
Nonparametric Estimation of Risk-Neutral Densities

Maria Grith², Wolfgang Karl Härdle¹ and Melanie Schienle³

¹ Professor at Humboldt-Universität zu Berlin, Ladislaus von Bortkiewicz Chair of Statistics and Director of CASE - Center for Applied Statistics and Economics, Spandauer Straße 1, 10178 Berlin, Germany and National Central University, Department of Finance, No. 300, Jhongda Rd., Jhongli City, Taoyuan County 32001, Taiwan, (R.O.C.), Email: haerdle@wiwi.hu-berlin.de

² Research Associate at Humboldt-Universität zu Berlin, Ladislaus von Bortkiewicz Chair of Statistics, Spandauer Straße 1, 10178 Berlin, Germany, Email: gritmari@wiwi.hu-berlin.de

³ Professor at Humboldt-Universität zu Berlin, Chair of Econometrics and CASE - Center for Applied Statistics and Economics, Spandauer Straße 1, 10178 Berlin Germany, Email: melanie.schienle@wiwi.hu-berlin.de

Summary. This chapter deals with nonparametric estimation of the risk neutral density. We present three different approaches which do not require parametric functional assumptions on the underlying asset price dynamics nor on the distributional form of the risk neutral density. The first estimator is a kernel smoother of the second derivative of call prices, while the second procedure applies kernel type smoothing in the implied volatility domain. In the conceptually different third approach we assume the existence of a stochastic discount factor (pricing kernel) which establishes the risk neutral density conditional on the physical measure of the underlying asset. Via direct series type estimation of the pricing kernel we can derive an estimate of the risk neutral density by solving a constrained optimization problem. The methods are compared using European call option prices. The focus of the presentation is on practical aspects such as appropriate choice of smoothing parameters in order to facilitate the application of the techniques.

Keywords: Risk neutral density; Pricing kernel; Kernel smoothing; Local polynomials; Series methods

1 Introduction

Most of our economic understanding of investment under uncertainty is based on pure Arrow-Debreu Securities (Arrow 1964, Debreu, 1959), which pay one unit of currency at the end of a period if a state of nature is realized and zero

otherwise. Their theoretical state-contingent prices are the starting point for pricing any security in an economic equilibrium under uncertainty. In a continuum of states, the prices of the Arrow-Debreu securities are characterized by the state-price density, which yields one dollar if the final state is in the interval $[x, x+dx]$ when starting from any point x . While one way to justify existence and form of a state-price density is from preference-based equilibrium models (Lucas, 1978), we focus here on the reasoning from arbitrage-based models (Merton, 1973). In these models the state-price density is called risk neutral density (RND) under the assumption that the underlying market is dynamically complete, which we will adopt in the following. The RND also uniquely characterizes the equivalent martingale measure under which all asset prices discounted at the risk-free rate are martingales.

Under the above stated assumption, the price of a European call is obtained by discounting the expected payoff, where the expectation is taken with respect to the risk neutral measure

$$C(X, \tau, r_{t,\tau}, \delta_{t,\tau}, S_t) = e^{-r_{t,\tau}\tau} \int_0^\infty (S_T - X)^+ q(S_T|\tau, r_{t,\tau}, \delta_{t,\tau}, S_t) dS_T. \quad (1)$$

Here S_t is the underlying asset price at time t , X the strike price, τ the time to maturity, $T = t + \tau$ the expiration date, $r_{t,\tau}$ the deterministic risk free interest rate for that maturity T , $\delta_{t,\tau}$ corresponding dividend yield of the asset, $q(S_T|\tau, r_{t,\tau}, \delta_{t,\tau}, S_t)$ is the conditional risk neutral density. We assume that these state variables contain all essential information needed for estimation of C and q while quantities such as stochastic market volatility, trading volumes, ask-bid spreads are negligible. We write $q(S_T)$ instead of $q(S_T|\cdot)$ to keep notation simple. The risk neutral density can be derived from (1) as

$$q(S_T) = e^{r_{t,\tau}\tau} \left\{ \frac{\partial^2 C}{\partial X^2} \right\}_{X=S_T}. \quad (2)$$

The relation is due to Breeden and Litzenberger (1978) and is the starting point for two possible estimation strategies: first, estimate a continuous twice-differentiable call function in all its arguments from traded options by smoothing in the call price or secondly, by smoothing in the implied volatility space. In addition to these standard approaches we also introduce a third indirect way via estimation of the empirical pricing kernel. Assuming that all the variables other than X are fixed, the price of the European call option with strike price X expiring in τ years under the historical measure p is given by

$$\begin{aligned} C(X) &= e^{-r_{t,\tau}\tau} \int_0^\infty (S_T - X)^+ \frac{q(S_T)}{p(S_T)} p(S_T) dS_T \\ &= e^{-r_{t,\tau}\tau} \int_0^\infty (S_T - X)^+ m(S_T) p(S_T) dS_T, \end{aligned} \quad (3)$$

where p is the subjective density of the stock price at the expiration of the option, at time T and m is the so called pricing kernel characterizing the change of measure from q to p .

In standard option pricing models such as Merton (1976), Heston (1993) or Bates (1996), estimation of the risk neutral density crucially depends on underlying model assumptions such as the underlying asset price dynamics and the statistical family of distributions that the risk neutral density is assumed to belong to. Consumption based asset pricing models prespecify the preference of the representative agent and condition therefore the shape of the pricing kernel (Lucas, 1978 and Rubinstein, 1976). Recent empirical findings, however, question the validity of these popular specifications which drive the overall result (Campbell, Lo and McKinlay, 1997). Nonparametric estimation offers an alternative by avoiding possibly biased parametric restrictions and therefore reducing the respective misspecification risk. Since nonparametric estimation techniques require larger sample sizes for the same accuracy as a parametric estimation procedure, increasing availability of large data sets as intraday traded option prices have raised their feasibility. On the other hand, due to their flexibility, many existing nonparametric risk neutral density estimation techniques are afflicted by irregularities such as data sparsity in the tails, negative probabilities and integrability to one. We will address these problems by appropriate choices of bandwidths and kernels, suggesting semiparametric techniques or imposing relevant constraints.

We use kernel based methods (local polynomial regression) in each of the first two cases. These methods are flexible and yield point estimates as opposed to series expansion, sieve methods or splines. Though shape constraints such as convexity or monotonicity of the call price are hard to incorporate directly in the estimation step. Therefore in particular in small samples they are not satisfied leading to problems with economic theory. Thus in the estimation via EPK in Section 3 we use series approximation for directly controlling constraints in the estimation.

The rest of the chapter is organized as follows: section 2 describes kernel based regression methods for direct estimation of the RND from the call prices function, section 3 introduces the concept of pricing kernel and explains the indirect method of estimating RND, section 4 concludes. Throughout the chapter empirical studies using EUREX DAX Index based European option data illustrate the methods, comparing their performance.

2 Estimation of RND based on the second derivative

The standard approach for a nonparametric estimator for the risk neutral density is by estimating the second derivative of the call price with respect to the strike X . Then an estimate for q is obtained by discounting according to (2). Therefore in the following we focus on estimation of the second derivative of a regression function.

Call the d -dimensional vector of covariates \mathbf{Z} which comprises all estimation relevant variables of $(X, \tau, r_{t,\tau}, \delta_{t,\tau}, S_t)$ from (1) and denote call prices as response Y . From paired observations Y_i and $\mathbf{Z}_i = (Z_{ik})_{k=1}^d$, for $i = 1, \dots, n$ we want to estimate the following general, possibly nonlinear relationship

$$Y_i = C(\mathbf{Z}_i) + \varepsilon_i, \quad i = 1, \dots, n \quad (4)$$

where $C(\cdot) : \mathbb{R}^d \rightarrow \mathbb{R}$ is a smooth function in all d dimensions and ε_i is i.i.d. with $E[\varepsilon_i | \mathbf{Z}_i] = 0$.

