated at $X_1,...,X_k$; i.e., $R_{ij} = <r_{X_i},r_{X_j}>_{Sob} = r_{X_i}(X_j) = r_{X_j}(X_i)$. (For relevant details on Sobolev spaces, calculation of representors and related concepts, see Yatchew and Box (1997).)

If one solves:

$$\min_{c} \ \frac{1}{n} \left[y - BRc \right]^{\top} \Sigma^{-1} \left[y - BRc \right] \qquad \text{s.t.} \qquad c^{\top} Rc \leq L \qquad (2.6)$$

where c is a $k\times 1$ vector, then the minimum value is equal to that obtained by solving (2.5). Furthermore, there exists a solution to (2.5) of the form:

$$\hat{m}(X) = \sum_{j=1}^{k} \hat{c}_j \, r_{X_j}(X) \tag{2.7}$$

where $\hat{c} = \left(\hat{c}_1, \ldots, \hat{c}_k\right)^\top$ solves (2.6). First and second derivatives may be estimated by differentiating (2.7):

$$\hat{m}^{(1)}(X) = \sum_{j=1}^{k} \hat{c}_j \, r_{X_j}^{(1)}(X) \tag{2.8}$$

and

$$\hat{m}^{(2)}(X) = \sum_{j=1}^{k} \hat{c}_j \, r_{X_j}^{(2)}(X) \tag{2.9}$$

Now let $X = \left(X_1, \ldots, X_k\right)$ be the vector of distinct strike prices. Define $R^{(1)}$ to be the $k \times k$ matrix whose columns (and rows) are the first derivatives of the representors evaluated at $X_1, \ldots, X_k$. Define $R^{(2)}$ in a similar fashion. Then the estimates of the call function and its first two derivatives at the vector of observed strike prices is given by $\hat{m}(X) = R\hat{c}$, $\hat{m}^{(1)}(X) = R^{(1)}\hat{c}$ and $\hat{m}^{(2)}(X) = R^{(2)}\hat{c}$.

Imposing Constraints

What is convenient about this nonparametric setting is that additional constraints on the optimization problem can be readily incorporated. Suppose one wants to impose the constraint that m is monotone decreasing at each strike price. Then one restricts the first derivative (2.8) to be negative at these points. To impose convexity, one can require the second derivative (2.9) to be positive. Then the quadratic optimization problem (2.6) can be supplemented with the monotonicity constraints:

$$R^{(1)}c \;\leq\; 0 \tag{2.10}$$

13

and the convexity constraints:

$$R^{(2)}c \geq 0 \qquad (2.11)$$

Suppose the current asset price lies in the interval $\left[X_{j^*}, X_{j^*+1}\right]$ and that one wants to impose the constraint that the state price density is unimodal with the mode in this same interval. Since the SPD is (proportional to) the second derivative $m^{(2)}$, one needs to impose constraints on *its* derivative, that is on $m^{(3)}$. Define $R^{(3)}$ to be the $k \times k$ matrix whose columns (and rows) are the third derivatives of the representors evaluated at $X_1, \ldots, X_k$. Then one imposes the constraints:

$$\underset{\bar{j} \times k}{R^{(3)}\left[1:\bar{j}, 1:k\right]} \underset{k \times 1}{c} \geq 0 \qquad \underset{(k-\bar{j}) \times k}{R^{(3)}\left[\bar{j}+1:k, \ 1:k\right]} \underset{k \times 1}{c} \leq 0 \qquad (2.12)$$

The first set of $\bar{j}$ inequalities ensures that the SPD has a positive derivative at strike prices below the current asset price, the remaining $k - \bar{j}$ inequalities provide for a negative derivative at higher strike prices.

Finally, to impose the constraint that every valid density function must satisfy, that is $\int m^{(2)} = 1$ one inserts:

$$\sum_{j=1}^{k} c_j \left[r_{X_j}^{(1)}(b) - r_{X_j}^{(1)}(a)\right] = 1 \qquad (2.13)$$

Asymptotic Distribution of the Estimator

The estimator we propose in equation (2.7) is nonparametric in the sense that it is the closest function in the infinite-dimensional ball $\Im = \left\{ m \in \mathcal{H}^4 : \| m \|_{Sob}^2 \leq L \right\}$ to the data. On the other hand, if there is a fixed number of strike prices k, then the model is parametric in the sense that the unknown parameter vector c is finite.

Assumptions: i) Let $\pi_j > 0$ be the proportion of observations at strike price X_j. ii) Assume that the true call function m lies strictly inside the ball of functions $\Im$, i.e., $\| m \|_{Sob}^2 < L$. iii). Assume also that m is strictly monotone decreasing and strictly convex.

Assumption i) ensures there are ample observations at each strike price so that the option price can be estimated consistently without reference to data at other strike prices. We could allow the proportions to be functions of n so long as they do not vanish as n increases. Assumptions ii) and iii) greatly simplify asymptotic distribution and bootstrap results because they imply that the true function is strictly in the interior of the set of smooth, monotone and convex functions over which estimation is taking place.

Proposition 1: Let $\bar{y}(X) = \bar{y}(X_1,..,X_k)$ be the k-dimensional vector of average transactions prices at the k strike prices. Let $\hat{c}$ minimize (2.6) with the added constraints (2.10) and (2.11) and define $\hat{m}(X) = \hat{m}(X_1,..,X_k) = R\hat{c}$. Then $\lim_{n \to \infty} Prob\left[\hat{m}(X) = \bar{y}(X) \right] = 1$. ∎

Essentially, this proposition states that as data accumulate at each strike price, the inequalities implied by the smoothness constraint, and the monotonicity and convexity constraints eventually become non-binding and the estimator becomes

15

identical to the point mean estimator.

Proposition 2: Define $\hat{c}$ and $\hat{m}(X)$ as in Proposition 1. Let $\hat{m}^{(1)}(X) = R^{(1)}\hat{c}$ and $\hat{m}^{(2)}(X) = R^{(2)}\hat{c}$. Let $\Omega / n = Var(\hat{y}(X))$ be the $k \times k$ diagonal matrix of variances of the point mean estimators, i.e., $\Omega_{jj} = \sigma^2(X_j) / \pi_j$. Then

$$n^{\frac{1}{2}}(\hat{c} - c) \overset{D}{\to} N(0, R^{-1}\Omega R^{-1}) \tag{2.14a}$$

Furthermore,

$$n^{\frac{1}{2}}(\hat{m}(X) - m(X)) \overset{D}{\to} N(0, \Omega)$$

$$n^{\frac{1}{2}}(\hat{m}^{(1)}(X) - m^{(1)}(X)) \overset{D}{\to} N(0, R^{(1)}R^{-1}\Omega R^{-1}R^{(1)}) \tag{2.14b}$$

$$n^{\frac{1}{2}}(\hat{m}^{(2)}(X) - m^{(2)}(X)) \overset{D}{\to} N(0, R^{(2)}R^{-1}\Omega R^{-1}R^{(2)}) \quad \blacksquare$$

Proposition 2 provides for asymptotic scalar and vector confidence regions of the call function, its first derivative and the SPD. For example, if one is interested in confidence intervals at strike price X_s, the asymptotic pivots are:

$$\tau^{(0)} = n^{\frac{1}{2}} \frac{(\hat{m}(X_s) - m(X_s))}{[\hat{\Omega}]^{\frac{1}{2}}_{ss}} \tag{2.15a}$$

$$\tau^{(1)} = n^{\frac{1}{2}} \frac{(\hat{m}^{(1)}(X_s) - m^{(1)}(X_s))}{[R^{(1)}R^{-1}\hat{\Omega} R^{-1}R^{(1)}]^{\frac{1}{2}}_{ss}} \tag{2.15b}$$

16

$$\tau^{(2)} \quad = \quad n^{\frac{1}{2}} \frac{\left(\hat{m}^{(2)}\left(X_s\right) - m^{(2)}\left(X_s\right)\right)}{\left[R^{(2)}R^{-1}\hat{\Omega}R^{-1}R^{(2)}\right]_{ss}^{\frac{1}{2}}} \qquad (2.15c)$$

To estimate $\sigma^2\left(X_j\right)$ assemble the estimated residuals $\left\{\hat{\varepsilon}_{j1},...,\hat{\varepsilon}_{jn_j}\right\}$ at strike price X_j and calculate the usual mean squared deviation:

$$\hat{\sigma}^2\left(X_j\right) = \frac{1}{n_j}\sum_s^{n_j} \hat{\varepsilon}_{js}^2 - \left(\frac{1}{n_j}\sum_s^{n_j} \hat{\varepsilon}_{js}\right)^2 \qquad (2.16)$$

from which one immediately obtains an estimate of Ω, say $\hat{\Omega}$. Proposition 2 is also useful for establishing the asymptotic validity of bootstrap confidence sets.

