

Does hedging with implied volatility factors improve the hedging efficiency of barrier options? *

Szymon Borak[†]
Matthias R. Fengler
Wolfgang K. Härdle

CASE – Center for Applied Statistics and Economics
Humboldt-Universität zu Berlin,
Spandauer Straße 1, 10178 Berlin, Germany

February 11, 2009

Forthcoming: The Journal of Risk Model Validation

*We gratefully acknowledge financial support by the Deutsche Forschungsgemeinschaft and the Sonderforschungsbereich 649 “Ökonomisches Risiko”.

[†]Corresponding author: borak@wiwi.hu-berlin.de, TEL ++49 30 2093 5630 FAX ++49 30 2093 5649.

Does hedging with implied volatility factors improve the hedging efficiency of barrier options?

Abstract

The price of a barrier option depends on the shape of the entire implied volatility surface which is a high-dimensional dynamic object. Barrier options are hence exposed to nontrivial volatility risk. We extract the key risk factors of implied volatility surface fluctuations by means of a semiparametric factor model. Based on the factors we define a practical hedging procedure within a local volatility framework. The hedging performance is evaluated using DAX index options.

JEL classification codes: G11

Keywords: implied volatility surface, smile, local volatility, exotic options, semiparametric factor model, hedging

1 Introduction

In equity derivative markets barrier options are appealing instruments for investors looking for a partial protection of their equity allocation. From the perspective of an institution issuing barrier options this demand raises the need of efficient hedging strategies. This is a challenging task for at least two reasons. First, reverse barrier options, such as down-and-out puts and up-and-out calls, have discontinuous payoff profiles and knock out deep in-the-money thereby losing the maximum possible intrinsic value. Second, barrier options, as many other exotic options, are exposed to nontrivial volatility risk, since the knock-out probability strongly depends on the skew of the implied volatility smile. The latter effect also prevents simple Black-Scholes type formulae, such as those by Rubinstein and Reiner (1991), from being usable in practice.

Nowadays there is a plethora of models available that take the shape of the implied volatility surface (IVS) into account for option valuation. Potential candidates are: the local volatility (LV) model proposed by Dupire (1994), Derman and Kani (1994), and Rubinstein (1994), which introduces a nonparametric local volatility function that deterministically depends on the asset price and time; stochastic volatility models like Hull and White (1987), Stein and Stein (1991), Heston (1993), Carr et al. (2003); jump-diffusion models, such as Merton (1976), Bates (1996), and Kou (2002). When calibrated to the IVS, all these models are able to replicate the plain vanilla market to a similar extent, whereas their prices for barrier options may differ due to the different properties of the underlying asset price dynamics, see Hull and Suo (2002) and Hirt et al. (2003) on model risk for barrier options. The more challenging part is hedging. For it is straight forward to compute derivatives for the parameters of these models, but it is intricate to give the parameter greeks a meaning by mapping them on tradable instruments provided by the plain vanilla market. More seriously, since the prices of the hedging instruments, either over-the-counter or as listed options, are given in terms of implied volatility, they necessarily follow the dynamics of the IVS. Indeed it is in question whether the IVS dynamics inherent in the model that is calibrated to a static surface and used for pricing truly match the stylized facts of IVS dynamics, see Hagan et al. (2002) and Bergomi (2005) for such a discussion in context of the LV model and the Heston model, respectively. In contrast, the dynamics of the IVS are empirically well understood, see Skiadopoulos et al. (1999), Alexander (2001), Cont and da Fonseca (2002), Fengler et al.

(2003), Hafner (2004), Fengler et al. (2007) among others. The typical approach extracts the main driving factors like level, slope, or term structure movements and models these factors. It therefore appears natural to exploit this knowledge for hedging and portfolio risk management.

The aim of this paper is to study dynamic hedges of reverse barrier options built on factor functions of empirically observed IVS dynamics. We project the complex, high dimensional dynamics of the IVS on a low and finite dimensional space spanned by the semiparametric factor model (SFM)

$$\hat{\sigma}_t(\kappa, \tau) = \exp \left\{ \sum_{l=0}^L Z_{t,l} m_l(\kappa, \tau) \right\}, \quad (1)$$

where $\hat{\sigma}_t(\kappa, \tau)$ denotes the implied volatility of a certain moneyness κ and maturity τ observed in time t . The functions m are nonparametric components and invariant in time, while the time evolution is modelled by the latent factor series $Z_{t,l}$. In order to estimate (1) we apply an estimation technique suggested in Fengler et al. (2007). The SFM estimates the prevalent movements of the IVS in an $(L + 1)$ -dimensional function space.

Given the estimated factor functions \hat{m} , we construct hedges for barrier options priced in a LV model. We use a LV model, since by the nonparametric nature of the local volatility function it can match any arbitrage-free set of option prices to an arbitrarily precise degree. It will hence replicate the deformations of the IVS defined by the estimated factor functions and allow for a precise computation of factor greeks not prone to calibration error. Moreover, the LV model is numerically very efficient and allows for fast and accurate price valuations using the finite difference method. The factor hedges we obtain are more general than the usual vega hedges which are defined by a parallel shift of the IVS since they will take into account nontrivial surface movements, such as nonparallel up-and-down shifts, slope and term structure risks. Depending on the payoff profile of an exotic option, these risks can be substantial. Our approach is hence similar in spirit to Diebold et al. (2006) who define factor based duration measures and study the efficacy of these measures for the insurance of bond portfolios.

We note that strictly speaking it may not be necessary to vega hedge in an LV framework,

since it defines a complete market. This however is a theoretical perspective which does not correspond to market practice. When minimizing portfolio risk, traders are likely to set up vega hedges as soon as a liquid over-the-counter or listed option markets allow them to do so. In this sense our approach is e.g. similar to the practice of hedging a long dated plain vanilla option which are priced by means of a smile-adjusted Black-Scholes model by adding a short dated option to the portfolio.

The dynamic hedging performance of plain vanilla options in a LV model is studied in Dumas et al. (1998), Coleman et al. (2001), McIntyre (2001) and Vähämaa (2004), while the case of reverse barrier options is treated in Engelmann et al. (2006). Engelmann et al. (2006) implement hedging strategies that are delta ($\partial/\partial S$), vega ($\partial/\partial\sigma$) and vanna ($\partial^2/\partial\sigma\partial S$) neutral where vega and vanna are obtained by parallel shifts of the IVS and computing the difference quotient. We complement this analysis by defining sensitivities with respect to the most prevalent IVS movements motivated by model (1), namely ($\partial/\partial Z_1$), ($\partial/\partial Z_2$) and by constructing portfolios neutral to these greeks. For this purpose we establish a portfolio containing a reverse barrier option and hedge it on a daily basis with plain vanillas and the underlying asset using DAX data from January 3rd, 2000 to June 30th, 2004. We then study the distribution of the hedging errors across the different hedging strategies.

For completeness we remark that static hedging of barrier options is a competing way of portfolio insurance, see Derman et al. (1995), Carr and Chou (1997), Carr et al. (1998), Andersen et al. (2002), Tompkins (2002), Nalholm and Poulsen (2006a), Nalholm and Poulsen (2006b). For a static hedge one sets up a portfolio of plain vanillas which replicates the payoff of the barrier option as close as possible. The hedge is unwound in case of a knock-out or at expiry and no other adjustment of the hedge is necessary. In fact, Engelmann et al. (2007) and Maruhn et al. (2008) show that there are static hedges outperforming dynamic hedges. However, the practical use of static hedges is limited, since they may not always be implementable due to insufficient market depth of listed plain vanilla options.

The paper is structured as follows. In Section 2 we present the framework on which the empirical procedure is based. Section 3 concentrates on the description of the hedging method. In Section 4 we present the data, describe the empirical hedging design and discuss the empirical results. Section 5 concludes.

2 Models

2.1 Local Volatility Model

In the LV model the risk neutral price of the underlying asset is governed by the stochastic differential equation:

$$dS_t = r_t S_t dt + \sigma(S_t, t) S_t dW_t, \quad (2)$$

where W_t is a Wiener process and r_t denotes the instantaneous interest rate. Dividends are assumed to be zero, since the DAX, on which our empirical study is based, is a performance index. $\sigma(S_t, t)$ is the local volatility function which depends on the underlying price and time. This function has a unique representation if an arbitrage-free set of call options is given for all strikes and maturities, Dupire (1994). It can be shown that

$$\sigma^2(S_t, t) = \frac{2 \frac{\partial \hat{\sigma}(K, T)}{\partial T} + \frac{\hat{\sigma}(K, T)}{T} + 2K \int_0^T r_s ds \frac{\partial \hat{\sigma}(K, T)}{\partial K}}{K^2 \left\{ \frac{\partial^2 \hat{\sigma}(K, T)}{\partial K^2} - d_1 \sqrt{T} \left(\frac{\partial \hat{\sigma}(K, T)}{\partial K} \right)^2 + \frac{1}{\hat{\sigma}(K, T)} \left(\frac{1}{K\sqrt{T}} + d_1 \frac{\partial \hat{\sigma}(K, T)}{\partial K} \right)^2 \right\}} \Bigg|_{K=S_t, T=t} \quad (3)$$

where $d_1 = \frac{\log(S_0/K) + \int_0^T r_s ds + 0.5 \hat{\sigma}^2(K, T) T}{\hat{\sigma}(K, T) \sqrt{T}}$ and where $\hat{\sigma}(K, T)$ is the implied volatility at strike K and expiry T . Formula (3) gives a correspondence between local and implied volatility surfaces.

The LV model received much attention in the finance community since it achieves an almost exact fit of the observed vanilla market and is numerically and computationally very tractable. The price of the barrier option denoted by V with barrier B and expiry date T is obtained by numerically solving the partial differential equation

$$r_t V(S, t) = \frac{\partial V(S, t)}{\partial t} + \frac{1}{2} \sigma^2(S, t) S^2 \frac{\partial^2 V(S, t)}{\partial S^2} + r_t S \frac{\partial V(S, t)}{\partial S} \quad (4)$$

with additional boundary conditions, i.e. $V(B, t) = 0$ for $t < T$ and $V(S, T)$ equal to the payoff at expiry. For calibration of the model a number of methods are available, see Bouchouev and Isakov (1999) for comprehensive review. For example one may directly apply

the formula (3). Here we adopt the approach of Andersen and Brotherton-Ratcliffe (1997) which determines r and σ so that forwards, zero coupon bonds and plain vanilla options are priced correctly on each grid point. The finite difference method then gives barrier option prices and sensitivities very efficiently.

Yet the LV is also subject to criticism, see Fengler (2005, Chapter 3.11) for the details of this discussion. The severest objection was brought forward by Hagan et al. (2002) by showing that the LV model implies unrealistic smile dynamics and consequently wrong spot greeks. In practice this problem can be addressed by enforcing the desired smile dynamics when computing the greeks. Instead of calculating model-consistent LV greeks, one fixes the IVS in strikes (sticky-strike) or in moneyness (sticky-moneyness) and recalibrates the LV surface under the spot movements. Engelmann et al. (2006) find that the empirical performance of the dynamic hedges is negligible under different stickiness assumptions, if a vega hedge is implemented. Overall they find that the sticky-strike approach, which we will adopt here, performs best. We therefore believe that the LV model serves well for the purpose of this study.

2.2 The Semiparametric Factor Model

To model the IVS dynamics we employ the SFM which yields estimates of the IVS for each day of the sample and explains its dynamic behavior by extracting a small number of key driving factors of the surface movements. For this aim one could use any other factor model like the functional principal components model of Cont and da Fonseca (2002) or the parametric model of Hafner (2004). An alternative definition of the skew shifts can be also found in Taleb (1997). Our choice for the SFM is motivated by the flexible nonparametric structure, which allows to extract the most important factors along with a dimension reduction, and its adaptedness to the expiry behavior of implied volatility data, see Fengler et al. (2007) for details.

To describe the SFM denote by $Y_{t,j}$ the log-implied volatility observed on day $t = 1, \dots, T$. The index $j = 1, \dots, J_t$ counts the implied volatilities observed on day t . Let $X_{t,j}$ be a two-dimensional variable containing (forward) moneyness $\kappa_{t,j}$ and time to maturity $\tau_{t,j}$. We define the moneyness $\kappa_{t,j} \stackrel{\text{def}}{=} K_{t,j}/F_{\tau_{t,j}}$, where $K_{t,j}$ is a strike and $F_{\tau_{t,j}}$ the forward price of

the underlying asset at time t . The SFM regresses $Y_{t,j}$ on $X_{t,j}$ by:

$$Y_{t,j} = \sum_{l=0}^L Z_{t,l} m_l(X_{t,j}) + \varepsilon_{t,j}, \quad (5)$$

where m_l ($l = 1, \dots, L$) are nonparametric components and the $Z_{t,l}$ form a latent factor series depending on time t . The estimation error is denoted by $\varepsilon_{t,j}$. The basis functions m_0, \dots, m_L are constant in time, while the dynamic propagation of the IVS is modelled by the time varying weights $Z_{t,l}$.

The estimation procedure is based on minimizing the following least squares criterion ($\widehat{Z}_{t,0} \equiv 1$ for identification):

$$\sum_{t=1}^T \sum_{j=1}^{J_t} \int \left\{ Y_{t,j} - \sum_{l=0}^L \widehat{Z}_{t,l} \widehat{m}_l(u) \right\}^2 K_h(u - X_{t,j}) du, \quad (6)$$

where K_h denotes a two-dimensional kernel function. A possible choice for a two-dimensional kernel is a product of one-dimensional kernels $K_h(u) = k_{h_1}(u_1) \times k_{h_2}(u_2)$, where $h = (h_1, h_2)^\top$ are bandwidths and $k_h(v) = h^{-1}k(h^{-1}v)$ is a one dimensional kernel function. The minimization procedure searches across all functions $\widehat{m}_l : \mathbb{R}^2 \rightarrow \mathbb{R}$ ($l = 0, \dots, L$) and time series $\widehat{Z}_{t,l} \in \mathbb{R}$ ($t = 1, \dots, T; l = 1, \dots, L$). Details concerning the estimation algorithm can be found in Fengler et al. (2007) and Park et al. (2009). In the final step of the procedure one orthogonalizes the functions $\widehat{m}_1, \dots, \widehat{m}_L$ and orders them with respect to the variance explained. As a consequence the largest portion of variance is explained by the quantity $\widehat{Z}_{t,1}\widehat{m}_1$ and the second largest by $\widehat{Z}_{t,1}\widehat{m}_1 + \widehat{Z}_{t,2}\widehat{m}_2$ and so forth.

In order to illustrate the decomposition of the IVS dynamics achieved by the SFM we present in Figure 1 the results on DAX option data from January 3rd, 2000 till June 30th, 2004. The figure presents the estimated $\widehat{Z}_{t,l}$ time series in the upper panel and the estimates of the basis functions in the lower panel. The function \widehat{m}_0 is not presented to save space. It has no effect on the dynamics of the IVS but has to be included to set the correct level of the surface. The function \widehat{m}_1 is relatively flat and corresponds to the most important shocks. Changes in $\widehat{Z}_{t,1}$ result in up-and-down type of movements of the whole surface, but the deviations from a

flat basis function give different weight for each maturity-moneyness location. This effect is illustrated in Figure 2, where we plot several surfaces and one particular smile with different values of $\widehat{Z}_{t,1}$. The second factor function can be interpreted as a tilting of the smile. This can be inferred from the shape of \widehat{m}_2 and its influence on the IVS in the plots. The variation in $\widehat{Z}_{t,2}$ results in changing the slope of the smile by making it steeper or flatter while keeping roughly the same implied volatility levels.

We finally remark that the SFM has spurred further research on IVS dynamics and beyond. Brüggemann et al. (2008) study the statistical properties of the estimated factor series using a vector autoregressive framework and analyze the associated movements of macroeconomic variables. Giacomini and Härdle (2008) apply the modelling idea for an explanation of the dynamics of risk neutral densities. The CO₂ allowance term structure is studied in Trück et al. (2006) and electricity forward curves in Borak and Weron (2009).

3 Hedging Framework

Dynamic hedging of the asset V , in our case the reverse barrier option, is based on frequent adjustments of the hedge portfolio. This hedging strategy requires to construct a portfolio which is to first (or higher) order neutral to the relevant risk factors. Apart from standard delta hedging, a successful strategy requires hedging the vega, and possibly higher order greeks as pointed out by Ederington and Guan (2007).

For the LV framework Engelmann et al. (2006) study delta, delta-vega and delta-vega-vanna hedges. One knock-out option is hedged with the underlying asset and a set of plain vanilla options. Let the value of the barrier option be denoted by V and let HP_1 and HP_2 be portfolios of plain vanilla options. The corresponding hedge ratios are then given by solving

$$\begin{pmatrix} 1 & \frac{\partial HP_1}{\partial S} & \frac{\partial HP_2}{\partial S} \\ 0 & \frac{\partial HP_1}{\partial \sigma} & \frac{\partial HP_2}{\partial \sigma} \\ 0 & \frac{\partial^2 HP_1}{\partial \sigma \partial S} & \frac{\partial^2 HP_2}{\partial \sigma \partial S} \end{pmatrix} \cdot \begin{pmatrix} a_0 \\ a_1 \\ a_2 \end{pmatrix} = \begin{pmatrix} \frac{\partial V}{\partial S} \\ \frac{\partial V}{\partial \sigma} \\ \frac{\partial^2 V}{\partial \sigma \partial S} \end{pmatrix}. \quad (7)$$

Equation (7) reflects the full delta-vega-vanna hedge. Putting $a_2 = 0$ reduces (7) to the

delta-vega hedge and $a_1 = a_2 = 0$ to the pure delta hedge. Since good hedges have a large exposure to the risk factors to be hedged, one could use an at-the-money plain vanilla option for the HP_1 and for HP_2 a risk reversal. A risk reversal is a combination of a long out-of-the-money call and a short out-of-the-money put (or vice versa).

In order to compute the sensitivities one reprices the option under different scenarios and computes the greeks by a finite difference quotient. Following Engelmann et al. (2006), we make a sticky strike assumption for our greeks, i.e. the IVS remains constant in strikes. Vega and vanna are computed shifting the IVS in a parallel fashion. To be more specific, we compute

$$\frac{\partial V}{\partial S} \stackrel{\text{def}}{\approx} \frac{V(S + \Delta S, \hat{\sigma}) - V(S - \Delta S, \hat{\sigma})}{2\Delta S}, \quad (8)$$

$$\frac{\partial V}{\partial \hat{\sigma}} \stackrel{\text{def}}{\approx} \frac{V(S, \hat{\sigma} + \Delta \hat{\sigma}) - V(S, \hat{\sigma} - \Delta \hat{\sigma})}{2\Delta \hat{\sigma}}, \quad (9)$$

$$\frac{\partial^2 V}{\partial S \partial \hat{\sigma}} \stackrel{\text{def}}{\approx} \frac{\{V(S + \Delta S, \hat{\sigma} + \Delta \hat{\sigma}) - V(S + \Delta S, \hat{\sigma}) - V(S - \Delta S, \hat{\sigma} + \Delta \hat{\sigma}) + V(S - \Delta S, \hat{\sigma})\}}{2\Delta S \Delta \hat{\sigma}}. \quad (10)$$

With small abuse of notation $V(S, \hat{\sigma})$ denotes here the price obtained with spot S and IVS $\hat{\sigma}$, where we omit its arguments for simplicity. $\hat{\sigma} + \Delta \hat{\sigma}$ means the parallel shift of the whole surface.