Kernel based methods are local techniques for estimating the function C at any point \mathbf{z} in its domain; they use a weighted average of the Y_i -s to yield fitted values via

$$\hat{C}(\mathbf{z}) = \sum_{i=1}^n w_i(\mathbf{z}) Y_i \quad (5)$$

where the weights $w_i(\mathbf{z})$ assigned to each point of fit \mathbf{z} decline for increasing distance $|\mathbf{Z}_i - \mathbf{z}|$ and $\frac{1}{n} \sum_{i=1}^n w_i(\mathbf{z}) = 1$. Kernel regression methods use kernel functions to construct weights. A univariate kernel is a smooth, symmetric real-valued squared integrable function $K(u) : \mathbb{R} \rightarrow \mathbb{R}$ which integrates to one. We can think of a standard kernel function as a probability density with potentially compact support. Examples of such K are presented in Table 1.

Table 1. Kernel functions $K(u)$

Uniform	$\frac{1}{2} \mathbf{I}(u \leq 1)$
Triangle	$(1 - u) \mathbf{I}(u \leq 1)$
Epanechnikov	$\frac{3}{4} (1 - u^2) \mathbf{I}(u \leq 1)$
Quartic (Biweight)	$\frac{15}{16} (1 - u^2)^2 \mathbf{I}(u \leq 1)$
Triweight	$\frac{35}{32} (1 - u^2)^3 \mathbf{I}(u \leq 1)$
Gaussian	$\frac{1}{\sqrt{2\pi}} \exp(-\frac{1}{2}u^2)$
Cosine	$\frac{\pi}{4} \cos(\frac{\pi}{2}u) \mathbf{I}(u \leq 1)$

Furthermore, there exist more general types of kernel functions, so called higher order kernels, which can be used for bias refinements in the estimation, see Subsection 2.1. The order of a kernel $v > 0$ is defined as the first nonzero moment of the kernel, that is

$$\int u^l K(u) du = 0, \quad l = 1, \dots, v - 1 \quad (6)$$

$$\int u^v K(u) du = \kappa_v \neq 0$$

and $\kappa_v < \infty$. Solving the system of equations (6) for kernel functions integrating to unity for a fixed v , yields a v -th order kernel. The larger v , however,

the more "wiggly" the resulting kernel becomes - covering more and more negative areas. Here we mostly consider standard second order kernels, which are nonnegative functions.

Set $K_h(u) = \frac{1}{h} K\left(\frac{u}{h}\right)$ for all $u \in \mathbb{R}$ where h is the bandwidth, the smoothing parameter. In a d -dimensional space, for each pair \mathbf{z} and \mathbf{Z}_i the multivariate kernel function $\mathcal{K}(\mathbf{z} - \mathbf{Z}_i) : \mathbb{R}^d \rightarrow \mathbb{R}$ must analogously fulfil

$$\mathcal{K}_H(\mathbf{z} - \mathbf{Z}_i) = \frac{1}{\det(H)} \mathcal{K}\{H^{-1}(\mathbf{z} - \mathbf{Z}_i)\}$$

where $H = \text{diag}(\tilde{h})$ is the diagonal matrix of bandwidths $\tilde{h} = [h_1, \dots, h_d]$. The matrix H can in general also contain off-diagonal elements - but in practice such generality is not needed. Define the multidimensional kernel $\mathcal{K}_H(\mathbf{z} - \mathbf{Z}_i)$ as a product of univariate kernels

$$\mathcal{K}_H(\mathbf{z} - \mathbf{Z}_i) = \prod_{k=1}^d K\left(\frac{z_k - Z_{ik}}{h_k}\right).$$

For expositional simplicity we let $h_1 = \dots = h_d = h$. Details on how to choose the optimal bandwidths are addressed in the next section.

The simplest case of choosing w_i in (5) is to use Nadaraya-Watson weights

$$w_i(\mathbf{z}) = \frac{\mathcal{K}_h(\mathbf{z} - \mathbf{Z}_i)}{\sum_{i=1}^n \mathcal{K}_h(\mathbf{z} - \mathbf{Z}_i)}.$$

These are a special constant case of general local polynomial weights derived below. Besides other choices of weights such as in the k -nearest neighbour or the Gasser-Müller estimator are possible. Estimators of the second derivative of a function are constructed by twice differentiating the estimator of the function. In this case these estimators, however, have inferior statistical properties and are therefore not included here. References are e.g. Härdle et al. (2004) and Fan and Gijbels (1996). We focus on local polynomial estimation which directly yields estimates of derivatives. The idea of local polynomial regression is based on Taylor expansion approximating an unknown function C at a point \mathbf{z} . In order to keep notation simple we will first illustrate the method for the univariate case. The multivariate case is systematically the same and will be sketched afterwards.

Locally any sufficiently smooth function C can be approximated by a polynomial of degree p

$$\begin{aligned} C(Z_i) &= \sum_{j=0}^p \frac{C^{(j)}(z)}{j!} (Z_i - z)^j + \mathcal{O}\{(Z_i - z)^{p+1}\} \\ &= \sum_{j=0}^p \beta_j (Z_i - z)^j + \mathcal{O}\{(Z_i - z)^{p+1}\} \end{aligned} \tag{7}$$

with $i = 1, \dots, n$. Therefore by minimizing a locally weighted least squares regression

$$\min_{\beta} \sum_{i=1}^n \left\{ Y_i - \sum_{j=0}^p \beta_j(z)(z - Z_i)^j \right\}^2 \mathcal{K}_h(z - Z_i). \quad (8)$$

the solution $\hat{\beta}_0(z)$ provides an estimator of C at point z , while $j! \hat{\beta}_j(z)$, with $j = 1, \dots, p$ are the estimated derivatives at that point. Closed forms for $\hat{\beta}(z) = (\hat{\beta}_0(z), \dots, \hat{\beta}_p(z))$ can be obtained by solving (8) via equating the corresponding system of first order conditions to zero. As we are interested in an estimator for the second derivative of a function, we should choose $p \geq 2$. As will be outlined in the subsection below, for good statistical properties without requiring too much smoothness $p = 3$ will be a suitable choice.

In d -dimensional case, expansion (7) will include mixed terms which must be appropriately ordered. Then the interpretation of the coefficients is similar: $\hat{\beta}_0(\mathbf{z})$ is the estimator of C at point \mathbf{z} , while $j! \hat{\beta}_j(\mathbf{z}) = j! [\beta_{j_1}(\mathbf{z}), \dots, \beta_{j_\mu}(\mathbf{z})]$ with $\mu = 1, \dots, N_j$ is N_j -dimensional vector of j -th order derivatives of C evaluated at point \mathbf{z} . It is obvious that $N_0 = 1$ (β_0 is the local constant) and $N_1 = d$ (β_1 is the vector of partial derivatives) but for $j \geq 2$ the expansion contains cross order derivatives and the general formula for N_j is

$$N_j = \binom{d+j-1}{j-1} = \frac{(d+j-1)!}{(j-1)!}.$$

For example, when $j = 2$ we have $N_2 = d(d+1)/2$ distinct derivatives and, $\nabla^{(2)} \hat{C}(\mathbf{z}) = 2\hat{\beta}_2(\mathbf{z})$ is the estimate of

$$\nabla^{(2)} C(\mathbf{z}) = \begin{pmatrix} \frac{\partial^2 C(\mathbf{z})}{\partial z_1^2} \\ \frac{\partial^2 C(\mathbf{z})}{\partial z_1 \partial z_2} \\ \vdots \\ \frac{\partial^2 C(\mathbf{z})}{\partial z_d^2} \end{pmatrix}.$$

For the estimation of the RND we are interested in the second derivative of the call price with respect to the strike price X . In our notation with $\mathbf{z} = (X, \tau, r_{t,\tau}, \delta_{t,\tau}, S_t)$ this is $2\beta_{21}$. Thus

$$\hat{q}(S_T) = 2e^{r_{t,\tau}} \hat{\beta}_{21}(S_T, \mathbf{z}_{-1}) = e^{r_{t,\tau}} \left\{ \frac{\partial^2 \widehat{C}(\mathbf{z})}{\partial z_1^2} \right\}_{X=S_T}$$

with $\mathbf{z}_{-1} = (\tau, r_{t,\tau}, \delta_{t,\tau}, S_t)$.