Bootstrap Procedures

Percentile and percentile-*t* procedures are commonly used for constructing confidence intervals. The latter are often found to be more accurate when the statistic is an asymptotic pivot (see Hall (1992) for extensive arguments in support of this proposition). On the other hand, percentile methods might be better when the asymptotic approximation to the distribution of the pivot is poor as a result of small sample size or slow convergence.

In our case we are especially interested in confidence intervals for the call function and for the SPD. Our heuristic example depicted in Figures 1A,B,C suggests that the asymptotic approximation should be adequate in estimating the call function but poor when estimating its second derivative.

We continue by outlining a bootstrap procedure for obtaining confidence intervals

for the call function and its derivatives at one of the observed strike prices, say X_o. As there are multiple observations at each strike price, we can accommodate heteroskedasticity by resampling from the estimated residuals at each strike price or we can use the wild bootstrap (see Wu (1986) or Härdle (1990, p.106-8, 247)). Figures 2 contains the algorithm for constructing percentile confidence intervals for the call function and its first two derivatives at a confidence level of 95%. The procedures are applicable with the obvious modifications for a general confidence level α. (Algorithms for constructing percentile-t confidence intervals may be constructed with modest additional effort.)

Proposition 3: Suppose one constructs an α-confidence interval for $m(X_s)$, $m^{(1)}(X_s)$ or $m^{(2)}(X_s)$ using the bootstrap procedure in Figures 2. Then the probability that it contains the corresponding true value converges to α as $n \to \infty$. ∎

Let L_{CV} be the cross-validation estimate of the Sobolev bound L. Then,

Proposition 4: $L_{CV} \overset{P}{\to} c'Rc = m(X)'R^{-1}m(X) \le \|m\|_{Sob}^2$. ∎

To test stability, we will use conventional parametric and bootstrap tests on $\hat{c}$.

18

Figure 2: Percentile Confidence Intervals for $m(X_s)$, $m^{(1)}(X_s)$ **and** $m^{(2)}(X_s)$

1. Calculate $\hat{c}$ and $\hat{m}$ by solving (2.6) subject to (2.10) and (2.11). Calculate the estimated residuals $\{\hat{\varepsilon}_1, ..., \hat{\varepsilon}_n\}$.

2. a) Construct a bootstrap data-set $(x_i, y_1^B), ..., (x_n, y_n^B)$ where $y_i^B = \hat{m}(x_i) + \varepsilon_i^B$ and ε_i^B is obtained by sampling from $\{\hat{\varepsilon}_1, ..., \hat{\varepsilon}_n\}$ using the wild bootstrap.

   b) Using the bootstrap data-set obtain $\hat{c}^B$ by solving (2.6) subject to (2.10) and (2.11). Calculate and save $\hat{m}^B(X_s)$, $\hat{m}^{(1)^B}(X_s)$ and $\hat{m}^{(2)^B}(X_s)$.

3. Repeat steps 2 multiple times.

4. To obtain a 95% point-wise confidence intervals for $m(X_s)$, $m^{(1)}(X_s)$ and $m^{(2)}(X_s)$ obtain .025 and .975 quantiles of the corresponding bootstrap estimates.

3. Monte Carlo Results

Effects of Constraints

To perform their simulations, Aït-Sahalia and Duarte (2000) calibrate their model using characteristics of the S&P options market. We calibrate our experiments using DAX options in January 1999 which expire in March. At that time the DAX index was in the vicinity of 5000 (see Section 4. Below). The 25 distinct strike prices range over the interval 4400 to 5600 in 50 unit increments. The short term interest rate is set to 3.5% and the dividend yield to 2%. We assume the volatility smile function is linear in the strike price. We have already seen the 'true' call function, its first derivative and SPD plotted in Figures 1A,B, and C above. In each of the simulations below we assume there are 20 observations at each of the 25 strike prices for a total of 500 observations. At each strike price, the residual variance is set to about 10% of the option price.

Figure 3A, B and C illustrate the impact of constraints on estimation. The 'unconstrained' estimator consists of the point means at each strike price. The 'smooth' estimator imposes only the Sobolev constraint as in (2.5) and (2.6) with the degree of smoothness identical to the true Sobolev norm of the underlying function which is the square root of 1.812. Monotonicity and convexity constraints are imposed using equations (2.10) and (2.11). Finally, we impose 'unimodality' constraints which requires the estimated SPD to be non-decreasing over the lowest five strike prices and non-increasing over the highest five. Its purpose is to improve the estimator of the SPD at the boundaries.

As is evident from Figure 3A, the improvement in estimation of the call function resulting from adding constraints is barely discernible, though it is present as may be seen from the MSE's in the first column of Table 1.

Figure 3B illustrates the impacts on estimation of the first derivative. They are

clearly evident (note that the vertical scale narrows as one moves down the figure).

The most dramatic impact of the constraints is on estimation of the SPD as may be seen in Figure 3C. Smoothness alone produces a very broad band of estimates, so much so that the true SPD looks quite flat. Adding monotonicity and convexity improves the estimates substantially, though they are quite imprecise at low strike prices. This is in part due to the much larger variance there. The 'unimodality' constraints alleviate this problem. Average mean squared errors are summarized in Table 1 for the various estimators.

Table 1: Average MSE			
Model	Call Function	First Derivative	SPD
Unconstrained	200.02	.1219	1.49×10^{-4}
Smooth	25.68	.0052	1.26×10^{-6}
Smooth, Monotone, Convex	13.70	.0015	3.76×10^{-7}
Smooth, Monotone, Convex, Unimodal	13.26	.00097	2.00×10^{-7}
The unconstrained estimator consists of the point means at each strike price.			

Confidence Intervals

Next we turn to confidence interval estimation. Asymptotic confidence intervals can be estimated using Proposition 2. For bootstrap intervals we use the percentile method outlined in Figure 2. In each case we performed 100 bootstrap draws. Figure 4 contrasts asymptotic and bootstrap confidence intervals for our data generating mechanism. (Recall there are 25 distinct strike prices with 20

21

observations at each one.) In the top panel which corresponds to the call function, the asymptotic confidence intervals (dashed lines) are somewhat broader than the bootstrap intervals (dotted lines). The middle and lower panels correspond to the first derivative and SPD estimates. In these two cases, confidence intervals based on the asymptotic approximation are extremely poor relative to the bootstrap intervals.

Next we turn to the accuracy of the bootstrap intervals. We produced 500 samples and in each case obtained 100 bootstrap re-samples. The model which was estimated was the fully constrained version with smoothness, monotonicity, convexity and unimodality constraints. The smoothness constraint was set at the true norm of the function. Figure 5 plots the point-wise coverage probabilities for confidence intervals which were nominally set at 99%, 95%, 90% and 50%. The top panel corresponds to the call function itself. Bootstrap intervals are reasonably accurate except at strike prices between 4500 and 4600DM. The averages of the observed coverage frequencies across strike prices is .9795 for 99% confidence intervals, .94 for 95%, .899 for 90% and .5185 for 50% intervals.

