Journal of Econometrics 184 (2015) 242–261

Contents lists available at ScienceDirect

Journal of Econometrics

journal homepage: www.elsevier.com/locate/jeconom

# Semi-nonparametric estimation of the call-option price surface under strike and time-to-expiry no-arbitrage constraints[☆]

Matthias R. Fengler[a,*], Lin-Yee Hin[b]

[a] University of St. Gallen, Department of Economics, Bodanstrasse 6, CH-9000 St. Gallen, Switzerland

[b] Department of Mathematics and Statistics, Curtin University, Kent Street, Bentley WA 6102, Australia

ARTICLE INFO

*Article history:*
Received 26 April 2013
Received in revised form
1 September 2014
Accepted 1 September 2014
Available online 16 September 2014

*JEL classification:*
C14
C58
G13

*Keywords:*
B-splines
No-arbitrage constraints
Option pricing function
Semi-nonparametric estimation
Shape-constrained regression
State-price density

ABSTRACT

We suggest a semi-nonparametric estimator for the call-option price surface. The estimator is a bivariate tensor-product B-spline. To enforce no-arbitrage constraints across strikes and expiry dates, we establish sufficient no-arbitrage conditions on the control net of the B-spline surface. The conditions are linear and therefore allow for an implementation of the estimator by means of standard quadratic programming techniques. The consistency of the estimator is proved. By means of simulations, we explore the statistical efficiency benefits that are associated with estimating option price surfaces and state-price densities under the full set of no-arbitrage constraints. We estimate a call-option price surface, families of first-order strike derivatives, and state-price densities for S&P 500 option data.



## 1. Introduction

Option prices carry information on the risk factors that drive the underlying asset price process. This information can be exploited to price other, more complex contingent claims consistently with the market, to study policy events, and to learn about the risk perception and risk attitude of the representative agent in the market. Numerous strategies have therefore been suggested for estimating an option pricing function from randomly observed option price data in order to extract the relevant information.

Unlike many other financial data, option price data have the particular feature that a number of expiry dates across different exercise prices are traded concurrently. This feature allows the simultaneous study of cross-sections of option price data over various time horizons. Jackwerth and Rubinstein (1996) and Aït-Sahalia and Lo (1998), for example, compare the state-price density implied from option prices across different expiry dates; in a similar vein, Aït-Sahalia and Lo (2000), Jackwerth (2000) and Bliss and Panigirtzoglou (2004) study the empirical pricing kernel and implied risk aversion. An assumption that is implicit in many of these studies is that option prices observed contemporaneously over multiple time horizons are realizations from a smooth surface defined across exercise prices and expiry dates.

In this paper, we suggest an estimator for the pricing function of a European-style call-option that extends across all available expiry dates. In other words, we explicitly take account of two dimensions and estimate a call-option price surface. The estimator is a bivariate tensor product B-spline. It therefore belongs to the flexible class of series estimators also called semi-nonparametric estimators; see Gallant (1987). As financial theory requires, the estimator obeys shape constraints both across the strike and the expiry dimension to ensure that the resulting option price surface

☆ We thank Francesco Audrino, Wolfgang Härdle, Jens Jackwerth and the seminar participants at the Erasmus University Rotterdam, the Konstanz workshop "The Pricing Kernel Puzzle", the Frankfurt School of Finance and Management, the Humboldt-Universität zu Berlin, the CFE 2011 conference, the MathFinance Conference, the 7th World Congress of the Bachelier Finance Society, the 5th Annual SoFiE Conference and the EEA-ESEM 2012 for their comments. The paper has particularly benefited from valuable suggestions made by Patrick Gagliardini, Enno Mammen and three anonymous referees. Any remaining errors are ours.

* Corresponding author. Tel.: +41 71224 2457.
*E-mail addresses:* matthias.fengler@unisg.ch (M.R. Fengler), linyeehin@gmail.com (L.-Y. Hin).

http://dx.doi.org/10.1016/j.jeconom.2014.09.003
0304-4076/

is free of arbitrage. Since option data may be sparsely distributed, which could lead to an ill-posed estimation problem, we also study a regularized version of the estimator that is well-behaved.

Estimating an option price surface under shape constraints has rarely been achieved in the literature so far. Two-dimensional estimators, such as those suggested in Aït-Sahalia and Lo (1998), Cont and da Fonseca (2002), and Fengler et al. (2007), do not accommodate no-arbitrage constraints. It is therefore a common practice to estimate univariate option pricing functions for each expiry date independently, usually in combination with some interpolation of the data across the expiry dates.

The literature on flexible modeling techniques designed to estimate a univariate option price function and/or its second-order strike derivative (the state-price density) under shape constraints is vast. The most important rivals of our estimator, however, are the fully nonparametric approaches, such as kernel smoothers, and semi-nonparametric (SNP) regression techniques that are based on series estimators, such as polynomial regression splines or other series expansions. Within the nonparametric stream, the pioneering work of Aït-Sahalia and Duarte (2003) suggests a local linear smoother to estimate the option pricing function and its strike derivatives. Alternative kernel regression estimators are proposed by Birke and Pilz (2009) and Fan and Mancini (2009). As regards SNP techniques, a number of polynomial spline methods have been suggested to date: B-splines in Wang et al. (2004), Laurini (2011), and Corlay (2013); smoothing splines in Yatchew and Härdle (2006), Monteiro et al. (2008), and Fengler (2009); linear splines in Härdle and Hlávka (2009). Other SNP-type estimators are based on the Edgeworth expansion as in Jarrow and Rudd (1982), on Hermite polynomials as in Madan and Milne (1994) and Jondeau and Rockinger (2001), or on approximation methods, such as the positive convolution approximation as in Bondarenko (2003) and the nonparametric density mixtures as in Yuan (2009). Finally, there are flexible estimation approaches that are nonparametric in nature, but do not properly fit into either strand, such as the neural network used in Hutchinson et al. (1994), the maximum entropy method suggested in Stutzer (1996) and the regularized calibration approach devised by Jackwerth and Rubinstein (1996).[1]