2.1 Statistical Properties

Assume for simplicity that C is univariate and has continuous derivatives of total order $(p+1)$. The probability density function f_Z of (Z_1, \dots, Z_n) is

continuous and $(p + 1)$ times continuously differentiable ($f_Z \geq 0$) and also $K(u)$ is a bounded second order kernel having compact support. Let $\hat{C}^{(j)}$ denote the estimator of $C^{(j)}$ based on a p -th order local polynomial fit ($j \leq p$). The results below are standard and can be found for instance in Li and Racine (2007).

Theorem 1. *When $p - j$ is odd, the bias is*

$$\mathbb{E} \left[\hat{C}^{(j)}(z) \right] - C^{(j)}(z) = h^{p-j+1} c_{1,j,p} \left\{ \frac{\omega^{(p+1)}(z)}{(p+1)!} \right\} + o(h^{p-l+1}). \quad (9)$$

When $p - j$ is even, the bias is

$$\begin{aligned} \mathbb{E} \left[\hat{C}^{(j)}(z) \right] - C^{(j)}(z) &= h^{p-j+2} c_{2,j,p} \left\{ \frac{\omega^{(p+2)}(z)}{(p+2)!} \right\} \int u^{p+2} K(u) du \quad (10) \\ &+ h^{p-j+2} c_{3,j,p} \left\{ \frac{\omega^{(p+1)}(z) f^{(1)}(z)}{f(z)(p+1)!} \right\} \end{aligned}$$

where $\omega(z) = \left\{ \hat{C}(z) - C(z) \right\} f(z)$. In either case

$$\text{Var} \left(\hat{C}^{(j)}(z) \right) = \left\{ \frac{c_{4,j,p} \sigma^2(z)}{nh^{2j+1}} \right\} + o \left\{ (nh^{2j+1})^{-1} \right\}. \quad (11)$$

The exact form of the constants $c_{a,j,p}$ for $a = 1, 2, 3, 4$ can be found in Ruppert and Wand (1994).

For illustration consider the special case $p = 0$ and $j = 0$ - local constant estimation of a function. The bias is

$$\frac{h^2}{2} \left\{ C^{(2)}(z) + 2 \frac{C^{(1)}(z) f_X^{(1)}(z)}{f_Z(z)} \right\} \mu_2(K), \quad (12)$$

with $\mu_2(K) = \int u^2 K(u) du$. For $p = 1$ and $j = 0$ - local linear estimation of a function - the bias becomes

$$\frac{h^2}{2} \left\{ C^{(2)}(z) \right\} \mu_2(K). \quad (13)$$

Observe in general from (9) and (10) that the bias for $p - j$ even contains an additional design dependent term with factor $\frac{f_X^{(1)}(z)}{f_Z(z)}$ as opposed to the odd case. Sign and size of this quantity, however, depend on the shape of the underlying estimated function and the shape of f_Z . In particular at the boundary of the support of Z small values of f_Z inflate the entire term. Therefore odd values of $p - j$ are preferable avoiding such boundary bias problems and pertaining the same variance. In our case, we are interested in the second derivative. We therefore choose the design adaptive order $p = 3$ and not $p = 2$ according to Theorem 1.

With higher order kernels (6) of order v and corresponding higher smoothness assumptions the bias in Theorem 1 can be further reduced to be of order h^v for fixed p and j with $v > p - j + 2$ without changing the rate in the variance. In practice the order v , however, cannot be chosen too large as with increasing v the estimates have robustness problems in finite samples due to negative weights associated with the kernels (Müller, 1988).

Observe from Theorem 1 that kernel estimation of a derivative is harder than of the function itself. While the variance in the function estimation decreases with $\mathcal{O}(1/(nh))$ the corresponding rate in the second derivative is only $\mathcal{O}_P(1/(nh^5))$ which is much slower. Therefore the finite sample performance of second derivatives lacks the precision of the fit achieved for the function.

Rates of convergence can be judged according to the MSE. Assuming that $p - j$ is odd, it is

$$\begin{aligned} \text{MSE}(z, h, j) &= \mathbf{E} \left[\hat{C}^{(j)}(z) - C^{(j)}(z) \right]^2 \\ &= \mathcal{O} \left\{ \underbrace{h^{2(p-j+1)}}_{\text{bias}^2} + \underbrace{(nh^{2j+1})^{-1}}_{\text{var}} \right\}. \end{aligned} \quad (14)$$

For constructing confidence intervals of the nonparametric estimates use the following normality result

Theorem 2. *Under some additional moment assumptions it is for $p - j$ odd*

$$\sqrt{nh^{2j+1}} \{ \hat{C}^{(j)}(z) - C^{(j)}(z) \} \rightarrow \mathbf{N}(0, V_j) \quad (15)$$

with V_j as in Theorem 1.

For a precise statement of the standard moment conditions see Li and Racine (2007). Analogous results to Theorem 1 and 2 hold for d -dimensional functions. The only remarkable systematic difference is that the dimension of the regressors enters in the rate of the variance which is then $\mathcal{O}_P\{(nh^{2j+d})^{-1}\}$. Likewise the rate in the CLT also deteriorates with d and is nh^{2j+d} . This phenomenon is known in the literature as the curse of dimensionality capturing the fact that finite sample performance of nonparametric estimators decreases with an increasing number of regressors. Therefore in practice appropriate semiparametric dimension reduction techniques are used. They keep high modeling flexibility but yield better finite sample properties in regression settings with more than three regressors. See Subsection 2.3 for details.

2.2 Selection of smoothing parameter

Most important for good nonparametric estimation results is an appropriate choice of bandwidth. Other parameters like the selection of a suitable kernel K only have little influence on the final result in practice. This is because

kernel functions can be rescaled such that the difference between two kernel estimates using two different kernels is almost negligible (Marron and Nolan, 1988). For the choice of order p of the employed local polynomial estimator it is sufficient to follow the logic outlined above.




Fig. 1. MASE (simple line), squared bias (dashed line) and variance part (dotted line) for simulated data, weights $w(x) = I(x \in [0.05, 0.95])$

An optimal bandwidth should minimize both bias and variance of the estimator. Though according to Theorem 1 there is a tradeoff between these quantities as smaller bandwidths would reduce the bias but inflate the variance. Therefore selecting h by minimizing $MSE(\mathbf{z}, h, j)$ (14) balances bias and variance (see Figure 1 for an illustration in averages). However, such a choice depends on the location z . For a global choice use an integrated criterion like the weighted intergrated mean square error (WIMSE)

$$WIMSE(h, j) = \int MSE(\mathbf{z}, h, j)\psi(\mathbf{z})d\mathbf{z} = \int E[\hat{C}^{(j)}(\mathbf{z}) - C^{(j)}(\mathbf{z})]^2\psi(\mathbf{z})d\mathbf{z} \tag{16}$$

where $\psi(\mathbf{z})$ is a nonnegative weight function which ensures that WIMSE is well defined. Instead of an integrated criterion also an averaged criterion like the mean average squared error (MASE) can be used which replaces integration with summation in (16). When using a second order kernel straightforward calculations yield

$$h^* = \begin{cases} cn^{-1/(2p+d+2)} & \text{for } p-j \text{ odd} \\ c'n^{-1/(2p+d+4)} & \text{for } p-j \text{ even} \end{cases} \quad (17)$$

for the optimal bandwidth h^* in a multivariate setting. In our case of interest for $p = 3, j = 2$ and $d = 1$ it is $h^* = n^{-1/9}$. As for larger j also p must be enlarged, the optimal bandwidth for estimating the j -th derivative decreases in j . Note, however, that h^* is not feasible in practice because the constants c, c' in (17) contain unknown quantities such as higher derivatives of C and the regressors density f_Z as we can see from (9), (10) and (11).