The middle panel illustrates observed bootstrap coverage frequencies for the first derivative. The averages of the observed coverage frequencies across strike prices are .9835 for 99% confidence intervals, .9545 for 95%, .9095 for 90% and .4935 for 50%.

The bottom panel provides similar plots for the SPD. The averages of the observed coverage frequencies across strike prices is .994 for 99% confidence intervals, .978 for 95%, .953 for 90% and .575 for 50%. Thus, percentile bootstrap confidence intervals for the SPD in the fully constrained model are conservative.

Overall we found that while MSE improves and bootstrap confidence intervals narrow as one adds constraints to the 'smooth' model, bootstrap coverage accuracy

22

deteriorates moderately.

Cross-Validation

Finally, we performed simulations in which the Sobolev smoothness parameter was estimated as the minimum of the cross-validation function. We found that even with much smaller samples, e.g., with two observations at each of the 25 strike prices, the estimated cross-validation parameter was close to the true value.

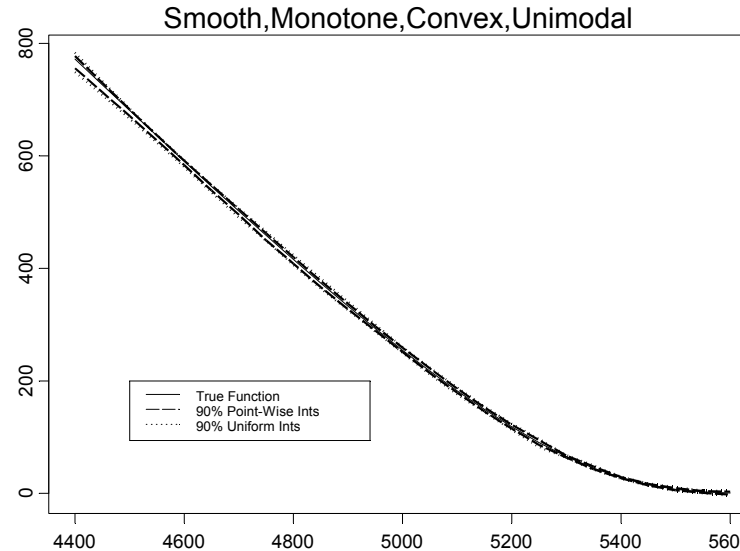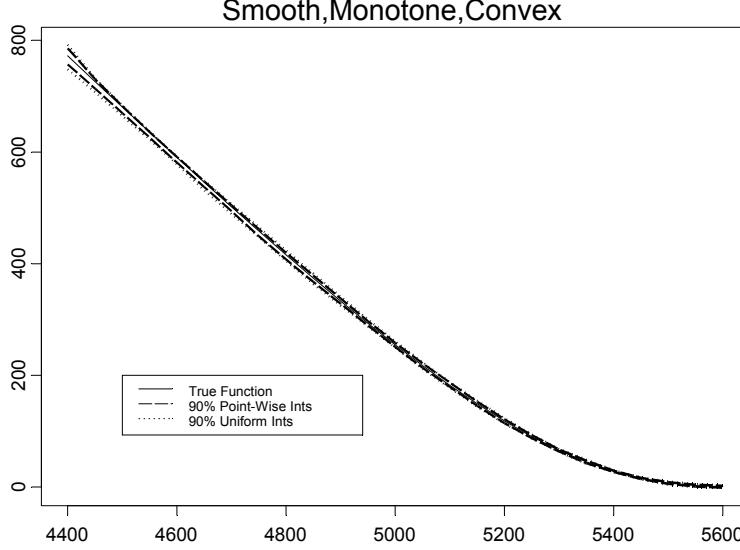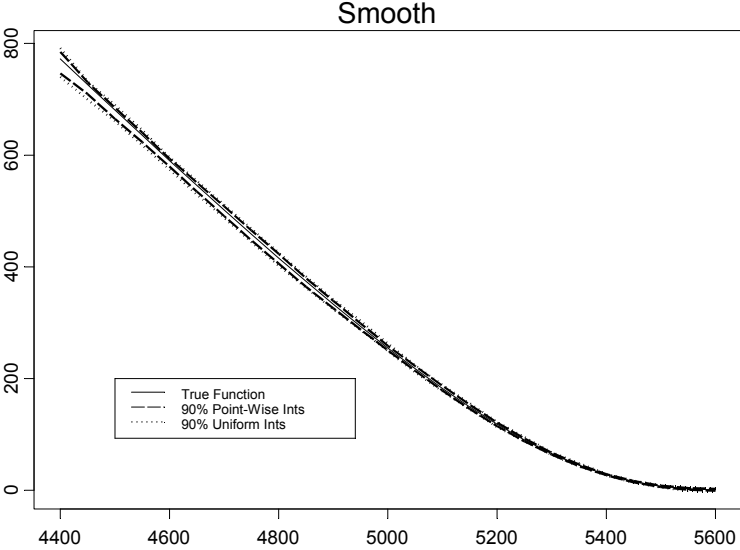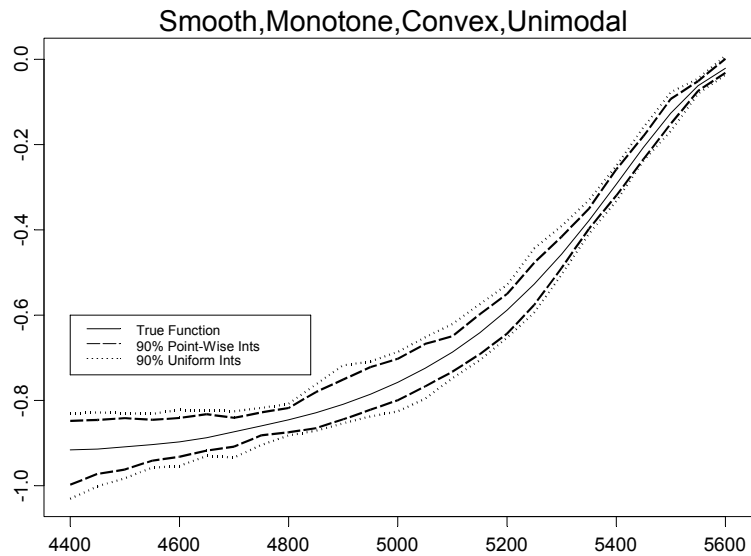**Figure 3A: Effects of Constraints on Call Function Estimates**

Smooth

Smooth,Monotone,Convex

Smooth,Monotone,Convex,Unimodal

Figure 3B: Effects of Constraints on First Derivative Estimates

Smooth

Smooth,Monotone,Convex

Smooth,Monotone,Convex,Unimodal
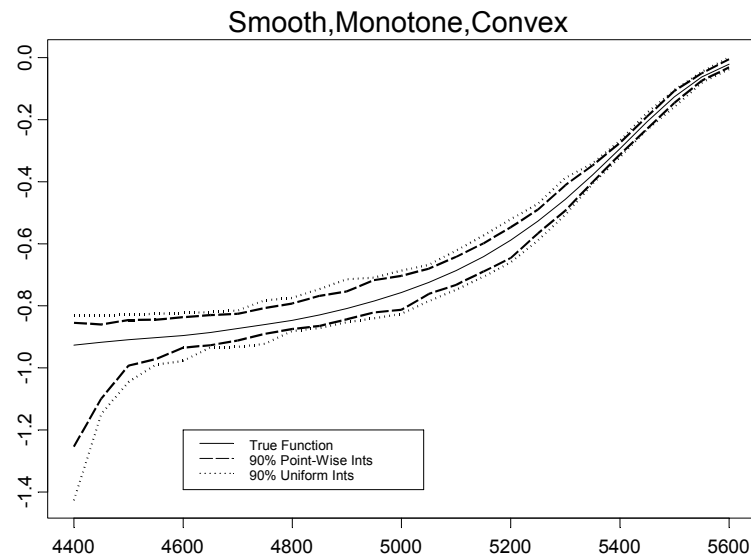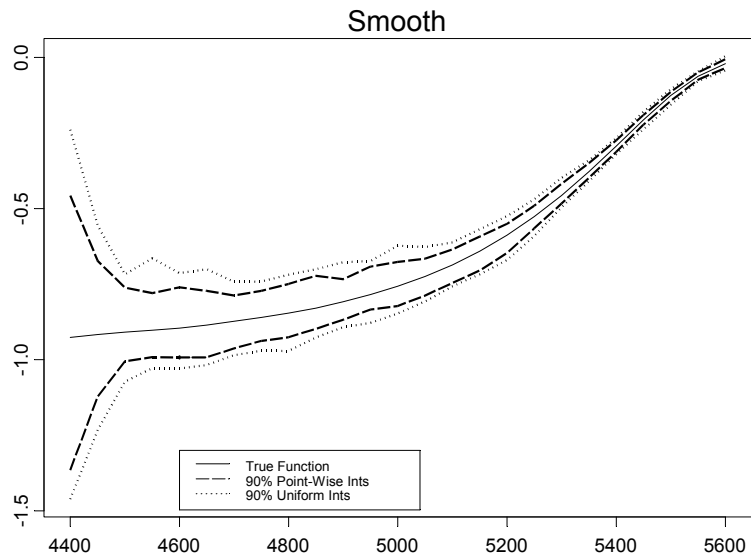
Figure 3C: Effects of Constraints on SPD Estimates