It is empirically widely confirmed that parallel shifts are the most prevalent movements of the IVS. It would be misleading, however, to conclude from this observation that other types of surface variations do only negligibly influence the prices of exotic derivatives, such as barrier options. Contrariwise a higher slope leads to a smaller price of an in-the-money down-and-out put. Consider an artificial example of two one year down-and-out put with strike 110, barrier 80 at the current spot level of 100. The first option is priced with the IVS observed on January 3rd, 2000 and the second one on January 2nd, 2001. Figure 3 shows the surfaces of these days. The LV prices of these options are 1.91% and 2.37% respectively (in percentage of the spot price), which is quite a difference. From the upper panel of Figure 1 one observes that the level related factor assumes similar values on these days, while the slope factor differs significantly. This price discrepancy stems mainly from the slope effect,

which is an exposure not directly hedged in traditional approaches. Our procedure will hedge such volatility shocks.

In our hedging framework we define new sensitivities with respect to the variation of the (log)-IVS, which we call ζ -greeks. Based on the results discussed in Section 2.2, the ζ_1 -greek ($\partial/\partial Z_{t,1}$) reflects an adjusted up-and-down shift, while the ζ_2 -greek ($\partial/\partial Z_{t,2}$) corresponds to the slope effect. Similarly to (7) we obtain the hedge ratios by

$$\begin{pmatrix} 1 & \frac{\partial HP_1}{\partial S} & \frac{\partial HP_2}{\partial S} \\ 0 & \frac{\partial HP_1}{\partial Z_{t,1}} & \frac{\partial HP_2}{\partial Z_{t,1}} \\ 0 & \frac{\partial HP_1}{\partial Z_{t,2}} & \frac{\partial HP_2}{\partial Z_{t,2}} \end{pmatrix} \cdot \begin{pmatrix} a_0 \\ a_1 \\ a_2 \end{pmatrix} = \begin{pmatrix} \frac{\partial V}{\partial S} \\ \frac{\partial V}{\partial Z_{t,1}} \\ \frac{\partial V}{\partial Z_{t,2}} \end{pmatrix}. \quad (11)$$

We call the full setting a $\zeta_1\zeta_2$ -hedge, the reduced one with $a_2 = 0$ a ζ_1 -hedge. As in the traditional hedge we use an at-the-money plain vanilla for HP_1 , again due to the high vega. For HP_2 , we employ risk reversals because they primarily respond to changes in the wings of the IVS. Moreover, by selecting appropriate strikes it can even be set up in a vega-neutral, i.e. ζ_1 -neutral, way.

We calculate the ζ -greeks by means of a difference quotient. As pricing input for the barrier options we do not use the estimate of the IVS obtained by the SFM, as it is necessarily subject to an estimation error. Instead, in order to avoid mispricings, we use the truly observed ones. Thus, by the definition of the ζ -greeks, the approximations are given by

$$\frac{\partial V}{\partial Z_{t,l}} \stackrel{\text{def}}{\approx} \frac{V(S, \hat{\sigma} \exp(\Delta Z_{t,l} \hat{m}_l)) - V(S, \hat{\sigma} \exp(-\Delta Z_{t,l} \hat{m}_l))}{2\Delta Z_{t,l}}. \quad (12)$$

In the practical implementation of (12) one faces a couple of numerical issues, which need to be addressed. First, the size of the $\Delta Z_{t,l}$ has to be chosen. An increment too small or too large can distort the meaning of the greeks. Moreover it cannot be unique for all $Z_{t,l}$, since the shift size depends on the basis functions \hat{m}_l and on the IVS on a particular day. Therefore we choose for each t a $\Delta Z_{t,l}$ such that the (absolute) mean upward (downward) shift amounts approximately to one volatility-point. Note that we do *not* use $\hat{Z}_{t,l}$ for these perturbations. Another challenge is an accurate calculation of the barrier greeks. To reduce

numerical errors we employ a constant grid in the pricing algorithm for calculating the ζ -greeks. Furthermore, the IVS $\hat{\sigma}$ needs to be arbitrage-free. However, the shifted surfaces do not necessarily possess this property. We thus additionally check no-arbitrage conditions before calculating the ζ -greeks and apply an algorithm due to Fengler (2008) in case of violations. This method estimates the option price function by means of a natural smoothing spline under no-arbitrage constraints, i.e. under convexity, monotonicity and bounds on the price function and on the first order strike derivatives. The resulting estimate is then converted back to implied volatility. The algorithm is not applied when computing vega and vanna since parallel shifts do typically not result into arbitrage violations.

The aforementioned greeks are demonstrated in Figure 4 for the down-and-out put with half a year to expiry. The plot displays the greeks as a function of spot and keeps other characteristics of the barrier option unchanged. It has to be noted that the SFM, i.e. $\hat{Z}_{t,l}$ and \hat{m}_l , can only be identified up to sign. The sign of the ζ -greeks therefore has no particular meaning. Hence vega and ζ_1 display similar patterns. For the spot values close to the barrier level vega is negative and approaches zero as it becomes a delta product. For out-of-the money options vega is positive since the option then resembles a plain vanilla contract. A similar behavior is observed for ζ_2 and vanna, but the vanna is discontinuous at the barrier as it is derived from the delta.

4 Empirical Results

4.1 Data

The data set covers DAX index options traded at the EUREX from January 3rd, 2000 till June 30th, 2004 which give 1135 trading days. We use settlement prices, which are prices published by the EUREX based on the last intra-day trades. The DAX index is a capital weighted performance index comprising 30 German blue chips. Since dividends less corporate tax are reinvested into the index, they do not need to be taken into account for option valuation.

We preprocess the data by eliminating implied volatilities bigger than 80% and maturities

smaller than 10 days. Arbitrage violations in the option data are removed by the arbitrage free smoothing procedure described in Fengler (2008). After smoothing, the data are converted into a regular grid of moneyness and time to maturities. For option pricing, the zero rates from EURIBOR quotes are linearly interpolated, see Dumas et al. (1998) for this practice.

4.2 Experimental Design

In our empirical study we assume no transaction costs, no restrictions on short selling and the possibility of trading each asset at arbitrary size. Each security is priced using the LV model calibrated to daily market data. We implement the hedging strategies described in Section 3, i.e. we focus exclusively on volatility and spot risks, leaving other risks like interest rate exposure unhedged.

In the first step of our experiment we estimate the SFM. As kernel function we use a product quartic kernel, where $k(u) = 15/16(1 - u^2)^2$ for $|u| < 1$ and 0 otherwise. For a data driven bandwidth choice and the model size selection, we refer to Fengler et al. (2007). The basic idea is to estimate the model for different combinations of L and h and compare various information criteria. For the moneyness direction we finally use a bandwidth of 0.04, but we slightly oversmooth the surfaces in the time to maturity direction in order to reduce numerical errors for the subsequent price computations. More precisely, we use a local bandwidth modelled by an arctangent function which increases monotonously from 0.02 to 0.15 (expressed in years). Since in the hedging procedure only two main factors are included, we set $L = 2$. With this choice the model describes sufficiently well the IVS dynamics, since the measure of explained variation is close to 98%.

For each day up to one year before the last observation date in the sample, a long position in the reverse barrier option is created. This is to evaluate all initiated hedges at market prices within the sample. We use up-and-out calls with strikes at 80% of the spot and barriers at 140% and down-and-out put with strikes at 80% and barriers at 110%. These specifications correspond to typically traded contracts. Based on the calibrated LV model, ζ -greeks, delta, vega and vanna are calculated and the hedging strategies as described in Section 3 are set up. We concentrate on vega, vanna, ζ_1 and $\zeta_1\zeta_2$ strategies since the pure delta hedge is of

inferior quality. As HP_1 we use at-the-money puts for the up-and-out calls and at-the-money calls for the down-and-out puts. The risk reversal are structured by taking 80% and 120% strikes of the current spot.

Positions that have not knocked are updated on a daily basis. This choice is motivated by the results of Engelmann et al. (2006) who do not obtain different rankings of the strategies for other re-balancing frequencies. For each day we calculate the greeks to solve (7) and (11) and adjust the hedge ratios a_0, a_1, a_2 . The hedges are financed from the cash account and if the barrier is breached or the barrier option expires we unwind the hedge and record the hedging error. All positions are traded at market prices. In case of a knock-out event, the hedging error pays or earns interest until expiry in order to render the results comparable. Also the cash account bears interest or is financed at the riskless short rate of the concurrent trading day. Summing up, we have a collection of hedging errors for the two types of barrier options with four different hedging strategies for each of them.

One could object that the experimental design suffers from an in-sample problem, since the SFM is estimated on the same data set as the hedging experiment. It is however a common finding in the empirical literature, either on interest rates or on the IVS, that eigenvectors or eigenfunctions are remarkably stable across time. Formal tests on IVS data between the years 1995 to 2001 confirming this hypothesis are provided by Fenger (2005, Chapter 5.2.3). Even if we made use of a training-sample, we would therefore recover very similar factor functions. Thus the issue will not seriously affect the results.

4.3 Results

For evaluating the performance we use a pool of 885 hedging errors (1135 trading days less 250 days, since products issued thereafter would not expire within the sample). In order to make them comparable we normalize by the spot price at the time when the hedge is initiated. This normalization is common in practice and is meant to remove the dependence from the underlying's level. Another normalizing factor could be the option price itself, but since the risk reversal has a market price close to zero, measuring errors with respect to the spot appears to be more natural.

The aim of hedging is to replicate the payoff of the option. In the ideal case the hedge portfolio should have zero variance and zero mean, but for obvious reasons this cannot be realized in practice. Our aim is to give a comparative analysis of the hedging error distributions in order to check how the volatility factors affect the hedging performance. We use traditional descriptive statistics to assess the location and dispersion of the errors. Clearly, a superior method would keep these quantities close to zero in absolute terms.

The empirical results are summarized in Tables 2 and 3 for up-and-out calls and down-and-out puts respectively. We present the minimum, maximum, mean, median, standard deviation, and the absolute deviation around the median. The terminal hedging error distributions are given in the rows marked with a ‘0’. As can be inferred from the tables, the center of all distributions is located around zero, with means slightly below zero for the up-and-out calls and slightly above zero for the down-and-out puts. Thus the different hedges are hardly distinguishable in terms of the center of the distribution. This finding corresponds to our expectations: the volatility risk is removed, both for the vega and the ζ_1 -hedges, and vanna and $\zeta_1\zeta_2$ -hedges do not add any additional drift, since they are almost costless.

For evaluating the dispersion of the hedging errors we focus on the standard deviation and the absolute deviation around the median (madev.). The first observation is that hedges relying on higher order greeks tend to exhibit lower variance. In case of the down-and-out puts the vanna hedge has a slightly smaller dispersion than the $\zeta_1\zeta_2$ -hedge, and the traditional vega hedge performs very similar to the ζ_1 -hedge. For the up-and-out calls the ranking is reversed: the standard hedges are clearly outperformed by the factor hedges. How can this asymmetry be explained and how is the quality of the factor hedges to be judged?

There are two major sources of bias in the hedging strategies due to the behavior of the underlying. Observe that during the analyzed time period the DAX had a downward trend: 81% out of the down-and-out put options knocked out, but only 10% of the up-and-out call options, while 5% of the puts and 39% calls expired in-the-money, see Table 1. As a first issue consider the huge amount of up-and-out calls ending in-the-money. This gives rise to what is known among practitioners as ‘theta risk’. For explanation reconsider the case in Section 3, where we demonstrated that the prices for one-year down-and-out puts with a strike of 110% and barrier at 80% were less than 3% in the two scenarios. In contrast, when the put ends in-the-money it will pay out up to 30%. Consequently, the value of an in-the-money reverse

barrier option increases sharply the nearer the expiry date draws (i.e. has a strong theta), rendering it more and more difficult for traders to earn the payoff by trading the gamma. Theta risk can thus lead to a more dispersed error distribution. A second issue is gap risk. We do not unwind the hedges at the barriers, but at the observed spots, since this is the more realistic scenario in practice. When a barrier is breached, one still owns the hedge and incurs unbalanced gains or losses. Again this leads to a more dispersed hedging error distribution. As is clear from Table 1, theta risk is dominating the risk in case of the calls and gap risk in case of the puts.

To receive a deeper insight, we refer once more to Tables 2 and 3. We report the statistics of the hedging experiment stopped at 1 day, 5 days and 25 days before the expiry. As is seen the dispersion measures increase the nearer expiry draws, and the distributions become less skewed and less heavy-tailed, while the location measures prove to remain stable. In terms of dispersion the relative order of the hedging strategies across the two products remains the same: for the down-and-out puts the strategies are comparable, while factor hedging remains superior for the up-and-out calls. This finding is confirmed in Figure 5, which displays the standard deviations of the hedging errors as a function through the options' life time. It is intuitive to expect this function to increase. Moreover there is a sharp jump just before the expiry date contributing a large portion of the overall cumulative hedging error in particular for the up-and-out calls. All these observations highlight the importance of the expiry effect relative to gap risk when interpreting the data.

We overall conclude two main findings. First, factor hedging is at least of similar quality as traditional hedging approaches. In particular the hedging efficiency does not deteriorate. This is a reassuring result given the huge computational effort that must be spent and that could easily come at the costs of accuracy. This result is obtained when the barrier options expire worthless or knock out early in life time. Second, when the option needs to be hedged till expiry and ends in-the-money, the factor hedging approach dominates clearly. From a trader's perspective the first situation is the 'easy one' unless the knock-out occurs close to expiry. The second one is much more intricate, because the intrinsic value needs to be earned. This is a strong case for volatility factor hedging.

5 Conclusion

We provide an empirical study on hedging reverse barrier options in the local volatility model. The main focus of this study is on risk factors arising from a decomposition of the dynamic behavior of the implied volatility surface, which are identified with a flexible semiparametric technique. The hedging framework is constructed as a natural extension to traditional vega hedging, where the sensitivity is measured with respect to the more complex surface movements.

Our empirical investigation shows that hedging higher order risk with risk reversals brings improvements to hedging with at-the-money plain vanillas only. This is consistent across the vanna hedge and the more complex factor based hedges, thus confirming evidence of Ederington and Guan (2007). Intuitively the vega hedge resembles a single factor based hedge since the first dynamic factor corresponds to a parallel type of shift. Adding a vanna hedge or another factor to the portfolio removes similar risks as can be inferred from the comparable hedging performance.

Measured in terms of the hedging error variance, factor hedging performs at least as good as the corresponding vega and vanna hedges, in certain cases it is superior. As is confirmed by hedging up-and-out call options and down-and-out put options, the first case occurs when options knock out early in life time or expire worthless, while the second occurs when the options need to be hedged up to expiry and end in-the-money. This evidence is present not only in the terminal hedging errors but also through the option's life time. From a trader's perspective the second case is the more interesting, making factor hedging a powerful alternative to traditional hedging.

These findings, however, are not necessarily similar for other complex derivatives sensitive to IVS movements, such as cliquets or long-dated forward starting options. Also a portfolio context may yield different findings. In particular, when a book of options contains assets with several maturities it could be beneficial to consider additional factors, such as those related to the term structure of the IVS. This exposure can be hedged by constructing the corresponding calendar spreads. Another application in a portfolio context could be stress test scenarios based on the volatility factors. This would provide a good understanding of

the volatility exposure of the portfolio. We leave these issues to future research.

References

- Alexander, C. (2001). Principles of the skew. *RISK*, 14(1):S29–S32.
- Andersen, L. B. G., Andreasen, J., and Eliezer, D. (2002). Static replication of barrier options: Some general results. *Journal of Computational Finance*, 5(4):1–25.
- Andersen, L. B. G. and Brotherton-Ratcliffe, R. (1997). The equity option volatility smile: An implicit finite-difference approach. *Journal of Computational Finance*, 1(2):5–37.
- Bates, D. S. (1996). Jumps and stochastic volatility: Exchange rate processes implicit in Deutsche Mark options. *Review of Financial Studies*, 9:69–107.
- Bergomi, L. (2005). Smile dynamics II. *RISK*, 18(10):67–73.
- Borak, S. and Weron, R. (2009). A semiparametric factor model for electricity forward curve dynamics. *The Journal of Energy Markets*, 1(3).
- Bouchouev, I. and Isakov, V. (1999). Uniqueness, stability and numerical methods for the inverse problem that arises in financial markets. *Inverse Problems*, 15:R95–R116.
- Brüggemann, R., Härdle, W., Mungo, J., and Trenkler, C. (2008). VAR modeling for dynamic loadings driving volatility strings. *Journal of Financial Econometrics*, 6:361–381.
- Carr, P. and Chou, A. (1997). Breaking Barriers. *Risk Magazine*, 10:139–145.
- Carr, P., Ellis, K., and Gupta, V. (1998). Static hedging of exotic options. *Journal of Finance*, 53(3):1165–1190.
- Carr, P., Geman, H., Madan, D., and Yor, M. (2003). Stochastic volatility for Lévy processes. *Mathematical Finance*, 13:345–382.
- Coleman, T. F., Kim, Y., Li, Y., and Verma, A. (2001). Dynamic hedging with a deterministic local volatility function model. *Journal of Risk*, 4(1):63–89.
- Cont, R. and da Fonseca, J. (2002). The dynamics of implied volatility surfaces. *Quantitative Finance*, 2(1):45–60.

- Derman, E., Ergener, D., and Kani, I. (1995). Static options replication. *Journal of Derivatives*, 2(4):78–95.
- Derman, E. and Kani, I. (1994). Riding on a smile. *RISK*, 7(2):32–39.
- Diebold, F., Ji, L., and Li, C. (2006). A three-factor yield curve model: Non-affine structure, systematic risk sources, and generalized duration. In Klein, L., editor, *Long-Run Growth and Short-Run Stabilization: Essays in Memory of Albert Ando*. Edward Elgar, Cheltenham, U.K.
- Dumas, B., Fleming, J., and Whaley, R. E. (1998). Implied volatility functions: Empirical tests. *Journal of Finance*, 80(6):2059–2106.
- Dupire, B. (1994). Pricing with a smile. *RISK*, 7(1):18–20.
- Ederington, L. and Guan, W. (2007). Higher order greeks. *Journal of Derivatives*, 14:7–34.
- Engelmann, B., Fengler, M., Nalholm, M., and Schwendner, P. (2007). Static versus Dynamic Hedges: An Empirical Comparison for Barrier Options. *Review of Derivatives Research*, 9(3):239–264.
- Engelmann, B., Fengler, M., and Schwendner, P. (2006). Better than its reputation: An empirical hedging analysis of the local volatility model for barrier options. Working paper, Available at SRRN.
- Fengler, M. R. (2005). *Semiparametric Modeling of Implied Volatility*. Lecture Notes in Finance. Springer-Verlag, Berlin, Heidelberg.
- Fengler, M. R. (2008). Arbitrage-free smoothing of the implied volatility surface. *Quantitative Finance*. Forthcoming.
- Fengler, M. R., Härdle, W., and Mammen, E. (2007). A semiparametric factor model for implied volatility surface dynamics. *Journal of Financial Econometrics*, 5(2):189–218.
- Fengler, M. R., Härdle, W., and Villa, C. (2003). The dynamics of implied volatilities: A common principle components approach. *Review of Derivatives Research*, 6:179–202.