A way to operationalize these are plug-in methods. They replace unknown quantities by pilot estimates and then calculate h^* via (17). The rule-of-thumb additionally uses normality assumptions in the distribution of the regressors and for the kernel to calculate exact constants. For $p = j = 0$ it is $h_k \approx s_k n^{-1/(4+d)}$ for $h = (h_1, \dots, h_d)$ with s_k the standard deviation of observations of covariate Z_k . It is an easy and fast way to obtain a rough estimate and can be used for pilot estimates in plug-in procedures. Nevertheless, a bandwidth choice based on these procedures yields only an asymptotically optimal selection as the employed criteria are asymptotic ones.

In small samples, however, there are better choices which can be made by data driven cross-validation methods. In general, these procedures yield precise finite sample results, but do not have closed form solutions. Therefore computation intensive numerical methods must be used in order to obtain h_{CV} . This can amount to a feasibility issue in particular for time series. We present a least squares cross-validation for local cubic estimation as our interest is in estimating the second derivative of C . Here select h_{CV} as minimizer of the sum of squared errors between obtained local cubic fit and observed response.

$$CV(\tilde{h}) = \sum_{i=1}^n \sum_{j \neq i}^n \left\{ Y_i - \hat{C}_{\tilde{h}, -i}(\mathbf{Z}_i) - \hat{C}_{\tilde{h}, -i}^{(1)}(\mathbf{Z}_i)(\mathbf{Z}_j - \mathbf{Z}_i) - \frac{1}{2} \hat{C}_{\tilde{h}, -i}^{(2)}(\mathbf{Z}_i)(\mathbf{Z}_j - \mathbf{Z}_i)^2 \right\}^2 M(\mathbf{Z}_i) \quad (18)$$

where $0 \leq M(\mathbf{Z}_i) \leq 1$ is a weight function that ensures existence of the limit for n large. and $(\hat{C}_{\tilde{h}, -i}, \hat{C}_{\tilde{h}, -i}^{(1)}, \hat{C}_{\tilde{h}, -i}^{(2)})$ denote the local cubic regression estimate obtained without using the i -th observation (\mathbf{Z}_i, C_i) . This way we ensure that the observations used for calculating $\hat{C}_{\tilde{h}, -i}(\cdot)$ are independent of \mathbf{Z}_i . It can be shown that asymptotically h_{CV} converges to the corresponding theoretical

bandwidth obtained from (17). This design driven choice of bandwidth is completely free of functional form assumptions and therefore most appealing in finite samples at the expense of potentially long computation time.

2.3 Dimension reduction techniques

While flexible, a high-dimension kernel regression requires large data samples for precise results in terms of confidence intervals. Ait-Sahalia and Lo (1998), for example, use a one year option data to empirically derive the call function based on five-dimensional kernel regression. Quite generally, larger samples make for smaller variances. Conversely, small samples make the bias-variance tradeoff even more acute. For a given bias, the variance of the model response is larger than for a model built from a larger sample. This is referred to as the "curse of dimensionality" (see Section 2.1. for theoretical details). Hence, there is a need to keep the dimension or equivalently the number of regressors low.

There has been lot of effort in developing methods which reduce the complexity of high dimensional regression problems resulting in better feasibility. In particular, the reduction of dimensionality is achieved by putting some structure on the model by e.g. imposing a parametric model. The resulting models are so-called semiparametric models, among which the additive models are the most feasible methods considered in practice. The basic idea of additive models is to take advantage of the fact that a regression surface may be of a simpler structure, that is a function of only certain linear combinations of the coordinates of the predictor variables such that high dimensional surface collapses down to one-dimensional surface. We refer to Mammen, Linton and Nielsen (1999) and Linton and Nielsen (1995) for detailed methods in this case. Here, however, we will focus on suitable parametric assumptions tailored to financial modeling.

One way is to use no-arbitrage arguments and collapse S_t , $r_{t,\tau}$ and $\delta_{t,\tau}$ into the forward price $F_t = S_t e^{(r_{t,\tau} - \delta_{t,\tau})\tau}$ in order to express the call pricing function as

$$C(S_t, X, \tau, r_{t,\tau}, \delta_{t,\tau}) = C(F_{t,\tau}, X, \tau, r_{t,\tau})$$

Alternatively use the non-arbitrage relation to estimate dividends and express the function in terms of the discounted stock price, that is either by $S_t^0 = S_t * e^{-\delta_{t,\tau}} = S_t - Div_{t,\tau}$ where $Div_{t,\tau}$ is the present value of the dividends to be paid before the expiration.

$$C(S_t, X, \tau, r_{t,\tau}, \delta_{t,\tau}) = C(S_t^0, X, \tau, r_{t,\tau})$$

A further reduction of the number of regressors is achieved by assuming that the call option function is homogeneous of degree one in S_t and X so that

$$C(S_t, K, \tau, r_{t,\tau}, \delta_{t,\tau}) = X C(S_t/K, \tau, r_{t,\tau}, \delta_{t,\tau}).$$

Combining the assumptions of the last two equations, the call price function can be further reduced to a function of three variables: moneyness $M_t = \frac{S_t^0}{K}$, maturity τ and risk free interest rate: $r_{t,\tau}$. Notice that by smoothing with respect to moneyness, rather than to the dividend adjusted index level we implicitly assume the theoretical option function is homogeneous of degree one with respect to the index and strike price. The basic Black-Scholes formula is an example of such a function, and as shown by Merton (1973) and discussed in Ingersoll (1987), a call price is homogeneous of degree one in the asset price and strike price if the asset's return distribution is independent of the level of the underlying index.

These dimension reduction techniques may be used both in direct estimation of RND from the call prices and indirectly via implied volatility. In the empirical study in the next subsection we will use only one regressor.

2.4 Application

We use tick data on the DAX index based European options prices maturing in one month (21 trading days), provided by EUREX for 20040121. The transformed data according to a methodology by Fengler (2005) contain date stamp, implied volatility, type of the option, maturity, strike price, option price, interest rate, intraday future price, average dividend rate.

The index stock price varies within one day and one needs to identify the price at which a certain transaction has taken place. Intraday DAX index prices are available on EUREX. Several authors (E.g. Jackwerth 2000) report that the change of the index price is stale and for every pair option/strike we use instead the prices of futures contracts closest to the time of the registered trade.

Original strike prices are given on an equidistant grid and in order to account for movements in the intraday price we use the following transformation $\frac{X_i}{F_i} S_t e^{r_{t,\tau} - \delta_{t,\tau}}$, where X_i and F_i are paired observations and S_t is the median intraday stock price, $r_{t,\tau}$ is the one month interest rate (linearly interpolated EURIBOR rates, for the desired maturity) and $\delta_{t,\tau}$ the average dividend. Conditional on these values we estimate q and interpret it as an average curve for the estimation date.

We use only at-the-money and out-of-the-money call options and in-the-money puts translated in call prices by using the put call parity

$$C_t - P_t = S_t e^{-\delta_{t,\tau}\tau} - X e^{-r_{t,\tau}\tau}$$

This guarantees that unreliable observations (high volatility) will be removed from our estimation samples. Since, as mentioned before, the intraday stock price varies, we use its median to compute the risk neutral density. For this price, we verify if our observations satisfy the arbitrage condition and delete for our sample those who do not satisfy it

$$S_t \geq C_i \geq \max(S_t - X_i e^{-r_t, \tau \tau}, 0).$$

Finally, if we have different call price observations for the same strike price we take their median at that point. In this way we ensure that we have a one to one relationship between every call and strike price.