Smooth

Smooth,Monotone,Convex

Smooth,Monotone,Convex,Unimodal
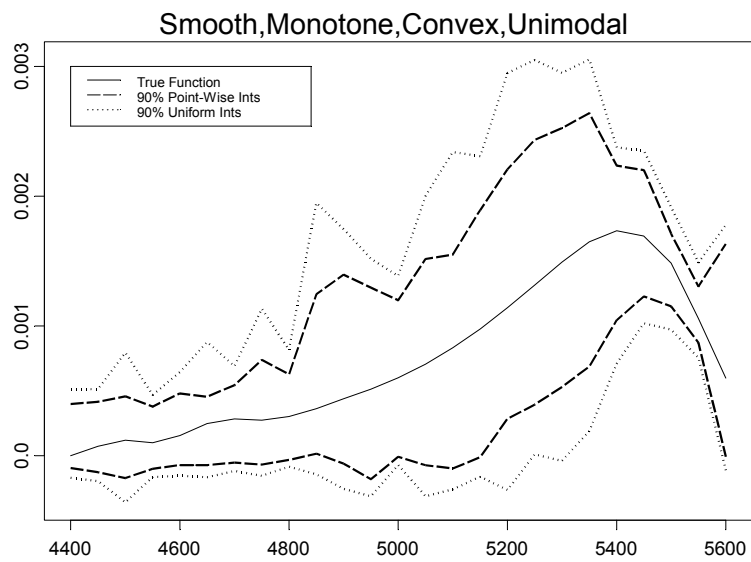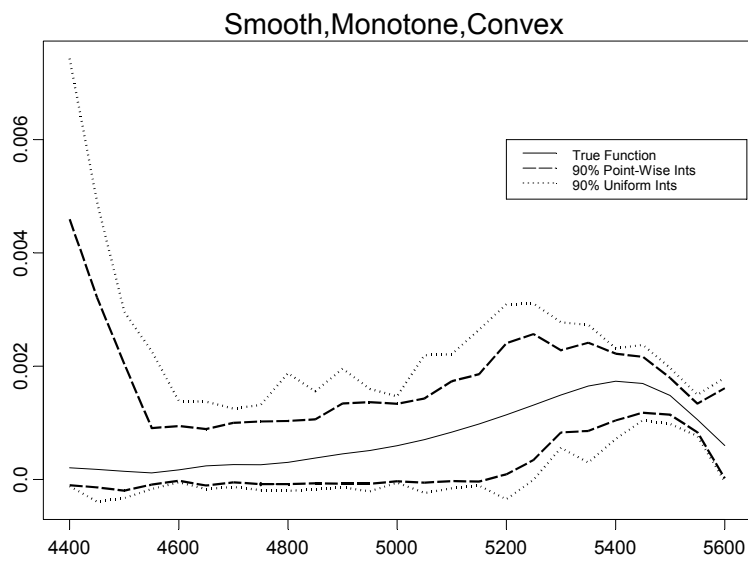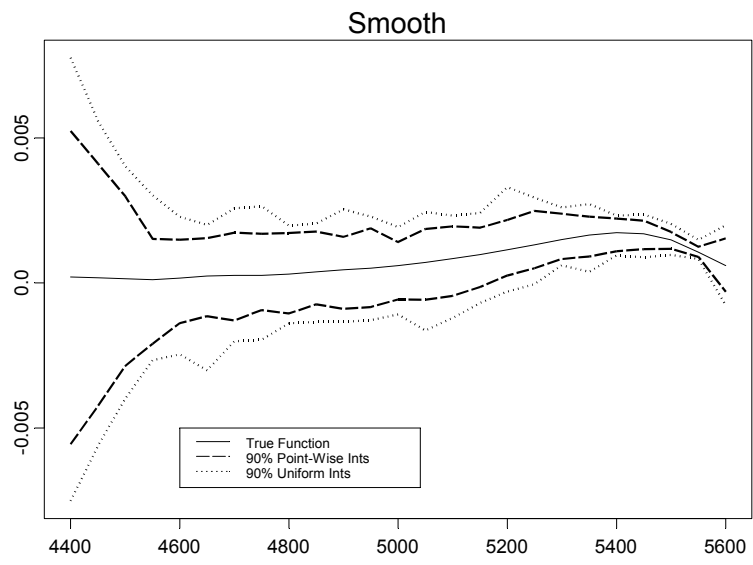
Figure 4: Asymptotic Vs. Bootstrap Confidence Intervals
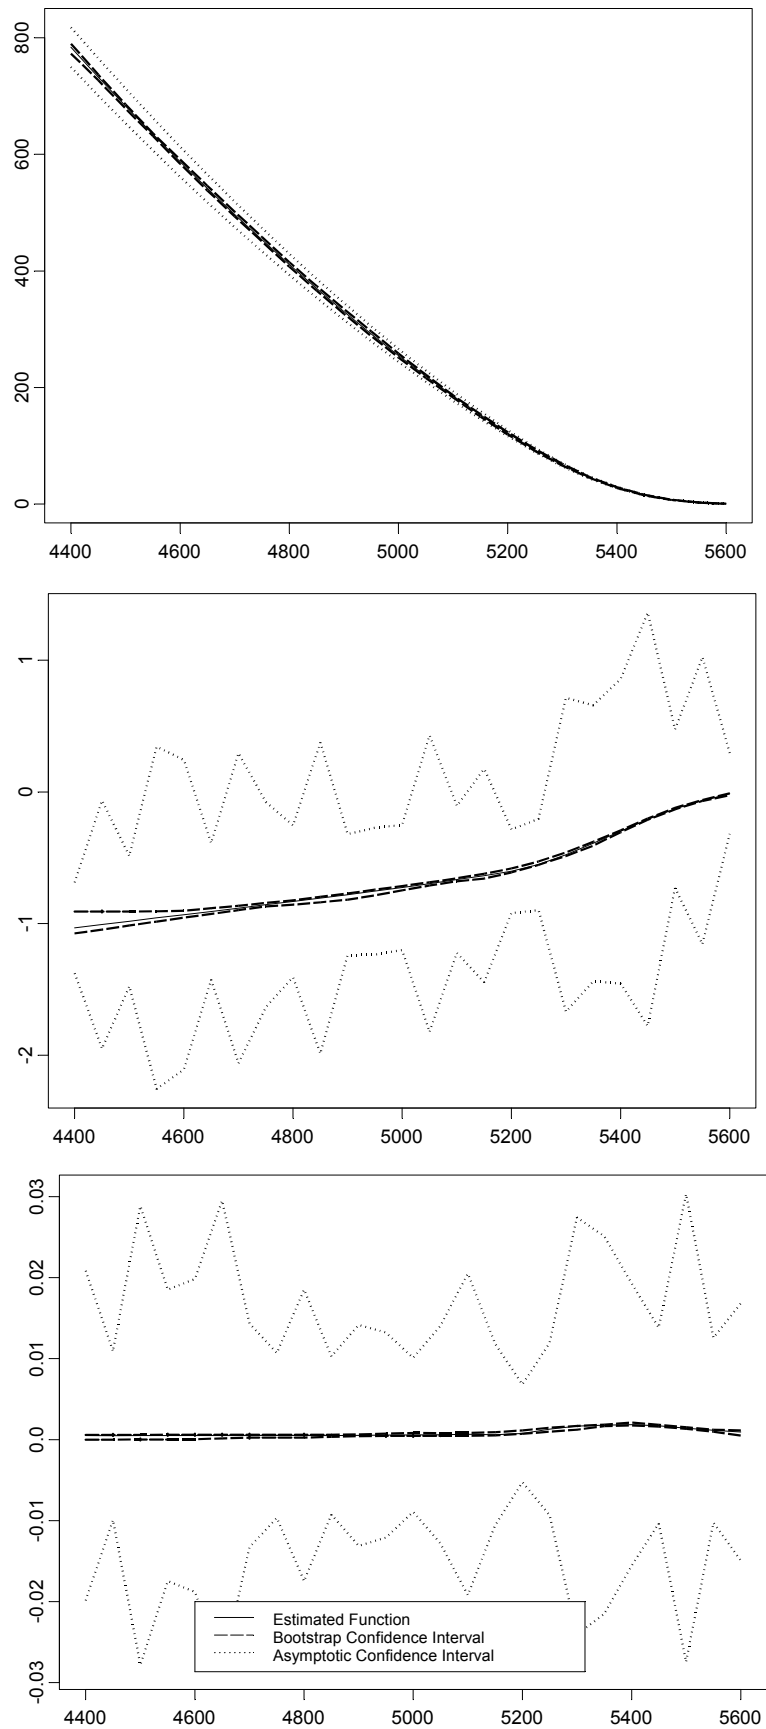
Figure 5: Accuracy of Bootstrap Confidence Intervals
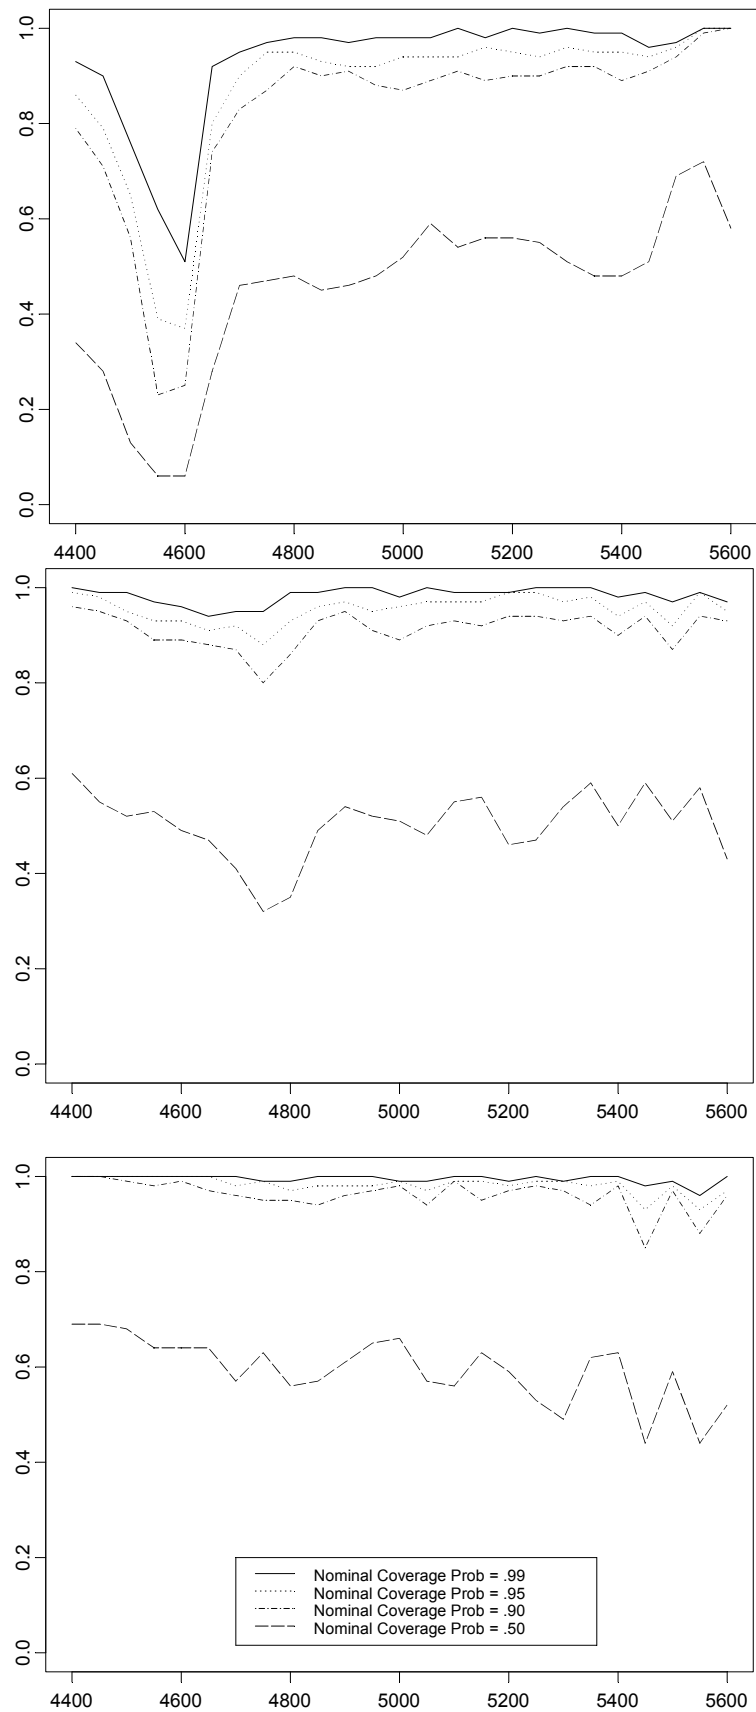
4. Applications to DAX Index Option Data