- Giacomini, E. and Härdle, W. (2008). Dynamic Semiparametric Factor Models in Pricing Kernels Estimation. In Dabo-Niang, S. and Ferraty, F., editors, *Functional and Operatorial Statistics*, pages 181–187. Physica-Verlag HD.
- Hafner, R. (2004). *Stochastic Implied Volatility*. Springer, Berlin.
- Hagan, P., Kumar, D., Lesniewski, A., and Woodward, D. (2002). Managing smile risk. *Wilmott magazine*, 1:84–108.
- 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.
- Hirsa, A., Courtadon, G., and Madan, D. (2003). The effect of model risk on the valuation of barrier options. *Journal of Risk Finance*, 4:47–55.
- Hull, J. and White, A. (1987). The pricing of options on assets with stochastic volatilities. *Journal of Finance*, 42:281–300.
- Hull, J. C. and Suo, W. (2002). A methodology for assessing model risk and its application to the implied volatility function model. *Journal of Financial and Quantitative Analysis*, 37(2):297–318.
- Kou, S. G. (2002). A jump-diffusion model for option pricing. *Management Science*, 48:1086–1101.
- Maruhn, J., Nalholm, M., and Fengler, M. R. (2008). Empirically robust static uncertain skew hedges for reverse barrier options. Working paper.
- McIntyre, M. L. (2001). Performance of Dupire’s implied diffusion approach under sparse and incomplete data. *Journal of Computational Finance*, 4(4):33–84.
- Merton, R. C. (1976). Option pricing when underlying stock returns are discontinuous. *Journal of Financial Economics*, 3:125–144.
- Nalholm, M. and Poulsen, R. (2006a). Static hedging and model risk for barrier options. *Journal of Future Markets*, 26:449–463.
- Nalholm, M. and Poulsen, R. (2006b). Static hedging of barrier options under general asset dynamics: Unification and application. *Journal of Derivatives*, 13:46–60.

- Park, B., Mammen, E., Härdle, W., and Borak, S. (2009). Time Series Modelling with Semiparametric Factor Dynamics. *Journal of the American Statistical Association*. Forthcoming.
- Rubinstein, M. (1994). Implied binomial trees. *Journal of Finance*, 49:771–818.
- Rubinstein, M. and Reiner, E. (1991). Breaking down the barrier. *RISK*, 4(9):28–35.
- Skiadopoulos, G., Hodges, S., and Clewlow, L. (1999). The dynamics of the S&P 500 implied volatility surface. *Review of Derivatives Research*, 3:263–282.
- Stein, E. M. and Stein, J. C. (1991). Stock price distributions with stochastic volatility: An analytic approach. *Review of Financial Studies*, 4:727–752.
- Taleb, N. (1997). *Dynamic Hedging: Managing Vanilla and Exotic Options*. John Wiley & Sons.
- Tompkins, R. (2002). Static versus dynamic hedging of exotic option: An evaluation of hedge performance via simulation. *The Journal of Risk Finance*, 3:6–34.
- Trück, S., Borak, S., Härdle, W., and Weron, R. (2006). Convenience yields for CO₂ emission allowance futures contracts. Discussion Paper 2006-076, SFB 649, Humboldt-Universität zu Berlin.
- Vähämaa, S. (2004). Delta hedging with the smile. *Financial Markets and Portfolio Management*, 18(3):241–255.




Figure 1: *The estimates of the SFM obtained from IVS data from January 3rd, 2000 till June 30th, 2004 for $L = 2$. Upper panel: estimated latent factor series \hat{Z}_1 and \hat{Z}_2 . Lower panel: estimates of \hat{m}_1 , the non-uniform up-and-down shift, and \hat{m}_2 , the slope risk.*

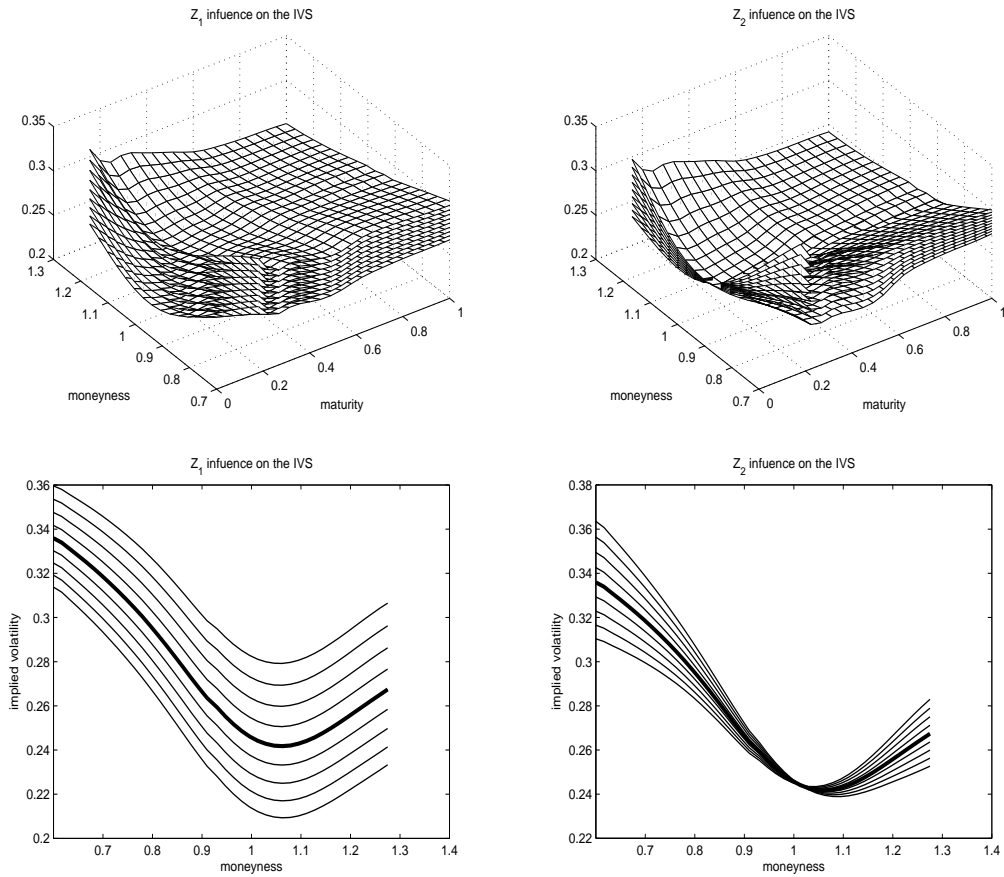


Figure 2: Impact of \widehat{Z}_1 and \widehat{Z}_2 on the IVS. Shocks in \widehat{Z}_1 trigger up-and-down movements while shocks in \widehat{Z}_2 tilt the smile around at-the-money point. Upper panel: a visualization of the shocks for the entire surface. Lower panel: the impact presented on one particular smile.




Figure 3: *IVS* observed on January 3rd, 2000 (the steeper surface) and January 2nd, 2001 (the flatter one). DAX levels on these days were 6751 and 6290 respectively.

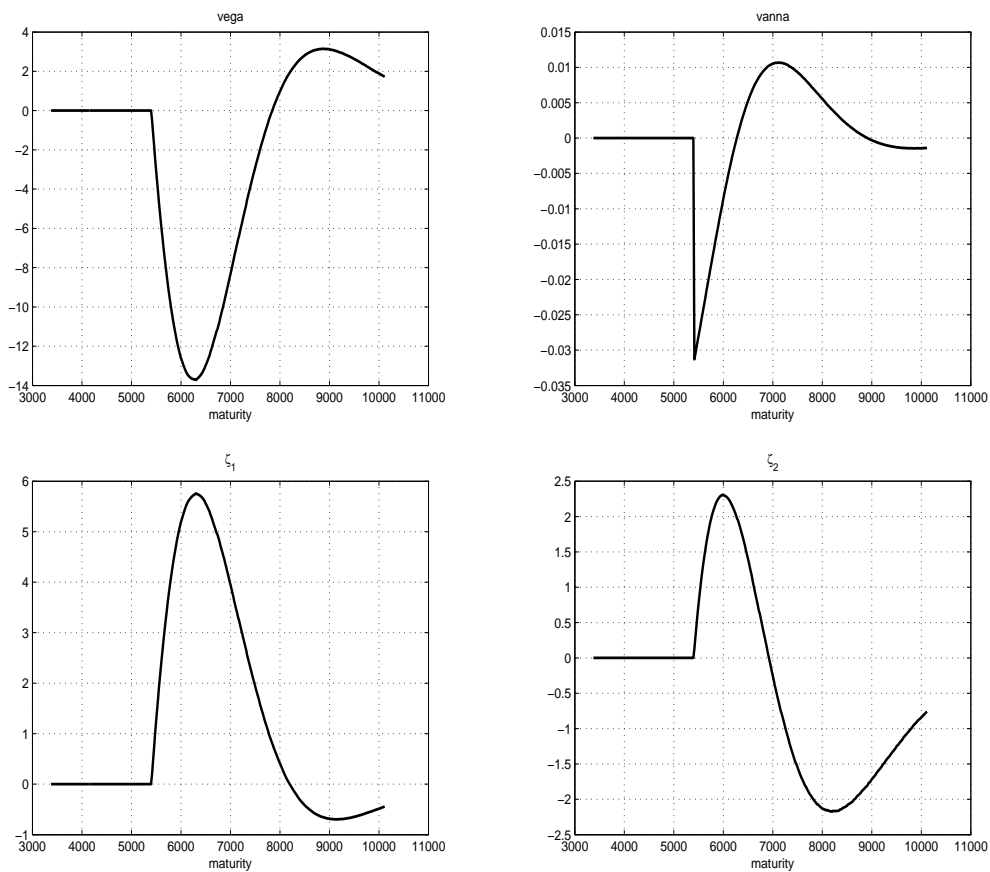


Figure 4: Greeks for a down-and-out put option with maturity 0.5 years with barrier 5400 strike 7425 as a function of the spot. Upper left panel: vega. Upper right panel: vanna. Lower right panel: ζ_1 . Lower right panel: ζ_2




Figure 5: *Standard deviations of the hedging errors as a function of time from option issuance. Solid lines represent the factor hedging methods motivated by the SFM. Dashed lines represent the vega and vanna hedges. Upper panel: up-and-out call. Lower panel: down-and-out put.*

option type	barrier	strike	knock-outs	in-the-money
up-and-out call	140%	80%	10%	39%
down-and-out put	80%	110%	81%	5%

Table 1: *Characteristics of the analyzed barrier options. Strikes and barriers are in percentage of spot at issuance. The column ‘knock-outs’ refers to the contracts that breached the barrier and ‘in-the-money’ to those yielding a positive payoff at expiry.*

	days	min	max	mean	median	std.	madev.	skew.	kurt.
vega	0	-0.1038	0.5813	-0.0165	-0.0175	0.0413	0.0209	7.2801	97.60
	1	-0.1038	0.2581	-0.0172	-0.0174	0.0314	0.0199	2.3526	19.32
	5	-0.1037	0.0970	-0.0181	-0.0169	0.0260	0.0183	0.3636	4.91
	25	-0.0827	0.0649	-0.0174	-0.0164	0.0249	0.0178	0.0587	3.74
ζ_1	0	-0.0752	0.5768	-0.0118	-0.0136	0.0387	0.0183	8.4877	119.04
	1	-0.0751	0.2332	-0.0125	-0.0134	0.0279	0.0172	2.8026	22.27
	5	-0.0749	0.0755	-0.0134	-0.0121	0.0216	0.0155	0.2846	4.60
	25	-0.0761	0.0573	-0.0134	-0.0127	0.0215	0.0161	0.0343	3.55
vanna	0	-0.1340	0.5310	-0.0081	-0.0138	0.0345	0.0151	8.6289	124.62
	1	-0.1340	0.1842	-0.0089	-0.0136	0.0239	0.0140	2.1325	17.40
	5	-0.1339	0.0807	-0.0099	-0.0131	0.0187	0.0121	0.2157	9.54
	25	-0.0582	0.0772	-0.0096	-0.0141	0.0173	0.0118	1.3367	6.22
$\zeta_1\zeta_2$	0	-0.0830	0.5684	-0.0066	-0.0119	0.0345	0.0137	10.6470	161.00
	1	-0.0829	0.2091	-0.0073	-0.0117	0.0226	0.0126	4.0718	31.51
	5	-0.0829	0.0710	-0.0083	-0.0113	0.0157	0.0106	1.3447	7.12
	25	-0.0370	0.0629	-0.0086	-0.0118	0.0152	0.0108	1.3559	5.72

Table 2: Hedging error distributions of the up-and-out calls. Given are descriptive statistics for the various hedging strategies. The rows present the statistics at 0, 1, 5 and 25 days before expiration.

	days	min	max	mean	median	std.	madev.	skew.	kurt.
vega	0	-0.0264	0.2799	0.0058	-0.0004	0.0213	0.0105	5.4903	51.91
	1	-0.0756	0.1172	0.0050	-0.0004	0.0166	0.0098	2.4531	12.38
	5	-0.0187	0.0882	0.0041	-0.0008	0.0147	0.0090	2.4682	10.92
	25	-0.0186	0.0749	0.0038	-0.0004	0.0124	0.0083	2.0267	8.73
ζ_1	0	-0.0210	0.2808	0.0080	0.0016	0.0214	0.0107	5.4775	51.59
	1	-0.0702	0.1215	0.0072	0.0015	0.0167	0.0100	2.4501	12.14
	5	-0.0137	0.0882	0.0063	0.0013	0.0147	0.0091	2.4112	10.29
	25	-0.0113	0.0798	0.0059	0.0014	0.0127	0.0085	2.0632	8.54
vanna	0	-0.0608	0.2072	0.0022	-0.0016	0.0178	0.0081	5.7326	53.32
	1	-0.0955	0.1309	0.0014	-0.0016	0.0137	0.0074	2.7735	23.76
	5	-0.0323	0.0649	0.0006	-0.0018	0.0114	0.0069	1.9306	9.98
	25	-0.0205	0.0582	0.0004	-0.0016	0.0093	0.0059	2.0273	9.23
$\zeta_1\zeta_2$	0	-0.0332	0.2676	0.0065	0.0008	0.0196	0.0092	6.0258	60.47
	1	-0.0824	0.1146	0.0057	0.0008	0.0149	0.0085	2.4463	14.06
	5	-0.0234	0.0774	0.0048	0.0007	0.0130	0.0079	2.4134	10.42
	25	-0.0121	0.0727	0.0045	0.0010	0.0110	0.0069	2.3842	10.27

Table 3: Hedging error distributions of the down-and-out puts. Given are descriptive statistics for the various hedging strategies. The rows present the statistics at 0, 1, 5 and 25 days before the expiration.

Adaptive pointwise estimation in time-inhomogeneous conditional heteroscedasticity models

P. ČÍŽEK[†], W. HÄRDLE[‡] AND V. SPOKOINY[§]

[†]*Department of Econometrics & OR, Tilburg University, P.O. Box 90153, 5000LE Tilburg, The Netherlands*

E-mail: P.Cizek@uvt.nl

[‡]*Humboldt-Universität zu Berlin and CASE, Spandauerstrasse 1, 10178 Berlin, Germany*

E-mail: haerdle@wiwi.hu-berlin.de

[§]*Weierstrass-Institute, Humboldt-Universität zu Berlin and CASE, Mohrenstrasse 39, 10117 Berlin, Germany*

E-mail: spokoiny@wias-berlin.de

First version received: April 2008; final version accepted: April 2009

Summary This paper offers a new method for estimation and forecasting of the volatility of financial time series when the stationarity assumption is violated. Our general, local parametric approach particularly applies to general varying-coefficient parametric models, such as GARCH, whose coefficients may arbitrarily vary with time. Global parametric, smooth transition and change-point models are special cases. The method is based on an adaptive pointwise selection of the largest interval of homogeneity with a given right-end point by a local change-point analysis. We construct locally adaptive estimates that can perform this task and investigate them both from the theoretical point of view and by Monte Carlo simulations. In the particular case of GARCH estimation, the proposed method is applied to stock-index series and is shown to outperform the standard parametric GARCH model.

Keywords: *Adaptive pointwise estimation, Autoregressive models, Conditional heteroscedasticity models, Local time-homogeneity.*

1. INTRODUCTION

A growing amount of econometrical and statistical research is devoted to modelling financial time series and their volatility, which measures dispersion at a point in time (i.e. conditional variance). Although many economies and financial markets have been recently experiencing many shorter and longer periods of instability or uncertainty such as the Asian crisis (1997), the Russian crisis (1998), the start of the European currency (1999), the ‘dot-Com’ technology-bubble crash (2000–02) or the terrorist attacks (September, 2001), the war in Iraq (2003) and the current global recession (2008), mostly used econometric models are based on the assumption of time homogeneity. This includes linear and non-linear autoregressive (AR) and moving-average models and conditional heteroscedasticity (CH) models such as ARCH (Engel, 1982)

and GARCH (Bollerslev, 1986), stochastic volatility models (Taylor, 1986), as well as their combinations such as AR-GARCH.

On the other hand, the market and institutional changes have long been assumed to cause structural breaks in financial time series, which was confirmed, e.g. in data on stock prices (Andreou and Ghysels, 2002, and Beltratti and Morana, 2004) and exchange rates (Herwatz and Reimers, 2001). Moreover, ignoring these breaks can adversely affect the modelling, estimation and forecasting of volatility as suggested e.g. by Diebold and Inoue (2001), Mikosch and Starica (2004), Pesaran and Timmermann (2004) and Hillebrand (2005). Such findings led to the development of the change-point analysis in the context of CH models; see e.g. Chen and Gupta (1997), Kokoszka and Leipus (2000) and Andreou and Ghysels (2006).

An alternative approach lies in relaxing the assumption of time homogeneity and allowing some or all model parameters to vary over time (Chen and Tsay, 1993, Cai et al., 2000, and Fan and Zhang, 2008). Without structural assumptions about the transition of model parameters over time, time-varying coefficient models have to be estimated non-parametrically, e.g. under the identification condition that their parameters are smooth functions of time (Cai et al., 2000). In this paper, we follow a different strategy based on the assumption that a time series can be locally, i.e. over short periods of time, approximated by a parametric model. As suggested by Spokoiny (1998), such a local approximation can form a starting point in the search for the longest period of stability (homogeneity), i.e. for the longest time interval in which the series is described well by the parametric model. In the context of the local constant approximation, this strategy was employed for volatility modelling by Härdle et al. (2003), Mercurio and Spokoiny (2004) and Spokoiny (2009a). Our aim is to generalize this approach so that it can identify intervals of homogeneity for any parametric CH model regardless of its complexity.

In contrast to the local constant approximation of the volatility of a process (Mercurio and Spokoiny, 2004), the main benefit of the proposed generalization consists in the possibility to apply the methodology to a much wider class of models and to forecast over a longer time horizon. The reason is that approximating the mean or volatility process by a constant is in many cases too restrictive or even inappropriate and it is fulfilled only for short time intervals, which precludes its use for longer-term forecasting. On the contrary, parametric models like GARCH mimic the majority of stylized facts about financial time series and can reasonably fit the data over rather long periods of time in many practical situations. Allowing for time dependence of model parameters offers then much more flexibility in modelling real-life time series, which can be both with or without structural breaks since global parametric models are included as a special case.