Fig. 2. $\hat{q}(S_T)$ by local polynomial smoother for the optimally chosen bandwidth $h = 114.34$ by cross-validation (simple line) and oversmoothing bandwidths $h = 227.59$ (dashed line) and $h = 434.49$ (dotted line)

Smoothing in call option space

As described in Section 2.1 local polynomial method allows us to compute the second derivative of the call price directly, in a single step. We use local polynomial smoothing of degree three and a quartic kernel. In the first step we rescale the call price by dividing it by S_t and we smooth in this direction. We use cross-validation to choose the optimal bandwidth; however, this bandwidth yields a wiggly estimator and we decide to increase further the bandwidth. Smoothing comes at the expense of higher bias as we can see in Figure 2.

Smoothing in implied volatility space

In practice, the smoothing is mainly done in the implied volatility span because call prices are a more volatile function of the underlying asset price. In the present context, implied volatility is the volatility that yields a theoretical value for the option equal to the observed market price of that option, when using the Black-Scholes pricing model. We then estimate a smooth function $\hat{\sigma}$ and recover the call price by a bijective function evaluated at some fixed values of the regressors and variable σ

$$\begin{aligned}\hat{C}(S_t, X, \tau, r_{t,\tau}, \delta_{t,\tau}) &= C_{BS}(\cdot; \hat{\sigma}(S_t, X, \tau, r_{t,\tau}, \delta_{t,\tau})) \\ &= e^{-\delta_{t,\tau}\tau} S_t \Phi(y + \sigma\sqrt{\tau}) - e^{-r_{t,\tau}\tau} X \Phi(y)\end{aligned}$$

where Φ is the distribution function of the standard normal distribution and

$$y = \frac{\log(\frac{S_t}{K}) + (b - \frac{1}{2}\sigma^2)\tau}{\sigma\sqrt{\tau}}.$$

In this chapter we use a method based on Rookley (1997) who shows how to improve the efficiency of the estimator by estimating σ and its first two derivatives by local polynomial regression and plugging them into a modified version of the Black-Scholes formula. Below we describe the method for fixed maturity of one month.

For each pair (C_i, X_i) we define the rescaled call option $c_i = C_i/S_t$ in terms of moneyness $M_i = S_t/X_i$ so that starting from the Black-Scholes formula for the call price we can write

$$\begin{aligned}c_i &= c\{M_i; \sigma(M_i)\} = \Phi(d_1) - \frac{e^{-r\tau}\Phi(d_2)}{M_i} \\ d_1 &= \frac{\log(M_i) + \{r_{t,\tau} + \frac{1}{2}\sigma(M_i)^2\}\tau}{\sigma(M_i)\sqrt{\tau}} \\ d_2 &= d_1 - \sigma(M_{it})\sqrt{\tau}.\end{aligned}$$

For simplification we drop the indices. The risk neutral density can be expressed in terms of rescaled call price

$$q(\cdot) = e^{r\tau} \frac{\partial^2 C}{\partial X^2} = e^{r\tau} S \frac{\partial^2 c}{\partial X^2}$$

with

$$\frac{\partial^2 c}{\partial K^2} = \frac{d^2 c}{dM^2} \left(\frac{M}{X}\right)^2 + 2 \frac{dc}{dM} \frac{M}{X^2}$$

and

$$\begin{aligned} \frac{d^2c}{dM^2} &= \varphi(d_1) \left\{ \frac{d^2d_1}{dM^2} - d_1 \left(\frac{dd_1}{dM} \right)^2 \right\} \\ &\quad - \frac{e^{-r\tau}\varphi(d_2)}{M} \left\{ \frac{d^2d_2}{dM^2} - \frac{2}{M} \frac{dd_2}{dM} - d_2 \left(\frac{dd_2}{dM} \right)^2 \right\} \\ &\quad - \frac{2e^{-r\tau}\Phi(d_2)}{M^3} \end{aligned}$$

where φ is the probability density function of the standard normal distribution. The results depend further on the following quantities, where $\sigma(M)$, $\sigma'(M)$, $\sigma''(M)$ are smooth functions in moneyness direction

$$\begin{aligned} \frac{d^2d_1}{dM^2} &= - \frac{1}{M\sigma(M)\sqrt{\tau}} \left\{ \frac{1}{M} + \frac{\sigma'(M)}{\sigma(M)} \right\} \\ &\quad + \sigma''(M) \left\{ \frac{\sqrt{\tau}}{2} - \frac{\log(M) + r\tau}{\sigma(M)^2\sqrt{\tau}} \right\} \\ &\quad + \sigma'(M) \left\{ 2\sigma'(M) \frac{\log(M) + r\tau}{\sigma(M)^3\sqrt{\tau}} - \frac{1}{M\sigma(M)^2\sqrt{\tau}} \right\} \end{aligned}$$

$$\begin{aligned} \frac{d^2d_2}{dM^2} &= - \frac{1}{M\sigma(M)\sqrt{\tau}} \left\{ \frac{1}{M} + \frac{\sigma'(M)}{\sigma(M)} \right\} \\ &\quad - \sigma''(M) \left\{ \frac{\sqrt{\tau}}{2} + \frac{\log(M) + r\tau}{\sigma(M)^2\sqrt{\tau}} \right\} \\ &\quad + \sigma'(M) \left\{ 2\sigma'(M) \frac{\log(M) + r\tau}{\sigma(M)^3\sqrt{\tau}} - \frac{1}{M\sigma(M)^2\sqrt{\tau}} \right\}. \end{aligned}$$

In order to estimate $\sigma(M)$ and its associated first and second derivatives with respect to moneyness we use univariate local polynomial kernel regression of degree three and quartic kernel. The optimal bandwidth has been computed using cross-validation criteria (18) for the implied volatility. By increasing the bandwidth in Figure 3 we observe that oversmoothing improves the tails while having little effects on the values of \hat{q} situated in the middle of the distribution. It follows that smoothing in implied volatility yields a more robust estimator to the changes in the bandwidth. It is because the implied volatility is less sensitive to the changes in strike price than the call price.

Problems and refinements

In applications the support of strike prices is mostly compact and thus bounded. As shown in Section 2.1. the quality of estimates in regions close to the boundary might be low due to small values of the regressors' density when




Fig. 3. $\hat{q}(S_T)$ by Rookley method with oversmoothing bandwidth $h = 372.42$

using even order polynomials. By using a polynomial of order 3, estimation is design adaptive for the second derivative avoiding this problem.

Furthermore, associated with the boundary, option data is characterized by scarce observations close to the bounds. In general, nonparametric techniques do not perform well in regions with sparse data and other methods are required. Parametrization of the tails using Pareto type distributions might be advantageous leaving however the question of how to join the two regions in order to assure that the resulting distribution integrates to one. Alternatively, Rookley (1997) proposes to further parametrize these distributions by matching them with an Edgeworth expansion type density

$$q(S_T) = \frac{1}{S_T \sigma} \Phi(Z) \{1 + \beta(Z^3 - 3Z) + \gamma(Z^4 - 6Z^2 + 3)\}$$

for $Z = \frac{\log(S_T) - \alpha}{\sigma}$, where α and σ are the conditional mean and standard deviation of $\log(S_T)$ implied by the risk neutral measure, and β and γ are coefficients related to the higher moments of $\log(S_T)$.

According to economic theory, a nonparametric estimate of the call price function C must satisfy certain high-level conditions: It should be (1) positive,

(2) decreasing in X , (3) convex, and (4) its second derivative should exist, be nonnegative and integrable. Given that the first derivative of C with respect to X is the (negative) discounted cumulative density function of q conditions (2) and condition (3) can be summarized by the following inequality

$$-e^{r_{t,\tau}} \leq \frac{\partial C(S_t, X, \tau, r_{t,\tau}, \delta_{t,\tau})}{\partial X} \leq 0.$$

Convexity requires

$$\frac{\partial^2 C(S_t, X, \tau, r_{t,\tau}, \delta_{t,\tau})}{\partial^2 X} \geq 0.$$

Nonparametric kernel estimates may violate these constraints, unless we deal with large samples of observations. Imposing constraints like monotonicity or convexity directly in the estimation leads to nontrivial optimization problems in topological cones. If it is crucial for the outcome to fulfill the shape restrictions in small samples, it is recommended use series type estimation methods which easily allow to incorporate them directly in the estimation. In general, these constraints must be applied directly to the call price, because theoretical properties of the implied volatility are not well known. For further references see Ait-Sahalia (2003). This will be illustrated in the next section.