In this section we use the tools we have described to analyze data on DAX index options over the two week period January 4-15, 1999. Figure 6 illustrates the closing daily values of the DAX over the period. During the first week, the index fluctuates in a range above 5200. In the early part of the second week it begins to decline and during the last three days of the sample period, the index remains below 5000.

Figure 6: DAX Jan 4-15, 1999

	4-Jan-99	5-Jan-99	6-Jan-99	7-Jan-99	8-Jan-99	11-Jan-99	12-Jan-99	13-Jan-99	14-Jan-99	15-Jan-99
Series1	5290	5263	5443	5346	5371	5267	5196	4982	4903	4974
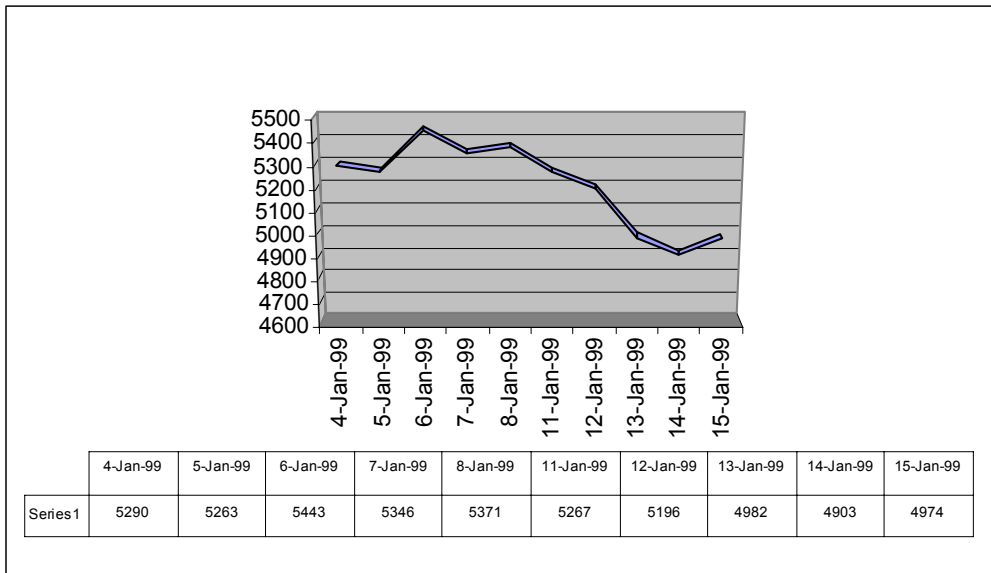
Our estimates will be restricted to call options which trade at 100 point intervals between 4500 and 5500 and expire in 1 month. Some trades do indeed take place outside this range, but there are few of them. We begin with data for January 4. Figure 7A illustrates the estimated call function, its first derivative and the SPD along with 95% point-wise bootstrap confidence intervals. In addition we have included uniform confidence bounds for the SPD based on the Bonferroni inequality. With the DAX closing at 5290, one can reasonably expect that there is some

probability mass beyond the 5500 level. Options at 5500 averaged 22DM suggesting that the market assigned a positive probability to the DAX exceeding 5500 at time of expiration of the options. Indeed our estimated SPD integrates to about .8 as may be inferred from the middle panel of Figure 7A.

Figure 7B depicts the estimated call function, its first derivative and SPD for January 14. By this time, the DAX had dropped to about 4900. The SPD is zero at both end points of our estimation range and constraint (2.13) which requires that the integral of the SPD not exceed one is binding and hence informative to the estimation process.

A test of equality of SPD's across these two days would clearly fail. A more interesting question is whether on two consecutive days when there is little change in the index, one would conclude that the market has significantly changed its probabilistic view of where the market will end up at some future date. We consider this question by testing equality of the SPD's for January 4[th] and 5[th] where there has been little change in the DAX. As we have indicated earlier, the test can be conducted by testing equality of the coefficient vectors $\hat{c}$. We use the bootstrap to obtain critical values. Figure 8 plots the SPD's for each day along with uniform confidence bounds which is strongly suggestive of rejection. The statistic yields a value of ___ which compares to a critical value of ___.