Moreover, the proposed adaptive local parametric modelling unifies the change-point and varying-coefficient models. First, since finding the longest time-homogeneous interval for a parametric model at any point in time corresponds to detecting the most recent change-point in a time series, this approach resembles the change-point modelling as in Bai and Perron (1998) or Mikosch and Starica (1999, 2004), for instance, but it does not require prior information such as the number of changes. Additionally, the traditional structural-change tests require that the number of observations before each break point is large (and can grow to infinity) as these tests rely on asymptotic results. On the contrary, the proposed pointwise adaptive estimation does not rely on asymptotic results and does not thus place any requirements on the number of observations before, between or after any break point. Second, since the adaptively selected time-homogeneous interval used for estimation necessarily differs at each time point, the model coefficients can arbitrarily vary over time. In comparison to varying-coefficient models assuming

smooth development of parameters over time (Cai et al., 2000), our approach however allows for structural breaks in the form of sudden jumps in parameter values.

Although seemingly straightforward, extending Mercurio and Spokoiny's (2004) procedure to the local parametric modelling is a non-trivial problem, which requires new tools and techniques. We concentrate here on the change-point estimation of financial time series, which are often modelled by data-demanding models such as GARCH. While the benefits of a flexible change-point analysis for time series spanning several years are well known, its feasibility (which stands in the focus of this work) is much more difficult to achieve. The reason is thus that, at each time point, the procedure starts from a small interval, where a local parametric approximation holds, and then iteratively extends this interval and tests it for time-homogeneity until a structural break is found or data exhausted. Hence, a model has to be initially estimated on very short time intervals (e.g. 10 observations). Using standard testing methods, such a procedure might be feasible for simple parametric models, but it is hardly possible for more complex parametric models such as GARCH that generally require rather large samples for reasonably good estimates.

Therefore, we use an alternative and more robust approach to local change-point analysis that relies on a finite-sample theory of testing a growing sequence of historical time intervals on homogeneity against a change-point alternative. The proposed adaptive pointwise estimation procedure applies to a wide class of time-series models, including AR and CH models. Concentrating on the latter, we describe in details the adaptive procedure, derive its basic properties, and focusing on the feasibility of adaptive estimation for CH models, study the performance in comparison to the parametric (G)ARCH by means of simulations and real-data applications. The main conclusion is two-fold: on one hand, the adaptive pointwise estimation is feasible and beneficial also in the case of data-demanding models such as GARCH; on the other hand, the adaptive estimates based on various parametric models such as constant, ARCH or GARCH models are much closer to each other (while being better than the usual parametric estimates), which eliminates to some extent the need for using too complex models in adaptive estimation.

The rest of the paper is organized as follows. In Section 2, the parametric estimation of CH models and its finite-sample properties are introduced. In Section 3, we define the adaptive pointwise estimation procedure and discuss the choice of its parameters. Theoretical properties of the method are discussed in Section 4. In the specific case of the ARCH(1) and GARCH(1,1) models, a simulation study illustrates the performance of the new methodology with respect to the standard parametric and change-point models in Section 5. Applications to real stock-index series data are presented in Section 6. The proofs are provided in the Appendix.

2. PARAMETRIC CONDITIONAL HETEROSCEDASTICITY MODELS

Consider a time series Y_t in discrete time, $t \in N$. The CH assumption means that $Y_t = \sigma_t \varepsilon_t$, where $\{\varepsilon_t\}_{t \in N}$ is a white noise process and $\{\sigma_t\}_{t \in N}$ is a predictable volatility (conditional variance) process. Modelling of the volatility process σ_t typically relies on some parametric CH specification such as the ARCH (Engle, 1982) and GARCH (Bollerslev, 1986) models:

$$\sigma_t^2 = \omega + \sum_{i=1}^p \alpha_i Y_{t-i}^2 + \sum_{j=1}^q \beta_j \sigma_{t-j}^2, \quad (2.1)$$

where $p \in N, q \in N$ and $\theta = (\omega, \alpha_1, \dots, \alpha_p, \beta_1, \dots, \beta_q)^\top$ is the parameter vector. An attractive feature of this model is that, even with very few coefficients, one can model most stylized facts of financial time series like volatility clustering or excessive kurtosis, for instance. A number of (G)ARCH extensions were proposed to make the model even more flexible; e.g. EGARCH (Nelson, 1991), QGARCH (Sentana, 1995) and TGARCH (Glosten et al., 1993) that account for asymmetries in a volatility process.

All such CH models can be put into a common class of generalized linear volatility models:

$$Y_t = \sigma_t \varepsilon_t = \sqrt{g(X_t)} \varepsilon_t, \tag{2.2}$$

$$X_t = \omega + \sum_{i=1}^p \alpha_i h(Y_{t-i}) + \sum_{j=1}^q \beta_j X_{t-j}, \tag{2.3}$$

where g and h are known functions and X_t is a (partially) unobserved process (structural variable) that models the volatility coefficient σ_t^2 via transformation $g : \sigma_t^2 = g(X_t)$. For example, the GARCH model (2.1) is described by $g(u) = u$ and $h(r) = r^2$.

Models (2.2)–(2.3) are time homogeneous in the sense that the process Y_t follows the same structural equation at each time point. In other words, the parameter θ and hence the structural dependence in Y_t is constant over time. Even though models like (2.2)–(2.3) can often fit data well over a longer period of time, the assumption of homogeneity is too restrictive in practical applications: to guarantee a sufficient amount of data for sufficiently precise estimation, these models are often applied over time spans of many years. On the contrary, the strategy pursued here requires only local time homogeneity, which means that at each time point t there is a (possibly rather short) interval $[t - m, t]$, where the process Y_t is well described by models (2.2)–(2.3). This strategy aims then both at finding an interval of homogeneity (preferably as long as possible) and at the estimation of the corresponding parameter values θ , which then enable predicting Y_t and X_t .

Next, we discuss the parameter estimation for models (2.2)–(2.3) using observations Y_t from some time interval $I = [t_0, t_1]$. The conditional distribution of each observation Y_t given the past \mathcal{F}_{t-1} is determined by the structural variable X_t , whose dynamics are described by the parameter vector $\theta : X_t = X_t(\theta)$ for $t \in I$ due to (2.3). We denote the underlying value of θ by θ_0 .

For estimating θ_0 , we apply the quasi-maximum likelihood (quasi-MLE) approach using the estimating equations generated under the assumption of Gaussian errors ε_t . This guarantees efficiency under the normality of innovations and consistency under rather general moment conditions (Hansen and Lee, 1994, and Francq and Zakoian, 2007). The log-likelihood for models (2.2)–(2.3) on an interval I can be represented in the form

$$L_I(\theta) = \sum_{t \in I} \ell\{Y_t, g[X_t(\theta)]\}$$

with log-likelihood function $\ell(y, v) = -0.5\{\log(v) + y^2/v\}$. We define the quasi-MLE estimate $\tilde{\theta}_I$ of the parameter θ by maximizing the log-likelihood $L_I(\theta)$,

$$\tilde{\theta}_I = \underset{\theta \in \Theta}{\operatorname{argmax}} L_I(\theta) = \underset{\theta \in \Theta}{\operatorname{argmax}} \sum_{t \in I} \ell\{Y_t, g[X_t(\theta)]\}, \tag{2.4}$$

and denote by $L_I(\tilde{\theta}_I)$ the corresponding maximum.

To characterize the quality of estimating the parameter vector $\theta_0 = (\omega, \alpha_1, \dots, \alpha_p, \beta_1, \dots, \beta_q)^\top$ by $\tilde{\theta}_I$, we now present an exact (non-asymptotic) exponential risk bound. This bound concerns the value of maximum $L_I(\tilde{\theta}_I) = \max_{\theta \in \Theta} L_I(\theta)$ rather than the point of maximum $\tilde{\theta}_I$. More precisely, we consider the difference $L_I(\tilde{\theta}_I, \theta_0) = L_I(\tilde{\theta}_I) - L_I(\theta_0)$. By definition, this value is non-negative and represents the deviation of the maximum of the log-likelihood process from its value at the ‘true’ point θ_0 . Later, we comment on how the accuracy of estimation of the parameter θ_0 by $\tilde{\theta}_I$ relates to the value $L_I(\tilde{\theta}_I, \theta_0)$. We will also see that the bound for $L_I(\tilde{\theta}_I, \theta_0)$ yields the confidence set for the parameter θ_0 , which will be used for the proposed change-point test. Now, the non-asymptotic risk bound is specified in the following theorem, which formulates corollaries 4.2 and 4.3 of Spokoiny (2009b) for the case of the quasi-MLE estimation of a CH model (2.2)–(2.3) at $\theta = \theta_0$. The result can be viewed as an extension of the Wilks phenomenon that the distribution of $L_I(\tilde{\theta}_I, \theta_0)$ for a linear Gaussian model is $\chi_p^2/2$, where p is the number of estimated parameters in the model.

THEOREM 2.1. *Assume that the process Y_t follows models (2.2)–(2.3) with the parameter $\theta_0 \in \Theta$, where the set Θ is compact. The function $g(\cdot)$ is assumed to be continuously differentiable with the uniformly bounded first derivative and $g(x) \geq \delta > 0$ for all x . Further, let the process $X_t(\theta)$ be sub-ergodic in the sense that for any smooth function $f(\cdot)$ there exists f^* such that for any time interval I*

$$E_{\theta_0} \left| \sum_I \{f(X_t(\theta)) - E_{\theta_0} f(X_t(\theta))\} \right|^2 \leq f^* |I|, \quad \theta \in \Theta.$$

Finally, let $E \exp\{\varkappa(\varepsilon_t^2 - 1) | \mathcal{F}_{t-1}\} \leq c(\varkappa)$ for some $\varkappa > 0, c(\varkappa) > 0$, and all $t \in N$. Then there are $\lambda > 0$ and $\varepsilon(\lambda, \theta_0) > 0$ such that for any interval I and $\mathfrak{z} > 0$

$$P_{\theta_0}(L_I(\tilde{\theta}_I, \theta_0) > \mathfrak{z}) \leq \exp\{\varepsilon(\lambda, \theta_0) - \lambda \mathfrak{z}\}. \tag{2.5}$$

Moreover, for any $r > 0$, there is a constant $\mathfrak{R}_r(\theta_0)$ such that

$$E_{\theta_0} |L_I(\tilde{\theta}_I, \theta_0)|^r \leq \mathfrak{R}_r(\theta_0). \tag{2.6}$$

REMARK 2.1. The condition $g(x) \geq \delta > 0$ guarantees that the variance process cannot reach zero. In the case of GARCH, it is sufficient to assume $\omega > 0$, for instance.

One attractive feature of Theorem 2.1, formulated in the following corollary, is that it enables constructing the non-asymptotic confidence sets and testing the parametric hypothesis on the basis of the fitted log-likelihood $L_I(\tilde{\theta}_I, \theta)$. This feature is especially important for our procedure presented in Section 3.

COROLLARY 2.1. *Under the assumptions of Theorem 2.1, let the value \mathfrak{z}_α fulfil $\varepsilon(\lambda, \theta_0) - \lambda \mathfrak{z}_\alpha < \log \alpha$ for some $\alpha < 1$. Then the random set $\mathcal{E}_I(\mathfrak{z}_\alpha) = \{\theta : L_I(\tilde{\theta}_I, \theta) \leq \mathfrak{z}_\alpha\}$ is an α -confidence set for θ_0 in the sense that $P_{\theta_0}(\theta_0 \notin \mathcal{E}_I(\mathfrak{z}_\alpha)) \leq \alpha$.*

Theorem 2.1 also gives a non-asymptotic and fixed upper bound for the risk of estimation $L_I(\tilde{\theta}_I, \theta_0)$ that applies to an arbitrary sample size $|I|$. To understand the relation of this result to the classical rate result, we can apply the standard arguments based on the quadratic expansion

of the log-likelihood $L(\tilde{\theta}, \theta)$. Let $\nabla^2 L(\theta)$ denote the Hessian matrix of the second derivatives of $L(\theta)$ with respect to the parameter θ . Then

$$L_I(\tilde{\theta}_I, \theta_0) = 0.5(\tilde{\theta}_I - \theta_0)^\top \nabla^2 L_I(\theta'_I)(\tilde{\theta}_I - \theta_0), \quad (2.7)$$

where θ'_I is a convex combination of θ_0 and $\tilde{\theta}_I$. Under usual regularity assumptions and for sufficiently large $|I|$, the normalized matrix $|I|^{-1} \nabla^2 L_I(\theta)$ is close to some matrix $V(\theta)$, which depends only on the stationary distribution of Y_t and is continuous in θ . Then (2.5) approximately means that $\|\sqrt{V(\theta_0)}(\tilde{\theta}_I - \theta_0)\|^2 \leq \mathfrak{z}/|I|$ with probability close to 1 for large \mathfrak{z} . Hence, the large deviation result of Theorem 2.1 yields the root- $|I|$ consistency of the MLE estimate $\tilde{\theta}_I$. See Spokoiny (2009b) for further details.

3. POINTWISE ADAPTIVE NON-PARAMETRIC ESTIMATION

An obvious feature of models (2.2)–(2.3) is that the parametric structure of the process is assumed constant over the whole sample and cannot thus incorporate changes and structural breaks at unknown times in the models. A natural generalization leads to models whose coefficients may change over time (Fan and Zhang, 2008). One can then assume that the structural process X_t satisfies the relation (2.3) at any time, but the vector of coefficients θ may vary with the time t , $\theta = \theta(t)$. The estimation of the coefficients as general functions of time is possible only under some additional assumptions on these functions. Typical assumptions are (i) varying coefficients are smooth functions of time (Cai et al., 2000) and (ii) varying coefficients are piecewise constant functions (Bai and Perron, 1998, and Mikosch and Starica, 1999, 2004).

Our local parametric approach differs from the commonly used identification assumptions (i) and (ii). We assume that the observed data Y_t are described by a (partially) unobserved process X_t due to (2.2), and at each point T , there exists a historical interval $I(T) = [t_0, T]$ in which the process X_t ‘nearly’ follows the parametric specification (2.3) (see Section 4 for details on what ‘nearly’ means). This local structural assumption enables us to apply well-developed parametric estimation for data $\{Y_t\}_{t \in I(T)}$ to estimate the underlying parameter $\theta = \theta(T)$ by $\hat{\theta} = \hat{\theta}(T)$. (The estimate $\hat{\theta} = \hat{\theta}(T)$ can then be used for estimating the value \hat{X}_T of the process X_t at T from equation (2.3) and for further modelling such as forecasting Y_{T+1} .) Moreover, this assumption includes the above-mentioned ‘smooth transition’ and ‘switching regime’ assumptions (i) and (ii) as special cases: parameters $\hat{\theta}(T)$ vary over time as the interval $I(T)$ changes with T and, at the same time, discontinuities and jumps in $\hat{\theta}(T)$ as a function of time are possible.

To estimate $\hat{\theta}(T)$, we have to find the historical interval of homogeneity $I(T)$, i.e. the longest interval I with the right-end point T , where data do not contradict a specified parametric model with fixed parameter values. Starting at each time T with a very short interval $I = [t_0, T]$, we search by successive extending and testing of interval I on homogeneity against a change-point alternative: if the hypothesis of homogeneity is not rejected for a given I , a larger interval is taken and tested again. Contrary to Bai and Perron (1998) and Mikosch and Starica (1999), who detect all change points in a given time series, our approach is local: it focuses on the local change-point analysis near point T of estimation and tries to find only one change closest to the reference point.

In the rest of this section, we first discuss the test statistics employed to test the time-homogeneity of an interval I against a change-point alternative in Section 3.1. Later, we rigorously describe the pointwise adaptive estimation procedure in Section 3.2. Its

implementation and the choice of parameters entering the adaptive procedure are described in Sections 3.2–3.4. Theoretical properties of the method are studied in Section 4.

3.1. Test of homogeneity against a change-point alternative

The pointwise adaptive estimation procedure crucially relies on the test of local time-homogeneity of an interval $I = [t_0, T]$. The null hypothesis for I means that the observations $\{Y_t\}_{t \in I}$ follow the parametric models (2.2)–(2.3) with a fixed parameter θ_0 , leading to the quasi-MLE estimate $\tilde{\theta}_I$ from (2.4) and the corresponding fitted log-likelihood $L_I(\tilde{\theta}_I)$.

The change-point alternative for a given change-point location $\tau \in I$ can be described as follows: process Y_t follows the parametric models (2.2)–(2.3) with a parameter θ_J for $t \in J = [t_0, \tau]$ and with a different parameter θ_{J^c} for $t \in J^c = [\tau + 1, T]$; $\theta_J \neq \theta_{J^c}$. The fitted log-likelihood under this alternative reads as $L_J(\tilde{\theta}_J) + L_{J^c}(\tilde{\theta}_{J^c})$. The test of homogeneity can be performed using the likelihood ratio (LR) test statistic $T_{I,\tau}$:

$$T_{I,\tau} = \max_{\theta_J, \theta_{J^c} \in \Theta} \{L_J(\theta_J) + L_{J^c}(\theta_{J^c})\} - \max_{\theta \in \Theta} L_I(\theta) = \{L_J(\tilde{\theta}_J) + L_{J^c}(\tilde{\theta}_{J^c}) - L_I(\tilde{\theta}_I)\}.$$

Since the change-point location τ is generally not known, we consider the supremum of the LR statistics $T_{I,\tau}$ over some subset $\tau \in \mathcal{T}(I)$; cf. Andrews (1993):

$$T_{I,\mathcal{T}(I)} = \sup_{\tau \in \mathcal{T}(I)} T_{I,\tau}. \quad (3.1)$$

A typical example of a set $\mathcal{T}(I)$ is $\mathcal{T}(I) = \{\tau : t_0 + m' \leq \tau \leq T - m''\}$ for some fixed $m', m'' > 0$.

3.2. Adaptive search for the longest interval of homogeneity

This section presents the proposed adaptive pointwise estimation procedure. At each point T , we aim at estimating the unknown parameters $\theta(T)$ from historical data $Y_t, t \leq T$; this procedure repeats for every current time point T as new data arrive. At the first step, the procedure selects on the base of historical data an interval $\hat{I}(T)$ of homogeneity in which the data do not contradict the parametric models (2.2)–(2.3). Afterwards, the quasi-MLE estimation is applied using the selected historical interval $\hat{I}(T)$ to obtain estimate $\hat{\theta}(T) = \tilde{\theta}_{\hat{I}(T)}$. From now on, we consider an arbitrary, but fixed time point T .

Suppose that a growing set $I_0 \subset I_1 \subset \dots \subset I_K$ of historical interval-candidates $I_k = [T - m_k + 1, T]$ with the right-end point T is fixed. The smallest interval I_0 is accepted automatically as homogeneous. Then the procedure successively checks every larger interval I_k on homogeneity using the test statistic $T_{I_k, \mathcal{T}(I_k)}$ from (3.1). The selected interval \hat{I} corresponds to the largest accepted interval $I_{\hat{k}}$ with index \hat{k} such that

$$T_{I_k, \mathcal{T}(I_k)} \leq \mathfrak{z}_k, \quad k \leq \hat{k}, \quad (3.2)$$

and $T_{I_{\hat{k}+1}, \mathcal{T}(I_{\hat{k}+1})} > \mathfrak{z}_{\hat{k}+1}$, where the critical values \mathfrak{z}_k are discussed later in this section and specified in Section 3.3. This procedure then leads to the adaptive estimate $\hat{\theta} = \tilde{\theta}_{\hat{I}}$ corresponding to the selected interval $\hat{I} = I_{\hat{k}}$.