3 Estimation of the RND via empirical pricing kernel

In the previous section, we studied nonparametric kernel methods for estimating q as the discounted second derivative of the call price function and discussed the problems associated with kernel type estimators in this setting. Now, we propose a new approach, based on series expansion of the pricing kernel.

In financial mathematics the relationship between the physical measure p and RND q of a financial asset can be represented via the pricing kernel m . Also called stochastic discount factor, the pricing kernel is the quotient of the Arrow security prices and the objective probability measure and summarizes information related to asset pricing. Based on this relation

$$q(S_T) = m(S_T)p(S_T). \tag{19}$$

From a behavioral economics perspective m describes risk preferences of a representative agent in an exchange economy. In many applications, the empirical pricing kernel is the object of interest. In most of the studies Ait-Sahalia and Lo (2000), Brown and Jackwerth (2004), Grith et al. (2009) it has been estimated as a ratio of two estimated densities: \hat{q} computed as the second derivative of a smooth call function (as described in Section 2) and \hat{p} based

on historical returns. This approach leads to difficulties in deriving the statistical properties of the estimator. In particular, the sample sizes for estimating p and q may differ substantially: p uses daily observations, whereas q is based on intraday high-frequency observations. On the other hand, methods for estimating p are in general much simpler and more stable compared to those for q for which typically nonparametric kernel estimation of a second derivative is required. Direct estimation of the pricing kernel can be seen as an improvement in this sense.

For estimating q , however, a series approach is additionally appealing, as high-level shape constraints are straightforward to incorporate in finite samples. Recall that for kernel type estimators this is not the case, see the end of Subsection 2.4.

We introduce the series expansion for the pricing kernel in (3). With an estimate of the physical measure from historical data and the pricing kernel m from option prices, these indirectly imply an estimate of q via (19). In statistical theory and also in practice, this indirect way of estimating q is much more stable than using series methods directly for q in (1). In particular, in (19) large values of S_T are downweighted by integrating over the physical measure, while they enter undamped in (1) leading to unreliable results.




Fig. 4. First five terms of the Laguerre polynomials

3.1 Direct estimation of pricing kernel via series methods

As for q , there are several factors which drive the form of the pricing kernel. Here, however, we focus on the projection of the pricing kernel on the set of available payoff functions m^* , which allows us to represent m in terms of S_T only. In practice this is a reasonable assumption. Thus we require that m and m^* are close in the following sense

$$\|m - m^*\|^2 = \int |m(x) - m^*(x)|^2 dx < \epsilon \tag{20}$$

with ϵ small. Further we assume that m^* has a Fourier series expansion

$$m^*(S_T) = \sum_{l=1}^{\infty} \alpha_l g_l(S_T), \tag{21}$$

where $\{\alpha_l\}_{l=1}^{\infty}$ are Fourier coefficients and $\{g_l\}_{l=1}^{\infty}$ is a fixed collection of basis functions. The functions g_l are taken as orthonormal with respect to a particular norm. Such a representation is possible if the function is absolutely integrable.

Based on (21), we can construct an estimator for m^* and thus m . If a finite number L of basis functions is sufficient for a good approximation of m then

$$\hat{m}(S_T) = \sum_{l=1}^L \hat{\alpha}_l g_l(S_T). \tag{22}$$

Estimates $\hat{\alpha}_l$ for the coefficients α_l could be obtained by least squares for fixed basis functions g_l if a direct response was observable. Clearly the choice of L controls the quality of the estimate. The larger L , the better the fit but the higher the computing cost and less robust the result. See Subsection 2.2 for a sophisticated way of selecting the smoothing parameter.

In financial applications the following polynomial basis functions are frequently used: e.g. Laguerre, Legendre, Chebyshev polynomials, see Subsection 3.3. While asymptotically equivalent, in finite samples their form will influence the size of L . In general, one would prefer to have g_l such that L small is sufficient. For a formal criterion on how to select between different basis options see Li and Racine (2007). They assess different candidate basis functions by comparing a *CV*-type criterion for fixed L .

Though the form of m is only indirectly determined by relating observable call prices Y_i to strike prices X_i for given T, τ via (3). A response to observed payoffs via the pricing kernel is not directly available in the data. But in sample an estimate of m should fulfill

$$\begin{aligned}
Y_i &= e^{-r_{t,\tau}\tau} \int_0^\infty (S_T - X_i)^+ \sum_{l=1}^L \hat{\alpha}_l g_l(S_T) p_t(S_T) dS_T + \varepsilon_i \\
&= \sum_{l=1}^L \hat{\alpha}_l \left\{ e^{-r_{t,\tau}\tau} \int_0^\infty (S_T - X_i)^+ g_l(S_T) p_t(S_T) dS_T \right\} + \varepsilon_i
\end{aligned} \tag{23}$$

with error ε such that $E[\varepsilon|X] = 0$. Set

$$\psi_{il} = \psi_l(X_i) = e^{-r_{t,\tau}\tau} \int_0^\infty (S_T - X_i)^+ g_l(S_T) p_t(S_T) dS_T. \tag{24}$$

Then for known p and fixed basis functions and fixed L , the vector $\hat{\alpha} = (\hat{\alpha}_1, \dots, \hat{\alpha}_L)^\top$ is obtained as

$$\operatorname{argmin}_\alpha \sum_{i=1}^n \left\{ Y_i - \sum_{l=1}^L \alpha_l \psi_l(X_i) \right\}^2 \tag{25}$$

In practice, however, p is not known and can only be estimated. Therefore instead of ψ_l in (24) we have only estimates $\hat{\psi}_l$ of the basis functions. There are two possible ways for constructing them. First, regard ψ as an expectation which can be estimated by sample averaging over J different payoffs at time T for fixed τ and given X

$$\hat{\psi}_{il} = e^{-r_{t,\tau}\tau} J^{-1} \sum_{s=1}^J (S_T^s - X_i)^+ g_l(S_T^s) \quad \text{with } . \tag{26}$$

How $(S_T^s)_{s=1}^J$ are obtained is explained in detail in the following subsection. Or alternatively, replace p by an estimator, e.g. a kernel density estimator. This gives

$$\hat{\psi}_{il} = e^{-r_{t,\tau}\tau} \int_0^\infty (S_T - X_i)^+ g_l(S_T) \hat{p}(S_T) dS_T. \tag{27}$$

Here some care is needed for the numerical integration to keep discretization errors negligible. Furthermore, for an appropriate choice of bandwidth in \hat{p} , both approaches are asymptotically equivalent. In finite samples, however, estimates for q might differ, see Figure 5).

In total we obtain a feasible estimator of α based on a feasible version of (25) as

$$\hat{\alpha} = (\hat{\Psi}^\top \hat{\Psi})^{-1} \hat{\Psi}^\top Y. \tag{28}$$

The elements of $\hat{\Psi}_{(n \times L)}$ are given either by (26) or (27) and $Y = (Y_1, \dots, Y_n)^\top$.