Figure 9 plots two perspectives on the procession of daily estimated SPD's over the full period January 4-18, 1999. During the first week expectations are that the DAX will be higher at time of expiry. By the end of the second week, expectations have moved sharply lower, consistent with the DAX index value, and the SPD's become more highly concentrated, possibly a result of the fact that the expiry date is now closer. The pattern may be seen perhaps more clearly in the Appendix which contains graphs of the estimated call functions, their first derivatives and the SPDs for each day individually.
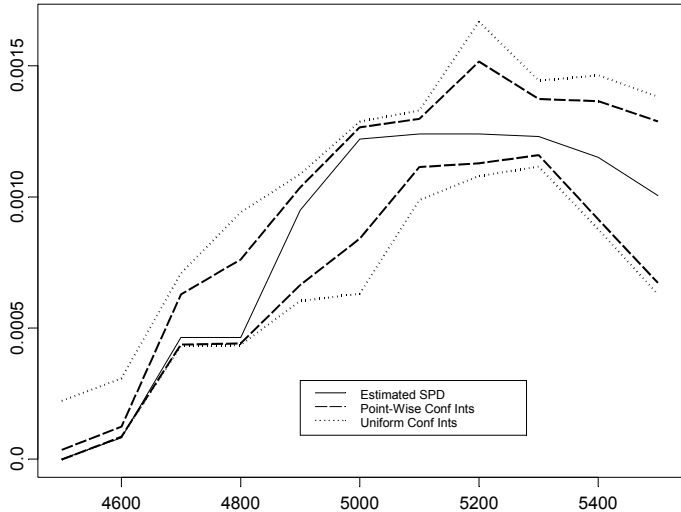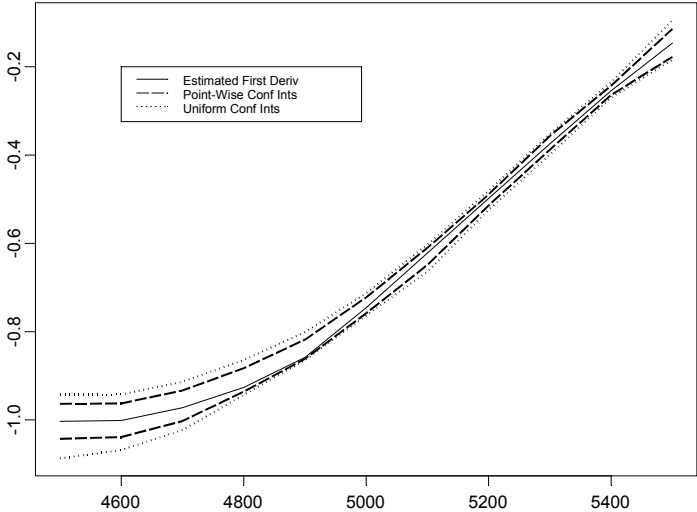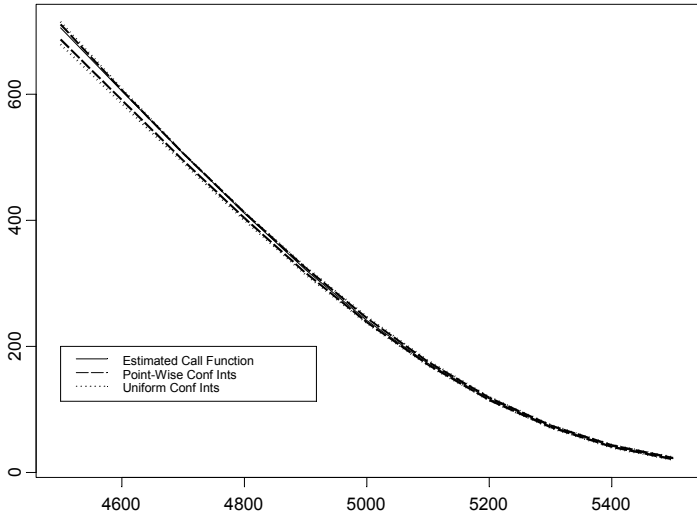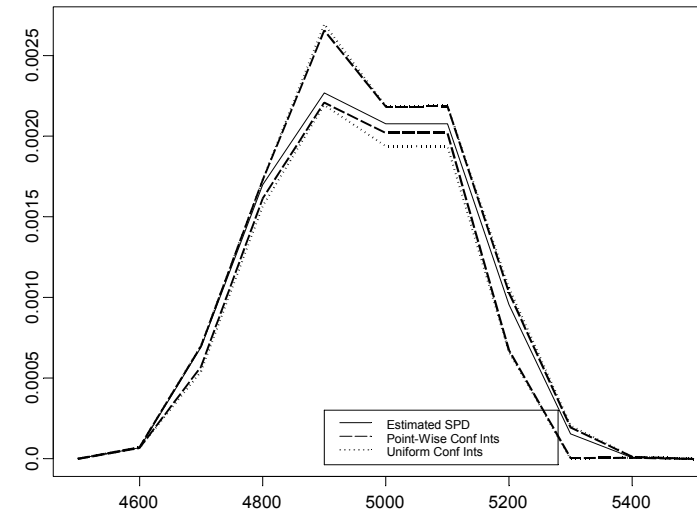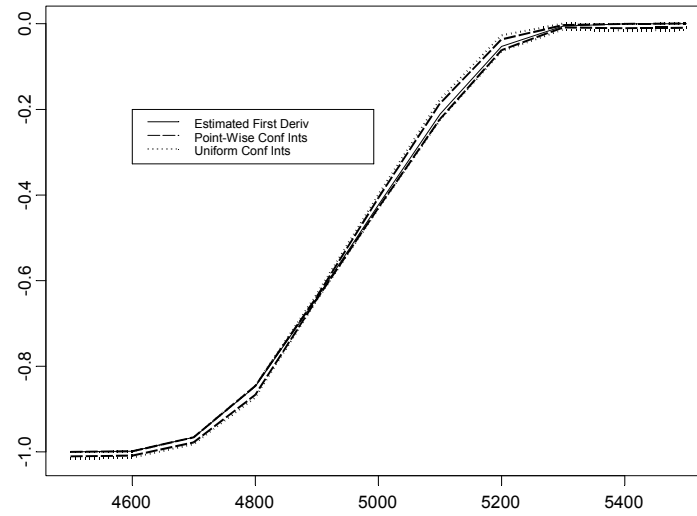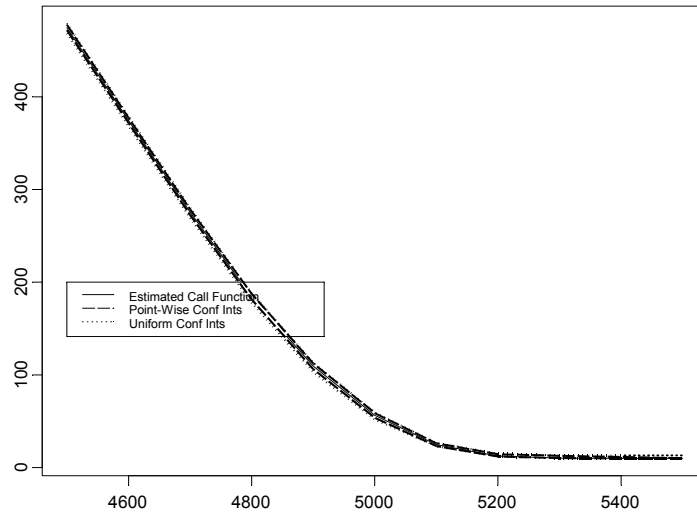
Figure 7A: January 4, 1999
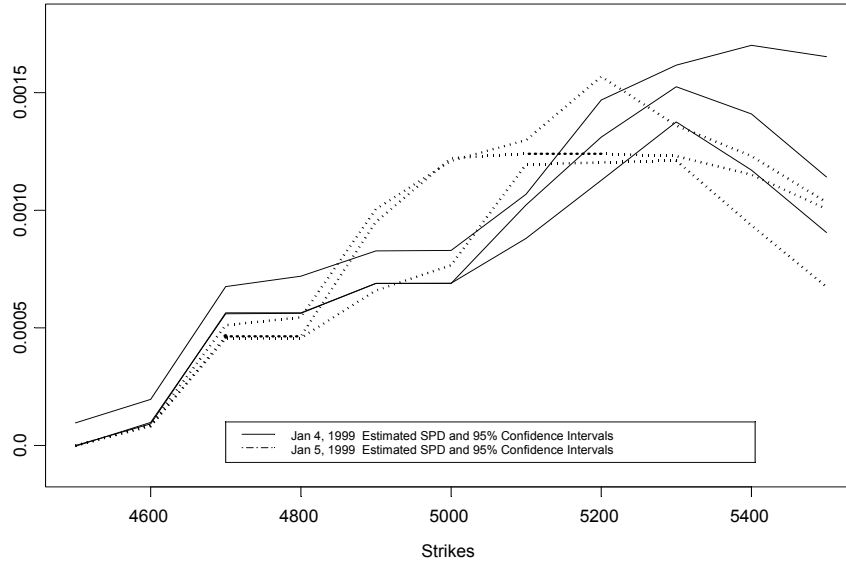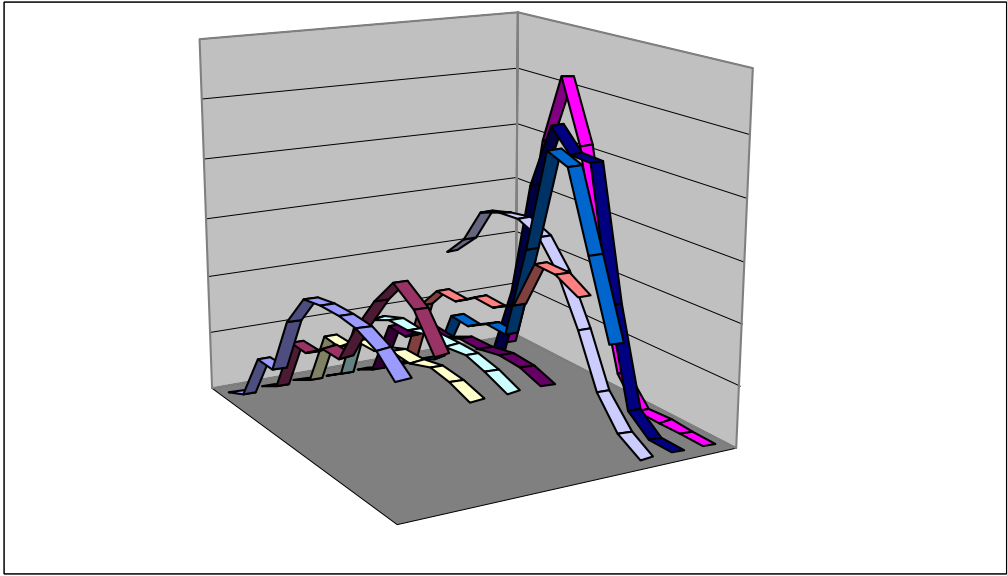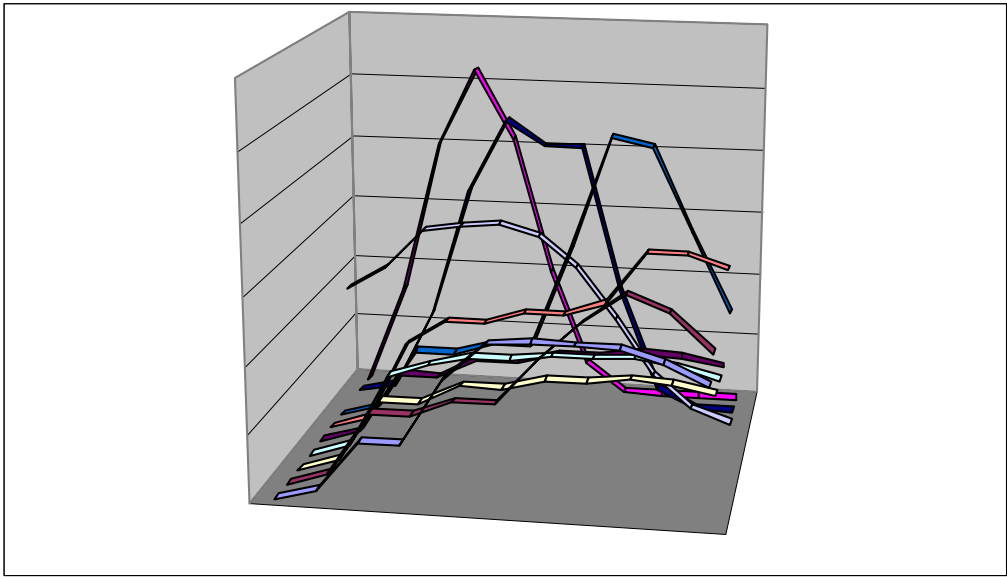
Figure 7B: January 14, 1999

Figure 8: January 4 and 5, 1999

Jan 4, 1999 Estimated SPD and 95% Confidence Intervals
Jan 5, 1999 Estimated SPD and 95% Confidence Intervals

Strikes

REFERENCES

Aït-Sahalia, Y. and A. Lo (1998): "Nonparametric Estimation of State-Price Densities Implicit in Financial Asset Prices", *Journal of Finance*, 53, 499-547.