The complete description of the procedure includes two steps. (A) Fixing the set-up and the parameters of the procedure. (B) Data-driven search for the longest interval of homogeneity.

(A) Set-up and parameters:

- 1 Select specific parametric models (2.2)–(2.3) [e.g. constant volatility, ARCH(1), GARCH(1,1)].
- 2 Select the set $\mathcal{I} = (I_0, \dots, I_K)$ of interval-candidates, and for each $I_k \in \mathcal{I}$, the set $\mathcal{T}(I_k)$ of possible change points $\tau \in I_k$ used in the LR test (3.1).
- 3 Select the critical values β_1, \dots, β_K in (3.2) as described in Section 3.3.

(B) Adaptive search and estimation: Set $k = 1$, $\hat{I} = I_0$ and $\hat{\theta} = \tilde{\theta}_{I_0}$.

- 1 Test the hypothesis $H_{0,k}$ of no change point within the interval I_k using test statistics (3.1) and the critical values β_k obtained in (A3). If a change point is detected ($H_{0,k}$ is rejected), go to (B3). Otherwise proceed with (B2).
- 2 Set $\hat{\theta} = \tilde{\theta}_{I_k}$ and $\hat{I}_k = \tilde{I}_k$. Further, set $k := k + 1$. If $k \leq K$, repeat (B1); otherwise go to (B3).
- 3 Define $\hat{I} = I_{k-1}$ = ‘the last accepted interval’ and $\hat{\theta} = \tilde{\theta}_{\hat{I}}$. Additionally, set $\hat{\theta}_{I_k} = \dots = \hat{\theta}_{I_K} = \hat{\theta}$ if $k \leq K$.

In step (A), one has to select three main ingredients of the procedure. First, the parametric model used locally to approximate the process Y_t has to be specified in (A1), e.g. the constant volatility or GARCH(1,1) in our context. Next, in step (A2), the set of intervals $\mathcal{I} = \{I_k\}_{k=0}^K$ is fixed, each interval with the right-end point T , length $m_k = |I_k|$, and the set $\mathcal{T}(I_k)$ of tested change points. Our default proposal is to use a geometric grid $m_k = [m_0 a^k]$, $a > 1$, and to set $I_k = [T - m_k + 1, T]$ and $\mathcal{T}(I_k) = [T - m_{k-1} + 1, T - m_{k-2}]$. Although our experiments show that the procedure is rather insensitive to the choice of m_0 and a (e.g. we use $m_0 = 10$ and $a = 1.25$ in simulations), the length m_0 of interval I_0 should take into account the parametric model selected in (A1). The reason is that I_0 is always assumed to be time-homogeneous and m_0 thus has to reflect flexibility of the parametric model; e.g. while $m_0 = 20$ might be reasonable for the GARCH(1,1) model, $m_0 = 5$ could be a reasonable choice for the locally constant approximation of a volatility process. Finally, in step (A3), one has to select the K critical values β_k in (3.2) for the LR test statistics $T_{I_k, \mathcal{T}(I_k)}$ from (3.1). The critical values β_k will generally depend on the parametric model describing the null hypothesis of time-homogeneity, the set \mathcal{I} of intervals I_k and corresponding sets of considered change points $\mathcal{T}(I_k)$, $k \leq K$, and additionally, on two constants r and ρ that are counterparts of the usual significance level. All these determinants of the critical values can be selected in step (A) and the critical values are thus obtained before the actual estimation takes place in step (B). Due to its importance, the method of constructing critical values $\{\beta_k\}_{k=1}^K$ is discussed separately in Section 3.3.

The main step (B) performs the search for the longest time-homogeneous interval. Initially, I_0 is assumed to be homogeneous. If I_{k-1} is negatively tested on the presence of a change point, one continues with I_k by employing test (3.1) in step (B1), which checks for a potential change point in I_k . If no change point is found, then I_k is accepted as time-homogeneous in step (B2); otherwise the procedure terminates in step (B3). We sequentially repeat these tests until we find a change point or exhaust all intervals. The latest (longest) interval accepted as time-homogeneous is used for estimation in step (B3). Note that the estimate $\hat{\theta}_{I_k}$ defined in (B2) and (B3) corresponds to the latest accepted interval \hat{I}_k after the first k steps, or equivalently, the interval selected out of I_1, \dots, I_k .

Moreover, the whole search and estimation step (B) can be repeated at different time points T without reiterating the initial step (A) as the critical values β_k depend only on the approximating parametric model and interval lengths $m_k = |I_k|$, not on the time point T (see Section 3.3).

3.3. Choice of critical values \mathfrak{z}_k

The presented method of choosing the interval of homogeneity \hat{I} can be viewed as multiple testing procedure. The critical values for this procedure are selected using the general approach of testing theory: to provide a prescribed performance of the procedure under the null hypothesis, i.e. in the pure parametric situation. This means that the procedure is trained on the data generated from the pure parametric time-homogeneous model from step (A1). The correct choice in this situation is the largest considered interval I_K and a choice $I_{\hat{k}}$ with $\hat{k} < K$ can be interpreted as a ‘false alarm’. We select the minimal critical values ensuring a small probability of such a false alarm. Our condition slightly differs though from the classical level condition because we focus on parameter estimation rather than on hypothesis testing.

In the pure parametric case, the ‘ideal’ estimate corresponds to the largest considered interval I_K . Due to Theorem 2.1, the quality of estimation of the parameter θ_0 by $\tilde{\theta}_{I_K}$ can be measured by the log-likelihood ‘loss’ $L_{I_K}(\tilde{\theta}_{I_K}, \theta_0)$, which is stochastically bounded with exponential and polynomial moments: $E_{\theta_0} |L_{I_K}(\tilde{\theta}_{I_K}, \theta_0)|^r \leq \mathfrak{R}_r(\theta_0)$. If the adaptive procedure stops earlier at some intermediate step $k < K$, we select instead of $\tilde{\theta}_{I_K}$ another estimate $\hat{\theta} = \tilde{\theta}_{I_k}$ with a larger variability. The loss associated with such a false alarm can be measured by the value $L_{I_k}(\tilde{\theta}_{I_k}, \hat{\theta}) = L_{I_k}(\tilde{\theta}_{I_k}) - L_{I_k}(\hat{\theta})$. The corresponding condition bounding the loss due to the adaptive estimation reads as

$$E_{\theta_0} |L_{I_k}(\tilde{\theta}_{I_k}, \hat{\theta})|^r \leq \rho \mathfrak{R}_r(\theta_0). \tag{3.3}$$

This is in fact an implicit condition on the critical values $\{\mathfrak{z}_k\}_{k=1}^K$, which ensures that the loss associated with the false alarm is at most the ρ -fraction of the log-likelihood loss of the ‘ideal’ or ‘oracle’ estimate $\tilde{\theta}_{I_K}$ for the parametric situation. The constant r corresponds to the power of the loss in (3.3), while ρ is similar in meaning to the test level. In the limit case when r tends to zero, this condition (3.3) becomes the usual level condition: $P_{\theta_0}(I_K \text{ is rejected}) = P_{\theta_0}(\tilde{\theta}_{I_k} \neq \hat{\theta}) \leq \rho$. The choice of the metaparameters r and ρ is discussed in Section 3.4.

A condition similar to (3.3) is imposed at each step of the adaptive procedure. The estimate $\hat{\theta}_{I_k}$ coming after the k steps of the procedure should satisfy

$$E_{\theta_0} |L_{I_k}(\tilde{\theta}_{I_k}, \hat{\theta}_{I_k})|^r \leq \rho_k \mathfrak{R}_r(\theta_0), \quad k = 1, \dots, K, \tag{3.4}$$

where $\rho_k = \rho k/K \leq \rho$. The following theorem presents some sufficient conditions on the critical values $\{\mathfrak{z}_k\}_{k=1}^K$ ensuring (3.4); recall that $m_k = |I_k|$ denotes the length of I_k .

THEOREM 3.1. *Suppose that $r > 0, \rho > 0$. Under the assumptions of Theorem 2.1, there are constants a_0, a_1, a_2 such that the condition (3.4) is fulfilled with the choice*

$$\mathfrak{z}_k = a_0 r \log(\rho^{-1}) + a_1 r \log(m_K/m_{k-1}) + a_2 \log(m_k), \quad k = 1, \dots, K.$$

Since K and $\{m_k\}_{k=1}^K$ are fixed, the \mathfrak{z}_k ’s in Theorem 3.1 have a form $\mathfrak{z}_k = C + D \log(m_k)$ for $k = 1, \dots, K$ with some constant C and D . However, a practically relevant choice of these constants has to be done by Monte Carlo simulations. Note first that every particular choice of the coefficients C and D determines the whole set of the critical values $\{\mathfrak{z}_k\}_{k=1}^K$ and thus the local change-point procedure. For the critical values given by fixed (C, D) , one can run the procedure and observe its performance on the simulated data using the data-generating process (2.2)–(2.3); in particular, one can check whether the condition (3.4) is fulfilled. For any (sufficiently large) fixed value of C , one can thus find the minimal value $D(C) < 0$ of D that ensures (3.4).

Every corresponding set of critical values in the form $\beta_k = C + D(C) \log(m_k)$ is admissible. The condition $D(C) < 0$ ensures that the critical values decreases with k . This reflects the fact that a false alarm at an early stage of the algorithm is more crucial because it leads to the choice of a highly variable estimate. The critical values β_k for small k should thus be rather conservative to provide the stability of the algorithm in the parametric situation. To determine C , the value β_1 can be fixed by considering the false alarm at the first step of the procedure, which leads to estimation using the smallest interval I_0 instead of the ‘ideal’ largest interval I_K . The related condition (used in Section 5.1) reads as

$$E_{\theta_0} |L_{I_K}(\tilde{\theta}_{I_K}, \tilde{\theta}_{I_0})|^r \mathbf{1}(T_{I_1, \mathcal{I}(I_1)} > \beta_1) \leq \rho \mathfrak{R}_r(\theta_0)/K. \tag{3.5}$$

Alternatively, one could select a pair (C, D) that minimizes the resulting prediction error; see Section 3.4.

3.4. Selecting parameters r and ρ

The choice of critical values using inequality (3.4) additionally depends on two ‘metaparameters’ r and ρ . A simple strategy is to use conservative values for these parameters and the corresponding set of critical values (e.g. our default is $r = 1$ and $\rho = 1$). On the other hand, the two parameters are global in the sense that they are independent of T . Hence, one can also determine them in a data-driven way by minimizing some global forecasting error (Cheng et al., 2003). Different values of r and ρ may lead to different sets of critical values and hence to different estimates $\hat{\theta}^{(r, \rho)}(T)$ and to different forecasts $\hat{Y}_{T+h|T}^{(r, \rho)}$ of the future values Y_{T+h} , where h is the forecasting horizon. Now, a data-driven choice of r and ρ can be done by minimizing the following objective function:

$$(\hat{r}, \hat{\rho}) = \arg \min_{r > 0, \rho > 0} PE_{\Lambda, \mathcal{H}}(r, \rho) = \arg \min_{r, \rho} \sum_T \sum_{h \in \mathcal{H}} \Lambda(Y_{T+h}, \hat{Y}_{T+h|T}^{(r, \rho)}), \tag{3.6}$$

where Λ is a loss function and \mathcal{H} is the forecasting horizon set. For example, one can take $\Lambda_r(v, v') = |v - v'|^r$ for $r \in [1/2, 2]$. For daily data, the forecasting horizon could be one day, $\mathcal{H} = \{1\}$, or two weeks, $\mathcal{H} = \{1, \dots, 10\}$.

4. THEORETIC PROPERTIES

In this section, we collect basic results describing the quality of the proposed adaptive procedure. First, the definition of the procedure ensures the performance prescribed by (3.4) in the parametric situation. We however claimed that the adaptive pointwise estimation applies even if the process Y_t is only locally approximated by a parametric model. Therefore, we now define a locally ‘nearly parametric’ process, for which we derive an analogy of Theorem 2.1 (Section 4.1). Later, we prove certain ‘oracle’ properties of the proposed method (Section 4.2).

4.1. Small modelling bias condition

This section discusses the concept of a ‘nearly parametric’ case. To define it rigorously, we have to quantify the quality of approximating the true latent process X_t , which drives the observed data Y_t due to (2.2), by the parametric process $X_t(\theta)$ described by (2.3) for some $\theta \in \Theta$. Below

we assume that the innovations ε_t in the model (2.2) are independent and identically distributed and denote the distribution of $\sqrt{v}\varepsilon_t$ by P_v so that the conditional distribution of Y_t given \mathcal{F}_{t-1} is $P_{g(X_t)}$. To measure the distance of a data-generating process from a parametric model, we introduce for every interval $I_k \in \mathcal{I}$ and every parameter $\theta \in \Theta$ the random quantity

$$\Delta_{I_k}(\theta) = \sum_{t \in I_k} \mathcal{K} \{g(X_t), g[X_t(\theta)]\},$$

where $\mathcal{K}(v, v')$ denotes the Kullback–Leibler distance between P_v and $P_{v'}$. For CH models with Gaussian innovations ε_t , $\mathcal{K}(v, v') = -0.5\{\log(v/v') + 1 - v/v'\}$. In the parametric case with $X_t = X_t(\theta_0)$, we clearly have $\Delta_{I_k}(\theta_0) = 0$. To characterize the ‘nearly parametric case’, we introduce a {small modelling bias} (SMB) condition, which simply means that, for some $\theta \in \Theta$, $\Delta_{I_k}(\theta)$ is bounded by a small constant with a high probability. Informally, this means that the ‘true’ model can be well approximated on the interval I_k by the parametric one with the parameter θ . The best parametric fit (2.3) to the underlying model (2.2) on I_k can be defined by minimizing the value $E\Delta_{I_k}(\theta)$ over $\theta \in \Theta$ and $\tilde{\theta}_{I_k}$ can be viewed as its estimate.

The following theorem claims that the results on the accuracy of estimation given in Theorem 2.1 can be extended from the parametric case to the general non-parametric situation under the SMB condition. Let $\varrho(\hat{\theta}, \theta)$ be any loss function for an estimate $\hat{\theta}$.

THEOREM 4.1. *Let for some $\theta \in \Theta$ and some $\Delta \geq 0$*

$$E\Delta_{I_k}(\theta) \leq \Delta. \tag{4.1}$$

Then it holds for an estimate $\hat{\theta}$ constructed from the observations $\{Y_t\}_{t \in I_k}$ that

$$E \log(1 + \varrho(\hat{\theta}, \theta)/E\theta\varrho(\hat{\theta}, \theta)) \leq 1 + \Delta.$$

This general result applied to the quasi-MLE estimation with the loss function $L_I(\tilde{\theta}_I, \theta)$ yields the following corollary.

COROLLARY 4.1. *Let the SMB condition (4.1) hold for some interval I_k and $\theta \in \Theta$. Then*

$$E \log\left(1 + |L_{I_k}(\tilde{\theta}_{I_k}, \theta)|^r / \mathfrak{R}_r(\theta)\right) \leq 1 + \Delta,$$

where $\mathfrak{R}_r(\theta)$ is the parametric risk bound from (2.6).

This result shows that the estimation loss $|L_I(\tilde{\theta}_I, \theta)|^r$ normalized by the parametric risk $\mathfrak{R}_r(\theta)$ is stochastically bounded by a constant proportional to e^Δ . If Δ is not large, this result extends the parametric risk bound (Theorem 2.1) to the non-parametric situation under the SMB condition. Another implication of Corollary 4.1 is that the confidence set built for the parametric model (Corollary 2.1) continues to hold, with a slightly smaller coverage probability, under SMB.

4.2. The ‘oracle’ choice and the ‘oracle’ result

Corollary 4.1 suggests that the ‘optimal’ or ‘oracle’ choice of the interval I_k from the set I_1, \dots, I_K can be defined as the largest interval for which the SMB condition (4.1) still holds (for a given small $\Delta > 0$). For such an interval, one can neglect deviations of the underlying

process from a parametric model with a fixed parameter θ . Therefore, we say that the choice k^* is the ‘oracle’ choice if there exists $\theta \in \Theta$ such that

$$E \Delta_{I_{k^*}}(\theta) \leq \Delta \tag{4.2}$$

for a fixed $\Delta > 0$ and that (4.2) does not hold for $k > k^*$. Unfortunately, the underlying process X_t and, hence, the value Δ_{I_k} is unknown and the oracle choice cannot be implemented. The proposed adaptive procedure tries to mimic this oracle on the basis of available data using the sequential test of homogeneity. The final oracle result claims that the adaptive estimate provides the same (in order) accuracy as the oracle one.

By construction, the pointwise adaptive procedure described in Section 3 provides the prescribed performance if the underlying process follows the parametric model (2.2). Now, condition (3.4) combined with Theorem 4.1 implies similar performance in the first k^* steps of the adaptive estimation procedure.

THEOREM 4.2. *Let $\theta \in \Theta$ and $\Delta > 0$ be such that $E \Delta_{I_{k^*}}(\theta) \leq \Delta$ for some $k^* \leq K$. Also let $\max_{k \leq k^*} E_{\theta} |L_{I_k}(\tilde{\theta}_{I_k}, \theta)|^r \leq \mathfrak{R}_r(\theta)$. Then*

$$E \log \left(1 + \frac{|L_{I_{k^*}}(\tilde{\theta}_{I_{k^*}}, \theta)|^r}{\mathfrak{R}_r(\theta)} \right) \leq 1 + \Delta \quad \text{and} \quad E \log \left(1 + \frac{|L_{I_{k^*}}(\tilde{\theta}_{I_{k^*}}, \hat{\theta}_{I_{k^*}})|^r}{\mathfrak{R}_r(\theta)} \right) \leq \rho + \Delta.$$

Similarly to the parametric case, under the SMB condition $E \Delta_{I_{k^*}}(\theta) \leq \Delta$, any choice $\hat{k} < k^*$ can be viewed as a false alarm. Theorem 4.2 documents that the loss induced by such a false alarm at the first k^* steps and measured by $L_{I_{k^*}}(\tilde{\theta}_{I_{k^*}}, \hat{\theta}_{I_{k^*}})$ is of the same magnitude as the loss $L_{I_{k^*}}(\tilde{\theta}_{I_{k^*}}, \theta)$ of estimating the parameter θ from the SMB (4.2) by $\tilde{\theta}_{I_{k^*}}$. Thus, under (4.2) the adaptive estimation during steps $k \leq k^*$ does not induce larger errors into estimation than the quasi-MLE estimation itself.

For further steps of the algorithm with $k > k^*$, where (4.2) does not hold, the value $\Delta' = E \Delta_{I_k}(\theta)$ can be large and the bound for the risk becomes meaningless due to the factor $e^{\Delta'}$. To establish the result about the quality of the final estimate, we thus have to show that the quality of estimation cannot be destroyed at the steps $k > k^*$. The next ‘oracle’ result states the final quality of our adaptive estimate $\hat{\theta}$.