Then an estimate of the pricing kernel at s is given by

$$\hat{m}(s) = g^L(s)^\top \hat{\alpha}, \tag{29}$$

where $g^L(s) = (g_1(s), \dots, g_L(s))^\top$. We see in figure () that the estimator or m is less stable for different approximations of ψ_{il} . Finally, the risk neutral density is estimated as

$$\hat{q}(s) = \hat{m}(s) \hat{p}(s). \tag{30}$$




Fig. 5. $\hat{q}(S_T)$ by Laguerre basis expansion with $L = 5$ based on approximation (26) (simple line) and (27) of Ψ (dashed line)

3.2 Estimation of the PDF of S_T

In the empirical study we use two different ways of obtaining DAX Index prices at time T . And we look at the sensitivity of \hat{q} w.r.t. \hat{p} . First, we extrapolate possible realizations of S_T in the future from historical log-returns. Based on a sample of historical DAX Index values we get

$$S_T^s = S_t e^{r_T^s}, \quad \text{for } r_T^s = \log(S_{t-s}/S_{t-(s+1)}).$$

Alternatively, we use a GARCH(1,1) specification for the log-returns to account for slowly decaying autocorrelation in the data exhibiting persistence. The model is specified as follows

$$\log(S_t/S_{t-1}) = \mu + u_t, \quad u_t \sim f(0, \sigma_t^r). \tag{31}$$

In equation (31), the returns consist of a simple constant, plus an uncorrelated, non-Gaussian disturbance. The conditional variance follows an ARMA(1,1) type specification




Fig. 6. $\hat{m}^*(S_T)$ by Laguerre basis expansion with $L = 5$ based on approximation (26) (simple line) and (27) of Ψ (dashed line)

$$(\sigma_t^r)^2 = a_1 + a_2 r_{t-1}^2 + a_3 (\sigma_{t-1}^r)^2. \tag{32}$$

We can estimate the parameters of the model (μ, a_1, a_2, a_3) and retrieve a time series of stochastic volatilities $\{\sigma_{t-s}^r\}_{s=1}^J$. The simulated index prices at time T are obtained as in (3.2) above for

$$r_T^s = r_{t-s} \frac{\sigma_T^r}{\sigma_{t-s}^r}$$

where we use for the forecasted volatility σ_T^r today's volatility σ_t .

Then the probability density p of S_T is estimated at each point S_T using a kernel density estimator

$$\hat{p}_h(S_T) = \frac{1}{Jh} \sum_{s=1}^J K\left(\frac{S_T^s - S_T}{h}\right) \tag{33}$$

where K is a kernel function and the bandwidth is selected similarly to the criteria introduced in Subsection 2.2. The two approaches are illustrated in

Figure 7. We observe that they differ mainly in the tails for p which carries over to q via (30).




Fig. 7. \hat{q} by Laguerre basis expansion with $L = 5$ (red) and \hat{p} based on log-returns (simple line) and weighted log-returns (dashed line)

3.3 Choice of tuning parameters

The quality of the obtained series estimators (29) and (30) depends on a suitable choice of the number $L(n) \rightarrow \infty$ for $n \rightarrow \infty$ for given basis functions. Note that the role of L (or L/n) is similar to that played by the smoothing parameter h for the kernel methods. There are three well-known procedures for a data-driven optimal selection of L . The first one is Mallows's C_L as proposed in Mallows (1973): Select L_M such that it minimizes

$$C_L = n^{-1} \sum_{i=1}^n \left\{ Y_i - \sum_{l=1}^L \hat{\alpha}_l \hat{\psi}_l(X_i) \right\}^2 + 2\sigma^2(L/n)$$

where σ^2 is the variance of ε . One can estimate σ^2 by

$$\hat{\sigma}^2 = n^{-1} \sum_{i=1}^n \hat{\varepsilon}_i^2$$

with $\hat{\varepsilon}_i = Y_i - \sum_l \hat{\alpha}_l \hat{\psi}_l(X_i)$.

A second option is selecting L according to generalized cross-validation suggested by Craven and Wahba (1979). Choose L_{GCV} minimizing

$$CV_L^G = \frac{n^{-1} \sum_{i=1}^n \left\{ Y_i - \sum_{l=1}^L \hat{\alpha}_l \hat{\psi}_l(X_i) \right\}^2}{\{1 - (L/n)\}^2}.$$

The last criterion is leave-one-out cross-validation according to Stone (1974): Select L_{CV} minimizing

$$CV_L = \sum_{i=1}^n \left\{ Y_i - \sum_{l=1}^L \hat{\alpha}_l^{-i} \hat{\psi}_l(X_i) \right\}^2$$

where $\hat{\alpha}_l^{-i}$ is the leave one estimate of α_l obtained by removing (X_i, Y_i) from the sample.

Li (1987) showed that each of the above three criteria leads to an optimally selected L in the sense that they all minimize the asymptotic weighted integrated squared error (see (16)). In this sense the obtained L are asymptotically equivalent.

3.4 Statistical properties

Series type estimators are designed to provide good approximations in an L_2 sense, see (20). Therefore asymptotic properties as consistency and rates of convergence should be derived from the asymptotic mean squared error. The rate of convergence for the indirect estimator of q via the pricing kernel depends on the two smoothing parameters h and L .

$$\begin{aligned} \int_0^\infty \{\hat{q}(S_T) - q(S_T)\}^2 dS_T &= \int_0^\infty \{\hat{m}(S_T)\hat{p}(S_T) - m(S_T)p(S_T)\}^2 dS_T \\ &= \int_0^\infty [\hat{m}(S_T)\{\hat{p}(S_T) - p(S_T)\}]^2 dS_T \\ &\quad + \int_0^\infty [p(S_T)\{\hat{m}(S_T) - m(S_T)\}]^2 dS_T \\ &\quad + \int_0^\infty 2\hat{m}(S_T)\{\hat{p}(S_T) - p(S_T)\}p(S_T)\{\hat{m}(S_T) - m(S_T)\} dS_T \end{aligned}$$

It easily follows from the law of iterated expectations that the third term equals zero. Consequently, the convergence of $\hat{q}(S_T)$ depends only on the first two terms. Since $\sup \hat{m}(s) = \mathcal{O}_P(1)$ under Assumption 1 given below, the order of convergence for the first term is dominated by $\{\hat{p}(S_T) - p(S_T)\}^2$.

Assumption 1 Suppose that p is twice continuously differentiable, K is a second order kernel and the bandwidth is chosen optimally as $h = cn^{-1/5}$, for a known constant c .

Then the asymptotic mean squared error for the kernel density estimator is

$$\|\hat{p}_h(x) - p(x)\|_2^2 = \mathcal{O}_P(n^{-4/5}) \tag{34}$$

This follows along the same logic as the results for local polynomials in Section 2.1. For further details see e.g. Härdle et al. (2004).

The order of convergence for the second term only depends on $\{\hat{m}(S_T) - m(S_T)\}^2$ since $\sup p(s) \leq 1$. The next assumption establishes consistency of $\hat{m}(S_T)$.

Assumption 2 $\{X_i, Y_i\}$ are i.i.d. observations of (X, Y) , $\text{Var}(Y|X)$ is bounded on S , the compact connected interval of support of X . Furthermore p is bounded away from zero and m is ν -times continuously differentiable on S . Choose L such that $L^3/n \rightarrow 0$ as $n \rightarrow \infty$.

Under Assumption 2 it is

$$\int_0^\infty \{\hat{m}(S_T) - m(S_T)\}^2 dS_T = \mathcal{O}_p(L/n + L^{-2\nu}). \tag{35}$$

This result is from Newey (1997) for fixed basis functions ψ_l . With estimated basis $(\hat{\psi}_l)_l$ the result still goes through as the convergence of $\hat{\psi}_l$ to the true ψ_l is at parametric rate. The i.i.d. assumption is for simplicity of the exposition only. It can be easily relaxed to mixing type of observations.

The theorem below puts (34) and (35) together for an asymptotic result for q .