First approaches to include no-arbitrage constraints for surface estimation are Benko et al. (2007) and Glaser and Heider (2012). Both studies build on local polynomials, but they do not analyze the asymptotic properties of the constrained estimators, nor do they make an attempt to assess the efficiency benefits that are associated with implementing no-arbitrage constraints in the time-to-expiry dimension. Here, we exploit a projection framework for constrained smoothing devised by Mammen et al. (2001) to prove the consistency of the estimators and to provide an upper bound for their rates of convergence. For both estimators, this upper bound is the optimal rate for regression estimation owed to Stone (1982). In addition, we show by means of simulations that substantial efficiency gains are to be expected for the estimation of the option price surface and its derivatives, if one implements calendar conditions as well as strike constraints.

As noted above, the use of splines to estimate the option pricing function is not new in itself. Moreover, splines have a long tradition in the statistical literature on smoothing under shape constraints like positivity, monotonicity, and convexity; see Delecroix and Thomas-Agnan (2000) for a survey. Two main avenues for constraints implementation in spline spaces can be distinguished. The first exploits the fact that in the considered spaces the shape constraints can be represented by a finite number of linear inequality constraints to achieve the desired shape constraint globally; see, inter alia, Hildreth (1958), Brunk (1970), Dierckx (1980), Ramsay (1988), He and Shi (1998), and Meyer (2008). Alternatively, one seeks only approximately to satisfy the constraints on a finite subset of the domain of the function; see, e.g., Villalobos and Wahba (1987) and Mammen and Thomas-Agnan (1999). Both strands exploit specific properties of the spline spaces under consideration and impose the conditions directly on the unknown regression function.[2] For example, to impose convexity in a cubic spline space, one can utilize the linearity of second-order derivatives, which in turn leads to conditions on the coefficients of the spline.

Our approach differs. We do not impose the no-arbitrage shape constraints directly on the unknown regression function, but derive sufficient conditions for no-arbitrage on the *control net* of the tensor product (TP) B-spline. The notion of the control net, which is a set of points with certain averages of the knot sequences as abscissae and the B-spline coefficients as ordinates, originates in the literature on computer-aided geometric design; see Prautzsch et al. (2002). In some senses, the control net spans the shape of the TP spline surface. As an important property, it is independent of the degree of the B-spline, thereby allowing B-splines of arbitrary degree to be used for estimation (within numerical limitations). This is in contrast to the aforementioned literature that seldom generalizes beyond the polynomial degree for which it is developed. Nevertheless, our no-arbitrage conditions on the control net are linear. Constrained estimation can therefore be carried out by standard quadratic programming techniques. As we will discuss in detail in Section 4.3, apart from mechanically forcing a polynomial call-option price surface to be free of arbitrage, our conditions also have an economic interpretation. This is because they embed as a special case the findings of Carr and Madan (2005) and Davis and Hobson (2007), which apply to the linear call-option price surface.

A natural question is why we should use polynomial regression splines rather than kernel methods like Aït-Sahalia and Duarte's (2003) local polynomial method, and why specifically B-splines. Concerning the first point, it is difficult to argue that one approach is uniformly better than the other. The local polynomial method fits a polynomial in a local neighborhood around a given point and the estimate is given by a sequence of such fits. Consequently, the asymptotic behavior of the kernel estimator at a point is very well understood; see Fan (1992, 1993). For regression splines, one fits a piecewise polynomial by minimizing a global loss criterion. This makes it hard to obtain precise asymptotic bias expressions and could be seen as a disadvantage; see Zhou et al. (1998) and Huang (2003). On the other hand, regression splines and their derivatives are exhaustively characterized by the coefficients in a basis expansion. Since the number of basis functions is smaller than the sample size, one obtains a complete yet parsimonious summary of the underlying data. This feature may be of practical importance. As we also confirm in our simulations and empirical applications, in most practical situations, however, the two methods are likely to deliver similar estimates. This observation has a theoretical justification since a spline estimator can be shown to be asymptotically equivalent to a certain kernel smoother; see Silverman (1984) and Huang and Studden (1993).

With regard to the second question, B-splines are attractive among the polynomial splines, since they are compactly supported functions that form a basis for a polynomial spline space with a given degree, smoothness, and domain partition. Moreover, by means of the de Boor recursion formula, a stable algorithm for evaluating splines in B-spline form is available; see Appendix A. Hence, B-splines behave in a numerically favorable way and are easy to

[1] This short overview is far from complete. Most importantly, we omit the fully parametric models; see Jackwerth (2004) for more references.

[2] All of the studies cited above, which use polynomial splines to estimate a shape-constrained option pricing function or state-price density, can be assigned to either strand.