THEOREM 4.3. *Let $E \Delta_{I_{k^*}}(\theta) \leq \Delta$ for some $k^* \leq K$. Then $L_{I_{k^*}}(\tilde{\theta}_{I_{k^*}}, \hat{\theta}) \mathbf{1}(\hat{k} \geq k^*) \leq \mathfrak{z}_{k^*}$ yielding*

$$E \log \left(1 + \frac{|L_{I_{k^*}}(\tilde{\theta}_{I_{k^*}}, \hat{\theta})|^r}{\mathfrak{R}_r(\theta)} \right) \leq \rho + \Delta + \log \left(1 + \frac{\mathfrak{z}_{k^*}^r}{\mathfrak{R}_r(\theta)} \right).$$

Due to this result, the value $L_{I_{k^*}}(\tilde{\theta}_{I_{k^*}}, \hat{\theta})$ is stochastically bounded. This can be interpreted as the oracle property of $\hat{\theta}$ because it means that the adaptive estimate $\hat{\theta}$ belongs with a high probability to the confidence set of the oracle estimate $\tilde{\theta}_{I_{k^*}}$.

5. SIMULATION STUDY

In the last two sections, we present simulation study (Section 5) and real data applications (Section 6) documenting the performance of the proposed adaptive estimation procedure. To verify the practical applicability of the method in a complex setting, we concentrate on the volatility estimation using parametric and adaptive pointwise estimation of constant volatility, ARCH(1) and GARCH(1,1) models (for the sake of brevity, referred to as the local constant,

local ARCH and local GARCH). The reason is that the estimation of GARCH models requires generally hundreds of observations for a reasonable quality of estimation, which puts the adaptive procedure working with samples as small as 10 or 20 observations to a hard test. Additionally, the critical values obtained as described in Section 3.3 depend on the underlying parameter values in the case of (G)ARCH.

Here we first study the finite-sample critical values for the test of homogeneity by means of Monte Carlo simulations and discuss practical implementation details (Section 5.1). Later, we demonstrate the performance of the proposed adaptive pointwise estimation procedure in simulated samples (Section 5.2). Note that, throughout this section, we identify the GARCH(1,1) models by triplets (ω, α, β) : e.g. (1, 0.1, 0.3)-model. Constant volatility and ARCH(1) are then indicated by $\alpha = \beta = 0$ and $\beta = 0$, respectively. The GARCH estimation is done using the GARCH 3.0 package (Laurent and Peters, 2006) and Ox 3.30 (Doornik, 2002). Finally, since the focus is on modelling the volatility σ_t^2 in (2.2), the performance measurement and comparison of all models at time t is done by the absolute prediction error (PE) of the volatility process over a prediction horizon \mathcal{H} : $\text{APE}(t) = \sum_{h \in \mathcal{H}} |\sigma_{t+h}^2 - \hat{\sigma}_{t+h|t}^2| / |\mathcal{H}|$, where $\hat{\sigma}_{t+h|t}^2$ represents the volatility prediction by a particular model.

5.1. Finite-sample critical values for the test of homogeneity

A practical application of the pointwise adaptive procedure requires critical values for the test of local homogeneity of a time series. Since they are obtained under the null hypothesis that a chosen parametric model (locally) describes the data, see Section 3, we need to obtain the critical values for the constant volatility, ARCH(1) and GARCH(1,1) models. Furthermore, for given r and ρ , the average risk (3.4) between the adaptive and oracle estimates can be bounded for critical values that linearly depend on the logarithm of interval length $|I_k|$: $\mathfrak{z}(|I_k|) = \mathfrak{z}_k = C + D \log(|I_k|)$ (see Theorem 3.1). As described in Section 3.3, we choose here the smallest C satisfying (3.5) and the corresponding minimum admissible value $D = D(C) < 0$ that guarantees the conditions (3.4).

We simulated the critical values for ARCH(1) and GARCH(1,1) models with different values of underlying parameters; see Table 1 for the critical values corresponding to $r = 1$ and $\rho = 1$. Their simulation was performed sequentially on intervals with lengths ranging from $|I_0| = m_0 = 10$ to $|I_K| = 570$ observations using a geometric grid with multiplier $a = 1.25$; see Section 3.2. (The results are, however, not sensitive to the choice of a .)

Unfortunately, the critical values depend on the parameters of the underlying (G)ARCH model (in contrast to the constant-volatility model). They generally seem to increase with the values of the ARCH and GARCH parameters keeping the other one fixed; see Table 1. To deal with this dependence on the underlying model parameters, we propose to choose the largest (most conservative) critical values corresponding to any estimated parameter in the analysed data. For example, if the largest estimated parameters of GARCH(1,1) are $\hat{\alpha} = 0.3$ and $\hat{\beta} = 0.8$, one should use $\mathfrak{z}(10) = 26.4$ and $\mathfrak{z}(570) = 14.5$, which are the largest critical values for models with $\alpha = 0.3, \beta \leq 0.8$ and with $\alpha \leq 0.3, \beta = 0.8$. (The proposed procedure is, however, not overly sensitive to this choice, as we shall see later.)

Finally, let us have a look at the influence of the tuning constants r and ρ in (3.4) on the critical values for several selected models (Table 2). The influence is significant, but can be classified in the following way. Whereas increasing ρ generally leads to an overall decrease of critical values (cf. Theorem 3.1), but primarily for the longer intervals, increasing r leads to an increase of

Table 1. Critical values $\mathfrak{z}_k = \mathfrak{z}(|I_k|)$ of the supremum LR test.

$\mathfrak{z}(I_k)$	α	$ I_k $	β									
			0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9
0.0		10	15.5	15.5	16.4	16.8	17.9	17.3	17.0	17.0	16.9	16.0
		570	5.5	7.2	7.0	7.0	7.5	7.5	7.4	7.3	7.0	6.7
0.1		10	16.3	14.5	15.1	15.9	16.4	15.9	16.1	16.0	16.0	
		570	8.6	9.0	9.1	9.6	9.8	10.7	11.5	12.5	14.0	
0.2		10	16.7	15.2	15.7	16.2	16.9	18.9	20.1	25.1		
		570	9.4	10.6	11.2	11.4	11.4	12.5	13.3	14.2		
0.3		10	18.5	16.4	16.7	16.9	18.1	21.8	26.4			
		570	9.7	10.8	12.0	12.4	12.9	13.5	14.5			
0.4		10	22.1	16.5	18.3	19.3	22.8	30.9				
		570	9.9	12.0	13.0	13.4	13.9	14.7				
0.5		10	26.2	19.1	19.5	25.4	38.1					
		570	10.7	12.6	13.8	14.0	14.6					
0.6		10	33.0	22.8	25.9	32.4						
		570	12.7	12.7	13.9	15.3						
0.7		10	41.1	24.8	29.1							
		570	16.8	14.7	16.1							
0.8		10	66.2	26.4								
		570	31.5	15.8								
0.9		10	88.6									
		570	60.9									

Note: $\omega = 1, r = 1$ and $\rho = 1$.


critical values mainly for the shorter intervals; cf. (3.4). In simulations and real applications, we verified that a fixed choice such as $r = 1$ and $\rho = 1$ performs well. To optimize the performance of the adaptive methods, one can however determine constants r and ρ in a data-dependent way as described in Section 3.3. We use here this strategy for a small grid of $r \in \{0.5, 1.0\}$ and $\rho \in \{0.5, 1.0, 1.5\}$ and find globally optimal r and ρ . We will document, though, that the differences in the average absolute PE (3.6) for various values of r and ρ are relatively small.

5.2. Simulation study

We aim (i) to examine how well the proposed estimation method is able to adapt to long stable (time-homogeneous) periods and to less stable periods with more frequent volatility changes and (ii) to see which adaptively estimated model—local volatility, local ARCH or local GARCH—performs best in different regimes. To this end, we simulated 100 series from two change-point GARCH models with a low GARCH effect ($\omega, 0.2, 0.1$) and a high GARCH effect ($\omega, 0.2, 0.7$). Changes in constant ω are spread over a time span of 1000 days; see Figure 1. There is a long stable period at the beginning (500 days ≈ 2 years) and end (250 days ≈ 1 year) of time series with several volatility changes between them.

Table 2. Critical values $\mathfrak{z}(|I_k|)$ of the supremum LR test for various values r and ρ .

Model (ω, α, β)		(0.1, 0.0, 0.0)		(0.1, 0.2, 0.0)		(0.1, 0.1, 0.8)	
r	ρ	$\mathfrak{z}(10)$	$\mathfrak{z}(570)$	$\mathfrak{z}(10)$	$\mathfrak{z}(570)$	$\mathfrak{z}(10)$	$\mathfrak{z}(570)$
1.0	0.5	16.3	7.3	17.4	11.2	18.7	17.1
1.0	1.0	15.4	5.5	16.7	9.4	16.0	14.0
1.0	1.5	14.9	4.5	15.9	8.3	15.2	13.4
0.5	0.5	10.7	7.1	11.7	10.1	11.7	10.1
0.5	1.0	8.9	5.5	10.3	8.5	10.3	8.5
0.5	1.5	7.7	4.6	9.3	7.5	9.3	7.5

**Figure 1.** GARCH(1,1) parameters of low (left panel) and high (right panel) GARCH-effect simulations.

5.2.1. Low GARCH effect. Let us now discuss simulation results from the low GARCH-effect model. First, we mention the effect of structural changes in time series on the parameter estimation. Later, we compare the performance of all methods in terms of absolute PE.

Estimating a parametric model from data containing a change point will necessarily lead to various biases in estimation. For example, Hillebrand (2005) demonstrates that a change in volatility level ω within a sample drives the GARCH parameter β very close to 1. This is confirmed when we analyse the parameter estimates for parametric and adaptive GARCH at each time point $t \in [250, 1000]$ as depicted on Figure 2, where the mean (solid line), the 10% and 90% quantiles (dotted lines), and the true values (thick dotted line) of the model parameters are provided. The parametric estimates are consistent before breaks starting at $t = 500$, but the GARCH parameter β becomes inconsistent and converges to 1 once data contain breaks, $t > 500$. The locally adaptive estimates are similar to parametric ones before the breaks and become rather imprecise after the first change point, but they are not too far from the true value on average and stay consistent (in the sense that the confidence interval covers the true values). The low precision of estimation can be attributed to rather short intervals used for estimation (cf. Figure 2 for $t < 500$).

Next, we would like to compare the performance of parametric and adaptive estimation methods by means of absolute PE: first for the prediction horizon of one day, $\mathcal{H} = \{1\}$, and later for prediction two weeks ahead, $\mathcal{H} = \{1, \dots, 10\}$. To make the results easier to decipher,




Figure 2. Parameter values estimated by the parametric (top row) and locally adaptive (bottom row) GARCH methods.

we present in what follows PEs averaged over the past month (21 days). The absolute-PE criterion was also used to determine the optimal values of parameters r and ρ (jointly across all simulations and for all $t = 250, \dots, 1000$). The results differ for different models: $r = 0.5$, $\rho = 0.5$ for local constant, $r = 0.5$, $\rho = 1.0$ for local ARCH, and $r = 0.5$, $\rho = 1.5$ for local GARCH.

Let us now compare the adaptively estimated local constant, local ARCH and local GARCH models with the parametric GARCH, which is the best performing parametric model in this set-up. Forecasting one period ahead, the average PEs for all methods and the median lengths of the selected time-homogeneous intervals for adaptive methods are presented on Figure 3 for $t \in [250, 1000]$. First of all, let us observe in the case of the simplest local constant model that even the (median) estimated interval of homogeneity at the end of the first homogeneous period, $1 \leq t < 500$, can actually be shorter than the true one. The reason is that the probability of some 5 or 10 subsequent observations used as I_0 having their sample variance very different from the underlying one increases with the length of the series.

Next, one can notice that all methods are sensitive to jumps in volatility, especially to the first one at $t = 500$: the parametric ones because they ignore a structural break, the adaptive ones because they use a small amount of data after a structural change. In general, the local GARCH performs rather similarly to the parametric GARCH for $t < 650$ because it uses all historical data. After initial volatility jumps, the local GARCH, however, outperforms the parametric one, $650 < t < 775$. Following the last jump at $t = 750$, where the volatility level returns closer to the initial one, the parametric GARCH is best of all methods for some time, $775 < t < 850$, until the adaptive estimation procedure detects the (last) break, and after it, ‘collects’ enough observations for estimation. Then the local GARCH and local ARCH become preferable to the parametric model again, $850 < t$. Interestingly, the local ARCH approximation performs almost as well as both GARCH methods and even outperforms them shortly after structural breaks (except for break at $t = 750$), $600 < t < 775$ and $850 < t < 1000$. Finally, the local constant




Figure 3. Left-hand panel: Low GARCH-effect simulations—absolute prediction errors one period ahead. Right-hand panel: The median lengths of the adaptively selected intervals.




Figure 4. Left-hand panel: Low GARCH-effect simulations—absolute prediction errors 10 periods ahead. Right-hand panel: High GARCH-effect simulations—absolute prediction errors one period ahead.

volatility is lacking behind the other two adaptive methods whenever there is a longer time period without a structural break, but keeps up with them in periods with frequent volatility changes, $500 < t < 650$. All these observations can be documented also by the absolute PE averaged over the whole period $250 \leq t \leq 1000$ (we refer to it as the global PE from now on): the smallest PE is achieved by local ARCH (0.075), then by local GARCH (0.079) and the worst result is from local constant (0.094).

Additionally, all models are compared using the forecasting horizon of 10 days. Most of the results are the same (e.g. parameter estimates) or similar (e.g. absolute PE) to forecasting one period ahead due to the fact that all models rely on at most one past observation. The absolute PEs averaged over one month are summarized for $t \in [250, 1000]$ on Figure 4, which reveals that the difference between local constant volatility, local ARCH and local GARCH models are smaller in this case. As a result, it is interesting to note that: (i) the local constant model becomes a viable alternative to the other methods (it has in fact the smallest global PE 0.107 from all adaptive methods) and (ii) the local ARCH model still outperforms the local GARCH (global

PEs are 0.108 and 0.116, respectively) even though the underlying model is GARCH (with a small value of $\beta = 0.1$ however).

5.2.2. High GARCH effect. Let us now discuss the high GARCH-effect model. One would expect much more prevalent behaviour of both GARCH models, since the underlying GARCH parameter is higher and the changes in the volatility level ω are likely to be small compared to overall volatility fluctuations. Note that the optimal values of tuning constant r and ρ differ from the low GARCH-effect simulations: $r = 0.5$, $\rho = 1.5$ for local constant; $r = 0.5$, $\rho = 1.5$ for local ARCH; and $r = 1.0$, $\rho = 0.5$ for local GARCH.

Comparing the absolute PEs for the one-period-ahead forecast at each time point (Figure 4) indicates that the adaptive and parametric GARCH estimations perform approximately equally well. On the other hand, both the parametric and adaptively estimated ARCH and constant volatility models are lacking significantly. Unreported results confirm, similarly to the low GARCH-effect simulations, that the differences among method are much smaller once a longer prediction horizon of 10 days is used.

6. APPLICATIONS

The proposed adaptive pointwise estimation method will be now applied to real time series consisting of the log-returns of the DAX and S&P 500 stock indices (Sections 6.1 and 6.2). We will again summarize the results concerning both parametric and adaptive methods by the absolute PEs one day ahead averaged over one month. As a benchmark, we employ the parametric GARCH estimated using the last two years of data (500 observations). Since we however do not have the underlying volatility process now, it is approximated by squared returns. Despite being noisy, this approximation is unbiased and provides usually the correct ranking of methods (Andersen and Bollerslev, 1998).

6.1. DAX analysis

Let us now analyse the log-returns of the German stock index DAX from January 1990 till December 2002 depicted at the top of Figure 5. Several periods interesting for comparing the performance of parametric and adaptive pointwise estimates are selected since results for the whole period might be hard to decipher at once.

First, consider the estimation results for years 1991 to 1996. Contrary to later periods, there are structural breaks practically immediately detected by all adaptive methods (July 1991 and June 1992; cf. Stapf and Werner, 2003). For the local GARCH, this differs from less pronounced structural changes discussed later, which are typically detected only with delays of several months. One additional break detected by all methods occurs in October 1994. Note that parameters r and ρ were $r = 0.5$, $\rho = 1.5$ for local constant, $r = 1.0$, $\rho = 1.0$ for local ARCH, and $r = 0.5$, $\rho = 1.5$ for local GARCH.

The results for the period 1991–96 are summarized in the left bottom panel of Figure 5, which depicts the PEs of each adaptive method relative to the PEs of parametric GARCH. First, one can notice that the local constant and local ARCH approximations are preferable till July 1991, where we have less than 500 observations. After the detection of the structural change in June 1991, all adaptive methods are shortly worse than the parametric GARCH due to the limited amount of data used, but then outperform the parametric GARCH till the next structural break in the second half of 1992. A similar behaviour can be observed after the break detected in October 1994,




Figure 5. Top panel: The log-returns of DAX series. Bottom panels: The absolute prediction errors of the pointwise adaptive methods relative to the parametric GARCH errors for predictions one period ahead.

where the local constant and local ARCH models actually outperform both the parametric and adaptive GARCH. In the other parts of the data, the performance of all methods is approximately the same, and even though the adaptive GARCH is overall better than the parametric one, the most interesting fact is that the adaptively estimated local constant and local ARCH models perform equally well. In terms of the global PE, the local constant is best (0.829), followed by the local ARCH (0.844) and local GARCH (0.869). This closely corresponds to our findings in simulation study with low GARCH effect in Section 5.2. Note that for other choices of r and ρ , the global PEs are at most 0.835 and 0.851 for the local constant and local ARCH, respectively. This indicates low sensitivity to the choice of these parameters.

Next, we discuss the estimation results for years 1999 to 2001 ($r = 1.0$ for all methods now). After the financial markets were hit by the Asian crisis in 1997 and the Russian crisis in 1998, the market headed to a more stable state in year 1999. The adaptive methods detected the structural breaks in the autumn of 1997 and 1998. The local GARCH detected them, however, with more than a one-year delay—only during 1999. The results in Figure 5 (right bottom panel) confirm that the benefits of the adaptive GARCH are practically negligible compared to the parametric GARCH in such a case. On the other hand, the local constant and ARCH methods perform slightly better than both GARCH methods during the first presented year (July 1999 to June 2000). From July 2000, the situation becomes just the opposite and the performance




Figure 6. Left-hand panel: The log-returns of S&P 500. Right-hand panel: The absolute prediction errors of the pointwise adaptive methods relative to the parametric GARCH errors for predictions one period ahead.

of the GARCH models is better (parametric and adaptive GARCH estimates are practically the same in this period since the last detected structural change occurred approximately two years ago). Together with previous results, this opens the question of model selection among adaptive procedures as different parametric approximations might be preferred in different time periods. Judging by the global PE, the local ARCH provides slightly better predictions on average than the local constant and local GARCH—despite the ‘peak’ of the PE ratio in the second half of year 2000 (see Figure 5). This, however, depends on the specific choice of loss Λ in (3.6).

Finally, let us mention that the relatively similar behaviour of the local constant and local ARCH methods is probably due to the use of ARCH(1) model, which is not sufficient to capture more complex time developments. Hence, ARCH(p) might be a more appropriate interim step between the local constant and GARCH models.

6.2. S&P 500

Now we turn our attention to more recent data regarding the S&P 500 stock index considered from January 2000 to December 2004; see Figure 6. This period is marked by many substantial events affecting the financial markets, ranging from September 11, 2001, terrorist attacks and the war in Iraq (2003) to the crash of the technology stock-market bubble (2000–02). For the sake of simplicity, a particular time period is again selected: year 2003 representing a more volatile period (the war in Iraq) and year 2004 being a less volatile period. All adaptive methods detected rather quickly a structural break at the beginning of 2003, and additionally they detected a structural break in the second half of 2003, although the adaptive GARCH did so with a delay of more than eight months. The ratios of monthly PE of all adaptive methods to those of the parametric GARCH from January 2003 to December 2004 are summarized on Figure 6 ($r = 0.5$ and $\rho = 1.5$ for all methods).