Theorem 3. *Assume that Assumptions 1 and 2 hold. Then the integrated square error (ISE) converges as*

$$\int_0^\infty \{\hat{q}(S_T) - q(S_T)\}^2 dS_T = \mathcal{O}_p(n^{-4/5} + L/n + L^{-2\nu}). \tag{36}$$

3.5 Practical aspects for implementation

We illustrate the method using the data described in Subsection 2.4. We consider the univariate regression of C on the strike price X for fixed maturity and fixed interest rate. We estimate q using three different systems of orthogonal basis: Laguerre, Legendre and Chebyshev. We found that the fit of the call price is almost identical for fixed L , while $\hat{\alpha}$ varies obviously with the series. This little sensitivity with respect to the choice of the basis function

holds also for the empirical pricing kernel and the implied risk neutral density. Based on the selection criteria for L from Section 3.3, we have chosen $L = 5$. We exemplify the method with Laguerre polynomials. The first two terms are

$$\begin{aligned} g_1(x) &= 1 \\ g_2(x) &= 1 - x \end{aligned}$$

The other terms are given by the recurrence relation

$$g_{l+1}(x) = \frac{1}{l} \{(2l - 1 - x)g_{l-1}(x) - (l - 1)g_{l-2}(x)\} \quad \text{for } l = 2, \dots, L.$$

Estimation results are displayed in Figure (5) through Figure (8).




Fig. 8. $\hat{q}(S_T)$ by local polynomial regression with $h = 227.59$ in call space (black), by Rookley method $h = 372.42$ in IV space (green), indirect estimation of the pricing kernel as Laguerre basis expansion with $L = 5$ (green)

4 Conclusions

We have studied three nonparametric approaches for estimating the risk neutral density. They are based on fundamentally different techniques: two of them use local features and the third one is based on global curve fitting. For these approaches we have described the estimation methodology and their performance in finite sample, in terms of robustness and stability. Statistical properties of all procedures have been derived and illustrated focusing on practically most relevant aspects.

In terms of comparative performance, figure 8 shows estimates of q using the three methods we discussed in this article for suitable choices of tuning parameters. While all three nonparametric methods yield similar results, there still are some peculiarities. As we have seen so far, kernel methods for the estimation of q in implied volatility space work much better than those which smooth in the call price space. Local polynomial methods applied to call prices yield estimates which are highly sensitive to the choice of the bandwidth. Furthermore, extrapolation in the tails gives wiggly estimates unless we increase the bandwidth considerably. This comes at a high price of increased bias. In comparison to this, when we smooth in implied volatility space, the Rookley method yields a more stable estimate with respect to the choice of bandwidth and performs also better in the regions with sparse data. Estimation of risk neutral density based on the pricing kernel yields relatively stable results in small samples with respect to the choice of the basis function. The center distribution exhibits little sensitivity to the choice of basis functions, whereas the tails are highly sensitive due to scarcity of observations at the boundaries. This problem can be solved by imposing constraints on the estimator, as explained in the end of Section 2. Generally, series type methods should be preferred when performing constrained estimation.

References

1. Aït-Sahalia, Y. and Duarte, J. (2003) Nonparametric option pricing under shape restrictions. *Journal of Econometrics* 116: 947
2. Aït-Sahalia, Y. and Lo, A. W. (1998) Nonparametric estimation of state-price densities implicit in financial asset prices. *Journal of Finance* 53: 499-547
3. Aït-Sahalia, Y. and Lo, A. W. (2000) Nonparametric risk management and implied risk aversion. *Journal of Econometrics* 94: 9-51
4. Arrow, K. J. (1964) The role of securities in the optimal allocation of risk-bearing. *Review of Economic Studies* 31: 91-96
5. Bates, D. S. (1996) Jumps and stochastic volatility: Exchange rate processes implicit in deutsche mark options. *Review of Financial Studies* 9(1): 69-107
6. Black, F. and Scholes, M. (1973) The Pricing of Options and Corporate Liabilities. *Journal of Political Economy* 81: 637-654.
7. Breeden, D. T. and Litzenberger, R. H. (1978) Prices of state-contingent claims implicit in option prices. *The Journal of Business* 51(4): 621-651
8. Brown, D.P. and Jackwerth, J. C. (2004) The pricing kernel puzzle: Reconciling index option data and economic theory. Manuscript
9. Campbell, J., Lo, A. and McKinlay, A. (1997) *The econometrics of financial markets*. Princeton University Press, Princeton, New Jersey
10. Debreu, G. (1959) *Theory of value: An axiomatic analysis of economic equilibrium*. Yale University Press, New Haven
11. Engle, R. F. and Rosenberg, J. V. (2002) Empirical pricing kernels. *Journal of Financial Economics* 64: 341-372
12. Fan, J., and Gijbels, I. (1996) *Local polynomial modelling and its applications*. Chapman and Hall, London
13. Fengler, M. R. (2005) *Semiparametric modeling of implied volatility*, Springer Finance
14. Grith, M., Härdle, W. and Park, J. (2009) Shape invariant modelling pricing kernels and risk aversion. SFB649DP2009-041. Submitted to *Journal of Financial Econometrics* on 1 June 2009
15. Härdle, W. (1990) *Applied nonparametric regression*. Econometric Society Monographs No. 19. Cambridge University Press
16. Härdle, W. and Hlavka, Z. (2009) Dynamics of state price densities. *Journal of Econometrics* 150: 1-15
17. Härdle, W., Müller, M., Sperlich, S. and Werwatz, A. (2004) *Nonparametric and Semiparametric Models*. Springer Verlag, Heidelberg
18. Härdle, W., Okhrin, Y. and Wang, W. (2009) Uniform confidence for pricing kernels. SFB649DP2010-003. Submitted to *Econometric Theory* on 1 January 2010
19. Heston, S. (1993) A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Review of Financial Studies* 6: 327-343
20. Jackwerth, J. C. (1999) Option-implied risk-neutral distributions and implied binomial trees: a literature review. *Journal of Derivatives* 2: 66-82
21. Jackwerth, J. C. (2000) Recovering risk aversion from option prices and realized returns. *Review of Financial Studies* 13: 433-451
22. Li, K. C. (1987) Asymptotic optimality for c_p , c_l , cross-validation and generalized cross-validation: Discrete index set. *Annals of Statistics* 15: 958-975

23. Li, Q. and J. S. Racine (2007) *Nonparametric econometrics: Theory and practice*. Princeton University Press
24. Linton, O. and Nielsen, J. P. (1995). A kernel method of estimating structured nonparametric regression based on marginal integration. *Biometrika* 82: 93–100.
25. Longstaff, F. (1995) Option Pricing and the Martingale Restriction. *Review of Financial Studies* 8: 1091-1124.
26. Lucas, R.E. (1978) Asset prices in an exchange economy. *Econometrica* 46: 1429-1445
27. Mallows, C.L. (1973) Some comments on c_p . *Technometrics* 15: 661-675
28. Mammen, E., Linton, O. and Nielsen, J. (1999). The existence and asymptotic properties of a backfitting projection algorithm under weak conditions. *Annals of Statistics* 27 (5): 1443–1490.
29. Marron, J. S. and Nolan, D. (1988) Canonical kernels for density estimation. *Statistics and Probability Letters* 7(3): 195-199
30. Merton, R. C. (1973) Theory of rational option pricing. *Bell Journal of Economics, The RAND Corporation* 4(1): 141-183, Spring
31. Müller, H. G. (1988) *Nonparametric regression analysis of longitudinal data*. Lecture Notes in Statistics Vol. 46, New York, Springer-Verlag
32. W.K. Newey (1997) Convergence rates and asymptotic normality for series estimators. *Journal of Econometrics* 79: 147-168.
33. Renault, E., (1997) Econometric models of option pricing errors. In: Kreps, D.M. and Wallis, K.F. (eds) *Advances in Economics and Econometrics: Theory and Applications* vol. 3: 223-278. Cambridge University Press, Cambridge
34. Rookley, C. (1997) Fully exploiting the information content of intra day option quotes: Applications in option pricing and risk management, Working paper, University of Arizona
35. Rubinstein (1976) The valuation of uncertain income streams and the pricing of options. *Bell Journal of Economics* 7(2): 407-425
36. Ruppert, D. and Wand, M. P. (1994) Multivariate locally weighted least squares regression. *Annals of Statistics* 22(3): 1346-1370
37. Wahba, G. (1985) A comparison of GCV and GML for choosing the smoothing parameter in the generalized spline smoothing problem. *Annals of Statistics* 4: 1378-1402