Aït-Sahalia, Y. and A. Lo (2000): "Nonparametric Risk Management and Implied Risk Aversion", *Journal of Econometrics*, 94, 9-51.

Aït-Sahalia, Y. and J. Duarte (2000): "Nonparametric Option Pricing Under Shape Restrictions", manuscript, Princeton University.

Barle S. and N. Cakici (1998) "How to grow a smiling tree", *Journal Financial Engineering*, 7, 127 - 146,

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

Breeden, D. and R. Litzenberger (1978): "Prices of State-Contingent Claims Implicit in Option Prices", *Journal of Business*, 51, 621-651.

Cox, J. and S. Ross (1976): "The Valuation of Options for Alternative Stochastic Processes", *Journal of Financial Economics*, 3, 145-166.

Derman, E. and I. Kani (1994): "Riding on the Smile", *Risk*, 7, 32-39.

Hall, Peter, (1992): *The Bootstrap and Edgeworth Expansion*, Springer Verlag, New York

Härdle W. (1990): *Applied Nonparametric Regression*, Econometric Society Monograph Series, 19, Cambridge University Press.

Lucas, R. (1978): "Asset Prices in an Exchange Economy", *Econometrica*, 46, 1429-1446.

Mammen, E. (1991): "Estimating a Smooth Monotone Regression Function", *Annals of Statistics*, 19, 724-740. Mukarjee, Hari (1988): "Monotone Nonparametric Regression", *Annals of Statistics*, 16, 741-750.

Merton, R. (1973): "Rational Theory of Option Pricing", *Bell Journal of Economics and Management Science*, 4, 141-183.

Mukarjee, H. and S. Stern (1994): "Feasible Nonparametric Estimation of Multi-argument Monotone Functions", *Journal of the American Statistical Association*, 89, 77-80.

Ramsay, J. O. (1998): "Estimating Smooth Monotone Functions", *Journal of the Royal Statistical Society, Series B*, 60, 365-375.

Robertson, T., F.T. Wright and R.L. Dykstra (1988): *Order Restricted Statistical Inference,* New York: John Wiley and Sons.

Rubinstein, M. (1976): "The Valuation of Uncertain Income Streams and the Pricing of Options", *Bell Journal of Economics*, 7, 407-425.

Wright, I. and E. Wegman (1980): "Isotonic, Convex and Related Splines", *Annals of Statistics*", 8, 1023-1035.

Wu, Chien-Fu Jeff (1986): "Jackknife, Bootstrap and Other Resampling Methods in Regression Analysis" (with discussion), *Annals of Statistics,* 14, 1261-1350.

Yatchew, A.J. and Len Bos (1997) "Nonparametric Regression and Testing in Economic Models", *Journal of Quantitative Economics*, Vol. 13, 81-131.

Call Function, First Derivative and SPD January 4-8, 1999.
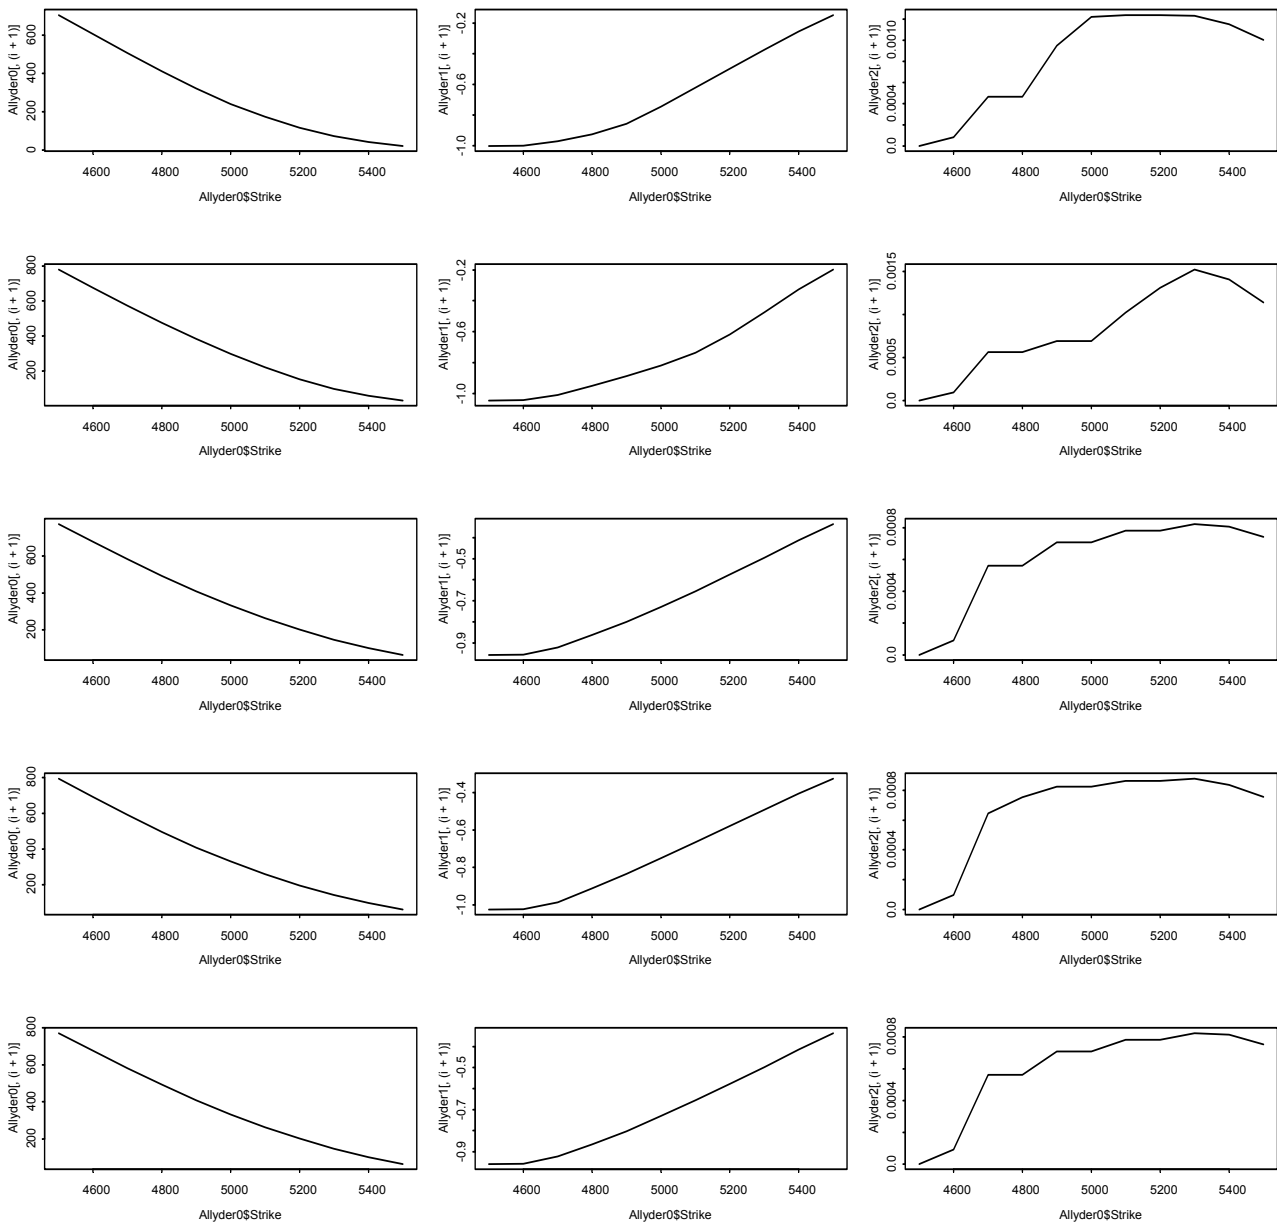
Call Function, First Derivative and SPD January 11-15, 1999.