In the beginning of year 2003, corresponding with 2002 to a more volatile period (see Figure 6), all adaptive methods perform as well as the parametric GARCH. In the middle of year 2003, the local constant and local ARCH models are able to detect another structural change (possibly less pronounced than the one at the beginning of 2003 because of its late detection by the adaptive GARCH). Around this period, the local ARCH shortly performs worse than the parametric GARCH. From the end of 2003 and in year 2004, all adaptive methods starts to outperform the parametric GARCH, where the reduction of the PEs due to the adaptive estimation amounts to 20% on average. All adaptive pointwise estimates exhibit a short period of instability in the first months of 2004, where their performance temporarily worsens to the level of parametric GARCH. This corresponds to ‘uncertainty’ of the adaptive methods about the length of the interval of homogeneity. After this short period, the performance of all adaptive methods is comparable, although the local constant performs overall best of all methods (closely followed by local ARCH) judged by the global PE.

Similarly to the low GARCH-effect simulations and to the analysis of DAX in Section 6.1, it seems that the benefit of pointwise adaptive estimation is most pronounced during periods of stability that follow an unstable period (i.e. year 2004) rather than during a presumably rapidly changing environment. The reason is that, despite possible inconsistency of parametric methods under change points, the adaptive methods tend to have a rather large variance when the intervals of time homogeneity become very short.

7. CONCLUSION

We extend the idea of adaptive pointwise estimation to parametric CH models. In the specific case of ARCH and GARCH, which represent particularly difficult cases due to high data demands and dependence of critical values on underlying parameters, we demonstrate the use and feasibility of the proposed procedure: on the one hand, the adaptive procedure, which itself depends on a number of auxiliary parameters, is shown to be rather insensitive to their choice, and on the other hand, it facilitates the global selection of these parameters by means of fit or forecasting criteria. The real-data applications highlight the flexibility of the proposed time-inhomogeneous models since even simple varying-coefficients models such as constant volatility and ARCH(1) can outperform standard parametric methods such as GARCH(1,1). Finally, the relatively small differences among the adaptive estimates based on different parametric approximations indicate that, in the context of adaptive pointwise estimation, it is sufficient to concentrate on simpler and less data-intensive models such as ARCH(p), $0 \leq p \leq 3$, to achieve good forecasts.

ACKNOWLEDGMENTS

This research was supported by the Deutsche Forschungsgemeinschaft through the SFB 649 ‘Economic Risk’.

REFERENCES

Andersen, T. G. and T. Bollerslev (1998). Answering the skeptics: yes, standard volatility models do provide accurate forecasts. *International Economic Review* 39, 885–905.

- Andreou, E. and E. Ghysels (2002). Detecting multiple breaks in financial market volatility dynamics. *Journal of Applied Econometrics* 17, 579–600.
- Andreou, E. and E. Ghysels (2006). Monitoring disruptions in financial markets. *Journal of Econometrics* 135, 77–124.
- Andrews, D. W. K. (1993). Tests for parameter instability and structural change with unknown change point. *Econometrica* 61, 821–56.
- Bai, J. and P. Perron (1998). Estimating and testing linear models with multiple structural changes. *Econometrica* 66, 47–78.
- Beltratti, A. and C. Morana (2004). Structural change and long-range dependence in volatility of exchange rates: either, neither or both? *Journal of Empirical Finance* 11, 629–58.
- Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics* 31, 307–27.
- Cai, Z., J. Fan and Q. Yao (2000). Functional coefficient regression models for nonlinear time series. *Journal of the American Statistical Association* 95, 941–56.
- Chen, J. and A. K. Gupta (1997). Testing and locating variance changepoints with application to stock prices. *Journal of the American Statistical Association* 92, 739–47.
- Chen, R. and R. J. Tsay (1993). Functional-coefficient autoregressive models. *Journal of the American Statistical Association* 88, 298–308.
- Cheng, M.-Y., J. Fan and V. Spokoiny (2003). Dynamic nonparametric filtering with application to volatility estimation. In M. G. Akritas and D. N. Politis (Eds.), *Recent Advances and Trends in Nonparametric Statistics*, 315–33. Amsterdam: Elsevier.
- Diebold, F. X. and A. Inoue (2001). Long memory and regime switching. *Journal of Econometrics* 105, 131–59.
- Doornik, J. A. (2002). Object-oriented programming in econometrics and statistics using Ox: a comparison with C++, Java and C#. In S. S. Nielsen (Ed.), *Programming Languages and Systems in Computational Economics and Finance*, 115–47. Dordrecht: Kluwer.
- Engle, R. F. (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica* 50, 987–1008.
- Fan, J. and W. Zhang (2008). Statistical models with varying coefficient models. *Statistics and Its Interface* 1, 179–95.
- Franco, C. and J.-M. Zakoian (2007). Quasi-maximum likelihood estimation in GARCH processes when some coefficients are equal to zero. *Stochastic Processes and their Applications* 117, 1265–84.
- Glosten, L. R., R. Jagannathan and D. E. Runkle (1993). On the relation between the expected value and the volatility of the nominal excess return on stocks. *Journal of Finance* 48, 1779–801.
- Hansen, B. and S.-W. Lee (1994). Asymptotic theory for the GARCH(1,1) quasi-maximum likelihood estimator. *Econometric Theory* 10, 29–53.
- Härdle, W., H. Herwatz and V. Spokoiny (2003). Time inhomogeneous multiple volatility modelling. *Journal of Financial Econometrics* 1, 55–99.
- Herwatz, H. and H. E. Reimers (2001). Empirical modeling of the DEM/USD and DEM/JPY foreign exchange rate: structural shifts in GARCH-models and their implications. 2001–83, Discussion Paper SFB 373, Humboldt-Universität zu Berlin, Germany.
- Hillebrand, E. (2005). Neglecting parameter changes in GARCH models. *Journal of Econometrics* 129, 121–38.
- Kokoszka, P. and R. Leipus (2000). Change-point estimation in ARCH models. *Bernoulli* 6, 513–39.
- Laurent, S. and J.-P. Peters (2006). *G@RCH 4.2, Estimating and Forecasting ARCH Models*. London: Timberlake Consultants Press.

Mercurio, D. and V. Spokoiny (2004). Statistical inference for time-inhomogeneous volatility models. *Annals of Statistics* 32, 577–602.

Mikosch, T. and C. Starica (1999). Change of structure in financial time series, long range dependence and the GARCH model. Working Paper, Department of Statistics, University of Pennsylvania. See <http://citeseer.ist.psu.edu/mikosch99change.html>.

Mikosch, T. and C. Starica (2004). Changes of structure in financial time series and the GARCH model. *Revstat Statistical Journal* 2, 41–73.

Nelson, D. B. (1991). Conditional heteroskedasticity in asset returns: a new approach. *Econometrica* 59, 347–70.

Pesaran, M. H. and A. Timmermann (2004). How costly is it to ignore breaks when forecasting the direction of a time series? *International Journal of Forecasting* 20, 411–25.

Sentana, E. (1995). Quadratic ARCH models. *Review of Economic Studies* 62, 639–61.

Spokoiny, V. (1998). Estimation of a function with discontinuities via local polynomial fit with an adaptive window choice. *Annals of Statistics* 26, 1356–78.

Spokoiny, V. (2009a). Multiscale local change-point detection with applications to value-at-risk. *Annals of Statistics* 37, 1405–36.

Spokoiny, V. (2009b). Parameter estimation in time series analysis. WIAS Preprint No. 1404, Weierstrass Institute for Applied Analysis and Stochastics, Berlin, Germany.

Stapf, J. and T. Werner (2003). How wacky is DAX? The changing structure of German stock market volatility. Discussion Paper 2003/18, Deutsche Bundesbank, Germany.

Taylor, S. J. (1986). *Modeling Financial Time Series*. Chichester: Wiley.

APPENDIX: PROOFS

Proof of Corollary 2.1: Given the choice of \mathfrak{z}_α , it directly follows from (2.5). □

Proof of Theorem 3.1: Consider the event $\mathcal{B}_k = \{\hat{I} = I_{k-1}\}$ for some $k \leq K$. This particularly means that I_{k-1} is accepted while $I_k = [T - m_k + 1, T]$ is rejected; i.e. there is $I' = [t', T] \subseteq I_k$ and $\tau \in \mathcal{I}(I_k)$ such that $T_{I_k, \tau} > \mathfrak{z}_k = \mathfrak{z}_{I_k, \mathcal{I}(I_k)}$. For every fixed $\tau \in \mathcal{I}(I_k)$ and $J = I_k \setminus [\tau + 1, T]$, $J^c = [\tau + 1, T]$, it holds by definition of $T_{I_k, \tau}$ that

$$T_{I_k, \tau} \leq L_J(\tilde{\theta}_J) + L_{J^c}(\tilde{\theta}_{J^c}) - L_I(\theta_0) = L_J(\tilde{\theta}_J, \theta_0) + L_{J^c}(\tilde{\theta}_{J^c}, \theta_0).$$

This implies by Theorem 2.1 that $\mathbf{P}_{\theta_0}(T_{I_k, \tau} > 2\mathfrak{z}) \leq \exp\{\epsilon(\lambda, \theta_0) - \lambda\mathfrak{z}\}$. Now,

$$\mathbf{P}_{\theta_0}(\mathcal{B}_k) \leq \sum_{t'=T-m_k+1}^{T-m_0} \sum_{\tau=t'+1}^{T-m_0+1} 2 \exp\{\epsilon(\lambda, \theta_0) - \lambda\mathfrak{z}_k/2\} \leq 2 \frac{m_k^2}{2} \exp\{\epsilon(\lambda, \theta_0) - \lambda\mathfrak{z}_k/2\}.$$

Next, by the Cauchy–Schwartz inequality

$$\begin{aligned} \mathbf{E}_{\theta_0} |L_{I_K}(\tilde{\theta}_{I_K}, \hat{\theta})|^r &= \sum_{k=1}^K \mathbf{E}_{\theta_0} [|L_{I_K}(\tilde{\theta}_{I_K}, \tilde{\theta}_{k-1})|^r \mathbf{1}(\mathcal{B}_k)] \\ &\leq \sum_{k=1}^K \mathbf{E}_{\theta_0}^{1/2} |L_{I_K}(\tilde{\theta}_{I_K}, \tilde{\theta}_{k-1})|^{2r} \mathbf{P}_{\theta_0}^{1/2}(\mathcal{B}_k). \end{aligned}$$

Under the conditions of Theorem 2.1, it follows similarly to (2.6) that

$$\mathbf{E}_{\theta_0} |L_{I_K}(\tilde{\theta}_{I_K}, \tilde{\theta}_{k-1})|^{2r} \leq (m_K/m_{k-1})^{2r} \mathfrak{R}_{2r}^*(\theta_0)$$

for some constant $\mathfrak{R}_{2r}^*(\theta_0)$ and $k = 1, \dots, K$, and therefore,

$$\mathbf{E}_{\theta_0} |L_{I_K}(\tilde{\theta}_{I_K}, \hat{\theta})|^r \leq [\mathfrak{R}_{2r}^*(\theta_0)]^{1/2} \sum_{k=1}^K m_k (m_K/m_{k-1})^r \exp\{\epsilon(\lambda, \theta_0)/2 - \lambda \mathfrak{z}_k/4\}$$

and the result follows by simple algebra provided that $a_1\lambda/4 \geq 1$ and $a_2\lambda/4 > 2$. □

LEMMA A.1. *Let \mathbf{P} and \mathbf{P}_0 be two measures such that the Kullback–Leibler divergence $\mathbf{E} \log(d\mathbf{P}/d\mathbf{P}_0)$, satisfies $\mathbf{E} \log(d\mathbf{P}/d\mathbf{P}_0) \leq \Delta < \infty$. Then for any random variable ζ with $\mathbf{E}_0\zeta < \infty$, it holds that $\mathbf{E} \log(1 + \zeta) \leq \Delta + \mathbf{E}_0\zeta$.*

Proof: By simple algebra one can check that for any fixed y the maximum of the function $f(x) = xy - x \log x + x$ is attained at $x = e^y$ leading to the inequality $xy \leq x \log x - x + e^y$. Using this inequality and the representation $\mathbf{E} \log(1 + \zeta) = \mathbf{E}_0\{Z \log(1 + \zeta)\}$ with $Z = d\mathbf{P}/d\mathbf{P}_0$ we obtain

$$\begin{aligned} \mathbf{E} \log(1 + \zeta) &= \mathbf{E}_0\{Z \log(1 + \zeta)\} \leq \mathbf{E}_0(Z \log Z - Z) + \mathbf{E}_0(1 + \zeta) \\ &= \mathbf{E}_0(Z \log Z) + \mathbf{E}_0\zeta - \mathbf{E}_0Z + 1. \end{aligned}$$

It remains to note that $\mathbf{E}_0Z = 1$ and $\mathbf{E}_0(Z \log Z) = \mathbf{E} \log Z$. □

Proof of Theorem 4.1: Lemma A.1 applied with $\zeta = \varrho(\hat{\theta}, \theta)/\mathbf{E}_\theta \varrho(\hat{\theta}, \theta)$ yields the result in the view of

$$\begin{aligned} \mathbf{E}_\theta(Z_{I,\theta} \log Z_{I,\theta}) &= \mathbf{E} \log Z_{I,\theta} = \mathbf{E} \sum_{t \in I} \log \frac{p[Y_t, g(X_t)]}{p[Y_t, g(X_t(\theta))]} \\ &= \mathbf{E} \sum_{t \in I} \mathbf{E} \left\{ \log \frac{p[Y_t, g(X_t)]}{p[Y_t, g(X_t(\theta))]} \middle| \mathcal{F}_{t-1} \right\} = \mathbf{E} \Delta_{I_k}(\theta). \end{aligned} \quad \square$$

Proof of Corollary 4.1: It is Theorem 4.1 formulated for $\varrho(\theta', \theta) = L_I(\theta', \theta)$. □

Proof of Theorem 4.2: The first inequality follows from Corollary 4.1, the second one from condition (3.4) and the property $x \geq \log x$ for $x > 0$. □

Proof of Theorem 4.3: Let $\hat{k} = k > k^*$. This means that I_k is not rejected as homogeneous. Next, we show that for every $k > k^*$ the inequality $T_{I_k, \tau} \leq T_{I_k, \mathcal{S}(I_k)} \leq \mathfrak{z}_k$ with $\tau = T - m_{k^*} = T - |I_{k^*}|$ implies $L_{I_{k^*}}(\tilde{\theta}_{I_{k^*}}, \tilde{\theta}_{I_k}) \leq \mathfrak{z}_{k^*}$. Indeed with $J = I_k \setminus I_{k^*}$, this means that, by construction, $\mathfrak{z}_k \leq \mathfrak{z}_{k^*}$ for $k > k^*$ and

$$\mathfrak{z}_k \geq T_{I_k, \tau} = L_{I_{k^*}}(\tilde{\theta}_{I_{k^*}}, \tilde{\theta}_{I_k}) + L_J(\tilde{\theta}_J, \tilde{\theta}_{I_k}) \geq L_{I_{k^*}}(\tilde{\theta}_{I_{k^*}}, \tilde{\theta}_{I_k}).$$

It remains to note that

$$|L_{I_{k^*}}(\tilde{\theta}_{I_{k^*}}, \hat{\theta})|^r \leq |L_{I_{k^*}}(\tilde{\theta}_{I_{k^*}}, \hat{\theta}_{I_{k^*}})|^r \mathbf{1}(\hat{k} < k^*) + \mathfrak{z}_{k^*}^r \mathbf{1}(\hat{k} > k^*),$$

which obviously yields the assertion. □

Dynamic semiparametric factor models in risk neutral density estimation

Enzo Giacomini · Wolfgang Härdle ·
Volker Krättschmer

Received: 1 March 2009 / Accepted: 31 August 2009 / Published online: 18 September 2009
© Springer-Verlag 2009

Abstract Dynamic semiparametric factor models (DSFM) simultaneously smooth in space and are parametric in time, approximating complex dynamic structures by time invariant basis functions and low dimensional time series. In contrast to traditional dimension reduction techniques, DSFM allows the access of the dynamics embedded in high dimensional data through the lower dimensional time series. In this paper, we study the time behavior of risk assessments from investors facing random financial payoffs. We use DSFM to estimate risk neutral densities from a dataset of option prices on the German stock index DAX. The dynamics and term structure of risk neutral densities are investigated by Vector Autoregressive (VAR) methods applied on the estimated lower dimensional time series.

Keywords Dynamic factor models · Dimension reduction · Risk neutral density

1 Introduction

Large datasets containing various samples of high dimensional observations became common in diverse fields of science with advances in measurement and computational techniques. In many applications the data come in curves, i.e., as observations of discretized values of smooth random functions, presenting evident functional structure. In these cases, it is natural to perform statistical inference using functional data analysis techniques.

E. Giacomini (✉) · W. Härdle · V. Krättschmer
CASE—Center for Applied Statistics and Economics, Humboldt-Universität zu Berlin,
Spandauerstr. 1, 10178 Berlin, Germany
e-mail: enzogiacomini@gmail.com

V. Krättschmer
Institute of Mathematics, Technische Universität Berlin, Straße des 17. Juni 136, 10623 Berlin,
Germany

Consider a dataset $\{(Y_{jt}, X_{jt})\}$, $j = 1, \dots, J_t$, $t = 1, \dots, T$, containing noisy samples of a real valued smooth random function $\mathcal{F} \in L_2(\mathcal{X})$, $\mathcal{X} \subseteq \mathbb{R}^d$, $d \in \mathbb{N}$, evaluated at unbalanced design points as

$$Y_{jt} = \mathcal{F}_t(X_{jt}) + \varepsilon_{jt}, \quad (1.1)$$

where ε_{jt} denote unknown zero-mean error terms and $\{\mathcal{F}_t\}$ are realizations of \mathcal{F} . Each sample $S_t = \{(Y_{jt}, X_{jt}) : j = 1, \dots, J_t\}$, $t = 1, \dots, T$, may correspond to observations on, e.g., different individuals, time periods or experimental conditions. Examples in biomedicine are measurements of growth curves and brain potentials across individuals, see Kneip and Gasser (1992) and Gasser and Kneip (1995), in econometrics such are expenditures across households and implied volatilities across trading days, see Kneip (1994) and Fengler et al. (2007).

A large branch of functional data analysis concentrates on approximating \mathcal{F} by lower dimensional objects. Distributions on function spaces are highly complex objects and dimension reduction techniques present a feasible and interpretable approach for investigating them. Functional principal components analysis (FPCA), based on the Karhunen–Loève expansion of \mathcal{F} is the most prominent and widely used dimension reduction technique, see Rice and Silverman (1991) and Ramsay and Dalzell (1991).

Asymptotic results on FPCA have been obtained by Dauxois et al. (1982) and Hall et al. (2006) for observed functional data $\{\mathcal{F}_t\}$. For non-observable data, the standard approach is to perform FPCA on presmoothed $\{\widehat{\mathcal{F}}_t\}$, see Benko et al. (2009) for recent developments. In practical applications, however, presmoothing may suffer from design-sparseness, see Cont and Fonseca (2002) and Fengler et al. (2007).

In general lines, previous literature combines PCA and dimension reduction with presmoothing for effective dimensional space at fixed time horizon. Various applications, however, involve the dynamics of the unobserved random functions, calling for dimension reduction techniques that smooth in space and are parametric in time.

In this paper, we investigate the dynamics of $\{\mathcal{F}_t\}$ by reducing dimensionality without presmoothing. \mathcal{F}_t is considered as a linear combination of $L + 1 \ll T$ unknown smooth basis functions $m_l \in L_2(\mathcal{X})$, $l = 0, \dots, L$:

$$\mathcal{F}_t(X_{jt}) = \sum_{l=0}^L Z_{lt} m_l(X_{jt}), \quad (1.2)$$

where $Z_t = (Z_{0t}, \dots, Z_{Lt})^\top$ is an unobservable random vector taking values on \mathbb{R}^{L+1} with $Z_{0t} = 1$. Defining the tuple of functions $m = (m_0, \dots, m_L)^\top$, the Dynamic Semiparametric Factor Model (DSFM) reads as

$$Y_{jt} = Z_t^\top m(X_{jt}) + \varepsilon_{jt}. \quad (1.3)$$

The basis functions are estimated nonparametrically avoiding specification issues. Their estimation is performed simultaneously with Z_t , i.e., the smoothing is transferred directly to m_l and design-sparseness issues become secondary. In addition, the random process $\{Z_t\}$ is allowed to be non-stationary. Park et al. (2009) show that

under (1.2) the autocorrelation structures of $\{\widehat{Z}_t\}$ and $\{Z_t\}$ are asymptotically equivalent; therefore, no loss is incurred by inferring the dynamics from the estimated $\{\widehat{Z}_t\}$, and there is no payment for not knowing the true $\{Z_t\}$. This result is essential for investigating cointegration between dynamical systems, see Brüggemann et al. (2008) for an econometric application.

Note that the common regressors model, Kneip (1994), also represents unobservable functions by (1.2). There are, however, crucial differences between the DSFM and common regressors:

1. In DSFM, $\{Z_t\}$ is a (non-stationary) random process with autocovariance structure inferable from $\{\widehat{Z}_t\}$.
2. DSFM is implementable in unbalanced designs.
3. DSFM avoids presmoothing by transferring the smoothing to the basis functions.

Thus DSFM goes beyond traditional dimension reductions techniques (FPCA and common regressors) as it captures structural dynamics embedded in the observations.

In economics, there is substantial interest in the behavior (over time) of investors facing risks and its relation to macroeconomic and financial indicators. The knowledge about the dynamics of risk assessments from investors is essential for many applications ranging from pricing of illiquid instruments to risk management.

Option prices contain information on risk assessments from investors facing future financial payoffs, summarized in the risk neutral densities q , see Ait-Sahalia and Lo (1998). An European call option with price C_t at time $0 \leq t \leq T$, maturity date $T > 0$ and strike $K > 0$ is a financial instrument that delivers the random payoff $(S_T - K)^+$ at time T where S_t is the price of an underlying asset at time t . Breeden and Litzenberger (1978) show that under no arbitrage assumptions the risk neutral density is obtained from the European call price function C_t through the relation

$$q_{t,T}(s_T | s_t) = e^{r(T-t)} \frac{\partial^2 C_t(s_t, r, K, T-t)}{\partial K^2} \Big|_{K=s_T}, \tag{1.4}$$

where $r > 0$ is interest rate, see Sect. 4 for details.

We estimate risk neutral densities based on observed intraday prices of calls on the German stock index (DAX). Each observation consists of a price Y_{jt} on a design point $X_{jt} = (\kappa_{jt}, \tau_{jt})^\top$ where $j = 1, \dots, J_t$, denote the transactions at day $t = 1, \dots, T$, κ is the moneyness, a monotone transformation of strikes K , and $\tau = T - t$ is the time to maturity associated with the option. Stock exchange regulations impose prespecified values for tradable maturities resulting in degenerated designs, see Fig. 1.

Following Ait-Sahalia and Lo (1998) and Fengler et al. (2007), call prices are transformed into log-implied volatilities $\widetilde{Y}_{jt} = \log C_{BS}^{-1}(Y_{jt})$, where C_{BS} is the Black–Scholes call price function defined in Sect. 4. These are assumed as discretized noisy values of the log-implied volatility surface evaluated at $\{X_{jt}\}$:

$$\widetilde{Y}_{jt} = \log \mathcal{V}_t(X_{jt}) + \varepsilon_{jt}, \tag{1.5}$$

where $\mathcal{V} \in L_2(\mathcal{X})$, $\mathcal{X} \subset \mathbb{R}_+^2$, is a smooth random function, called the implied volatility surface, and ε_{jt} is an error term. The realizations $\{\mathcal{V}_t\}$ are filtered out from the data with DSFM and, remarking that C_{BS} is a function of K , the risk neutral densities are




Fig. 1 Samples S_t , $t = 1, \dots, 22$, of DAX call prices traded on January 2001 (left). Corresponding unbalanced design $\{X_{jt}\}$ (right)

obtained by (1.4) with $C_{BS}(\widehat{V})$ as an estimator for C_t . The dynamics of the estimated $\{\widehat{q}_{t,T}\}$ is analyzed based on the autocorrelation structure of $\{\widehat{Z}_t\}$.

In the sequel, the DSFM estimation method and its asymptotic properties are described (Sect. 2). In Sect. 3, the risk neutral densities are defined, and in Sect. 4 they are estimated from observed prices of European call options on the DAX index

(ODAX dataset). Their dynamic structure is then analyzed by vector autoregressive models.

2 Estimation method

Consider a dataset $\{(Y_{jt}, X_{jt})\}$, $j = 1, \dots, J_t, t = 1, \dots, T$, such that

$$Y_{jt} = \sum_{l=0}^L Z_{lt} m_l(X_{jt}) + \varepsilon_{jt}, \tag{2.1}$$

where ε_{jt} are unknown error terms with $E[\varepsilon_{jt}] = 0$ and $E[\varepsilon_{jt}^2] < \infty$. The variables $X_{11}, \dots, X_{T,J_T}, \varepsilon_{1,1}, \dots, \varepsilon_{T,J_T}$ are independent. Here $Z_t = (Z_{0t}, \dots, Z_{Lt})^\top$ is an unobservable random vector taking values on \mathbb{R}^{L+1} with $Z_{0t} = 1$ and $m_l \in L_2(\mathcal{X})$, $l = 0, \dots, L$, are unknown smooth functions, called basis functions, mapping $\mathcal{X} \subseteq \mathbb{R}^d, d \in \mathbb{N}$, into real values.

Following Park et al. (2009), the basis functions are estimated using a series expansion. Defining K normed functions $\psi_k : \mathcal{X} \rightarrow \mathbb{R}, \int_{\mathcal{X}} \psi_k^2(x) dx = 1, k = 1, \dots, K$, and an $((L + 1) \times K)$ matrix of coefficients $\Gamma = (\gamma_{l,k}), \gamma_{l,k} \in \mathbb{R}$, the tuple of functions $m = (m_0, \dots, m_L)^\top$ is approximated by $\Gamma^\top \psi$ where $\psi = (\psi_1, \dots, \psi_K)^\top$. For simplicity of notation, we assume that $J_t = J$ does not depend on t . We define the least squares estimators as

$$(\widehat{\Gamma}, \widehat{Z}) = \arg \min_{\Gamma \in \mathcal{G}, Z \in \mathcal{Z}} \sum_{t=1}^T \sum_{j=1}^J \{Y_{jt} - Z_t^\top \Gamma \psi(X_{jt})\}^2, \tag{2.2}$$

where $\mathcal{G} = \mathcal{M}(L + 1, K), \mathcal{Z} = \{Z \in \mathcal{M}(T, L + 1) : Z_{0t} = 1\}$ and $\mathcal{M}(a, b)$ is the set of all $(a \times b)$ matrices. The basis functions m are estimated by $\widehat{m} = \widehat{\Gamma}^\top \psi$.

Theorem (2.1) gives the asymptotic behavior of the least squares estimators $(\widehat{\Gamma}, \widehat{Z})$. See Park et al. (2009) for the proof.

Theorem 2.1 *Suppose that DSFM holds and that $(\widehat{\Gamma}, \widehat{Z})$ is defined by (2.2). Under Assumptions (A1)–(A8), see Appendix, it holds for $K, J \rightarrow \infty$:*

$$\frac{1}{T} \sum_{1 \leq t \leq T} \|\widehat{Z}_t^\top \widehat{\Gamma} - Z_t^\top \Gamma^*\|^2 = \mathcal{O}_P(\delta_K^2 + \xi^2).$$

See (A5) and (A8) for the definitions of δ_K and ξ . Note that the model (2.1) is only identifiable up to linear transformations. Consider an $((L + 1) \times (L + 1))$ regular matrix $B = (b_{ij})$ with $b_{1j} = \delta_{1j}$ and $b_{i1} = \delta_{i1}$ for $i, j = 1, \dots, L + 1$, where $\delta_{ij} = 1(i = j)$. Define $Z_t^* = B^\top Z_t, m^* = B^{-1}m$. Then from (1.2)

$$\mathcal{F}_t(X) = Z_t^\top m(X) = Z_t^\top B B^{-1} m(X) = Z_t^{*\top} m^*(X)$$

for $X \in \mathcal{X}$. On the other hand, it is always possible to chose orthonormal basis functions by setting $m^* = Hm$ where H is an orthogonal matrix.

Theorem (2.2) states that for any \widehat{Z}_t there exists a random matrix B such that the autocovariances of $\{\widetilde{Z}_t\}$, $\widetilde{Z}_t = B^\top \widehat{Z}_t$, are asymptotically equivalent to the autocovariances of the true unobservable $\{Z_t\}$. This equivalence is transferred to classical estimation and testing procedures in the context of, e.g., vector autoregressive models and, in particular, justifies inference based on $\{\widetilde{Z}_t\}$ when $\{Z_t\}$ is a VAR process. Define for $H_t \in \mathcal{Z}$, $t = 1, \dots, T$: $\overline{H} = T^{-1} \sum_{t=1}^T H_t$, $H_{c,t} = H_t - \overline{H}$ and $H_{n,t} = (T^{-1} \sum_{s=1}^T H_{c,s} H_{c,s}^\top)^{-1/2} H_{c,t}$.

Theorem 2.2 *Suppose that DSFM holds and that $(\widehat{\Gamma}, \widehat{Z})$ is defined by (2.2). Under Assumptions (A1)–(A11), see Appendix, there exists a random matrix B such that for $h \neq 0$, $h_d = \max(1, 1 - h)$, $h_u = \max(T, T - h)$ and $T \rightarrow \infty$:*

$$\frac{1}{T} \sum_{t=h_d}^{h_u} \widetilde{Z}_{c,t} (\widetilde{Z}_{c,t+h} - \widetilde{Z}_{c,t})^\top - \frac{1}{T} \sum_{t=h_d}^{h_u} Z_{c,t} (Z_{c,t+h} - Z_{c,t})^\top = \mathcal{O}_P(T^{-1/2}),$$

where $\widetilde{Z}_t = B^\top \widehat{Z}_t$. Moreover,

$$\frac{1}{T} \sum_{t=h_d}^{h_u} \widetilde{Z}_{n,t} \widetilde{Z}_{n,t+h}^\top - \frac{1}{T} \sum_{t=h_d}^{h_u} Z_{n,t} Z_{n,t+h}^\top = \mathcal{O}_P(T^{-1/2}).$$

See Park et al. (2009) for the proof. Note that, in contrast to FPCA, DSFM does not require stationarity neither for $\{Z_t\}$ nor for $\{\varepsilon_t\}$, but only weak assumptions on the average behavior of Z_t , like being a martingale difference, see Appendix.

3 Risk neutral density estimation

3.1 Risk neutral densities

Consider a financial market with one risky asset and one riskless bond with constant interest rate $r > 0$. Let the price of the asset traded on the market be described by the real valued random process $\{S_t\}$, $t = [0, T]$, $T < \infty$, on a filtered probability space $(\Omega, \{\mathcal{F}_t\}, \mathbb{P})$ with $\mathcal{F}_t = \sigma(S_u, u \leq t)$ and $\mathcal{F}_0 = \{\emptyset, \Omega\}$. Assume further no arbitrage in the financial market in the sense that there exists a (risk neutral) probability measure \mathbb{Q} equivalent to \mathbb{P} under which the discounted price process $\{e^{-rt} S_t\}$ is a martingale.

A European call option at strike $K > 0$ is a financial instrument that pays $\Psi(S_T) = (S_T - K)^+$ at time T . By the risk-neutral valuation principle w.r.t. \mathbb{Q} , the price C_t of a European call option at time t is defined to be

$$C_t = e^{-r(T-t)} E^{\mathbb{Q}}[\Psi(S_T) | \mathcal{F}_t]. \tag{3.1}$$

Assuming that $\{S_t\}$ is a \mathbb{Q} -Markov process and denoting the \mathbb{P} -density of \mathbb{Q} by π , the price can be rewritten as

$$C_t = e^{-r(T-t)} E[\Psi(S_T) \mathcal{K}_\pi^t(S_t, S_T) | S_t],$$

where E denotes the expectation under \mathbb{P} and $\mathcal{K}_\pi^t(S_t, S_T) \stackrel{\text{def.}}{=} \frac{E[\pi|S_t, S_T]}{E[\pi|S_t]}$. The conditional risk neutral distribution of S_T is defined as

$$Q_{S_T|S_t=s_t}([S_T \leq x]) \stackrel{\text{def.}}{=} \int_{-\infty}^x \mathcal{K}_\pi^t(s_t, \cdot) dP_{S_T|S_t=s_t}, \tag{3.2}$$

where $P_{S_T|S_t=s_t}$ is the conditional distribution of S_T under $S_t = s_t$. Specializing to the following two factor model, we assume that the price process has dynamics given by

$$dS_t = S_t \mu(Y_t) dt + S_t \sigma(Y_t) dW_t^1,$$

here W^1 is a standard \mathbb{P} -Brownian motion and Y denotes an external economic factor process modeled by

$$dY_t = g(Y_t) + \rho dW_t^1 + \bar{\rho} dW_t^2,$$

where $\rho \in [-1, 1]$ is some correlation factor, $\bar{\rho} \stackrel{\text{def.}}{=} \sqrt{1 - \rho^2}$ and W^2 is a standard \mathbb{P} -Brownian motion independent of W^1 under \mathbb{P} . Market models of this type are popular in mathematical finance and economics, in particular, if Y follows an Ornstein-Uhlenbeck dynamics with mean reversion term $g(y) = \iota(\theta - y)$ for constants $\theta \geq 0$ and $\iota > 0$. Moreover, $\{S_t\}$ is a \mathbb{Q} -Markov process for any \mathbb{Q} , see Hernández-Hernández and Schied (2007) and the conditional risk neutral distribution $Q_{S_T|S_t=s_t}$ has a density function denoted by $q_{t,T}(\cdot|s_t)$. Hence, recalling (3.1), the call prices can be expressed as

$$C_t(s_t, r, K, T - t) = e^{-r(T-t)} \int (s_T - K)^+ q_{t,T}(s_T|s_t) ds_T.$$

We assume that the observed prices in the financial market are built based on the risk neutral valuation principle w.r.t. an unknown risk neutral measure \mathbb{Q} . Our interest lies in estimating the conditional risk neutral distribution $Q_{S_T|S_t=s_t}$, or equivalently the risk neutral density function $q_{t,T}(\cdot|s_t)$, implied by \mathbb{Q} through (3.2).

3.2 Estimation

Adapting Breeden and Litzenberger (1978), one can show that the risk neutral density function $q_{t,T}(\cdot|s_t)$ is obtained as the second derivative of the call price function C_t with respect to strike K

$$q_{t,T}(s_T|s_t) = e^{r\tau} \left. \frac{\partial^2 C_t(s_t, r, K, \tau)}{\partial K^2} \right|_{K=s_T}, \tag{3.3}$$

where $\tau = T - t$ is the time to maturity.

The unknown price function C_t might be smoothed out of price observations and used in (3.3) to recover risk neutral densities. Here we follow the semiparametric approach from Ait-Sahalia and Lo (1998) where the smoothing is carried out in the space of implied volatilities.

The implied volatility surface is the function $v_t : \mathbb{R}_+^2 \rightarrow \mathbb{R}_+$ satisfying for all $(K, \tau) \in \mathbb{R}_+^2$

$$C_t(s_t, r, K, \tau) = C_{BS}\{s_t, r, K, \tau, v_t(K, \tau)\}, \tag{3.4}$$

where $C_{BS}(s, r, K, \tau, v) = s\Phi(d_1) - Ke^{-r\tau}\Phi(d_2)$ is the Black–Scholes price of Ψ with strike K and maturity τ , $\Phi(\cdot)$ is the cumulative distribution function of the standard normal distribution, $d_1 = \{\log(\frac{s}{K}) + (r + \frac{1}{2}v^2)\tau\}/(v\sqrt{\tau})$ and $d_2 = d_1 - v\sqrt{\tau}$.

More generally, the implied volatility surface is considered a smooth random function $\mathcal{V} \in L_2(\mathcal{X})$ on the space $\mathcal{X} \subset \mathbb{R}^2$ of strikes K and maturities τ . Combining (3.3) and (3.4), the functional random variable $\mathcal{H} \in L_2(\mathcal{X})$, called the risk neutral (RN) surface, is defined as

$$\begin{aligned} \mathcal{H}(s, r, K, \tau, \mathcal{V}) &= e^{r\tau} D^2 C_{BS}(s, r, K, \tau, \mathcal{V}) \\ &= \varphi(d_2) \left\{ \frac{1}{K\sqrt{\tau}\mathcal{V}} + \frac{2d_1}{\mathcal{V}} D\mathcal{V} + K\sqrt{\tau} \frac{d_1 d_2}{\mathcal{V}} (D\mathcal{V})^2 + K\sqrt{\tau} D^2 \mathcal{V} \right\}, \end{aligned} \tag{3.5}$$

where D^m denotes the m th partial derivative with respect to K and $\varphi(\cdot)$ the probability density function of the standard normal distribution. The explicit derivation of (3.5) and a detailed treatment of implied volatilities can be found in Hafner (2004) and Fengler (2005). Clearly, lower dimension objects describing \mathcal{V} may be used to analyze the RN surface \mathcal{H} .

A functional dataset containing realizations of the implied volatility surface \mathcal{V} is, however, not available, as in an exchange only discretized values of \mathcal{V}_t corrupted by noise are registered from trades. On each day $t = 1, \dots, T$ there are J_t options traded, each intraday trade $j = 1, \dots, J_t$ corresponds to an observed option price Y_{jt} at a pair of moneyness κ and maturities τ , $X_{jt} = (\kappa_{jt}, \tau_{jt})^\top$ where $\kappa = e^{r\tau} K/s_t$. Let $C_{BS}(v) = C_{BS}(v; s, r, K, \tau)$ denote the Black–Scholes price as a function of v with all other arguments held constant. As $C_{BS}(v)$ is continuous and monotone in v with inverse C_{BS}^{-1} , the observed implied volatility associated with trade j at day t is then $v_{jt} = C_{BS}^{-1}(Y_{jt})$. Figure 2 shows the implied volatilities from options on the German Stock Index DAX traded on 2 May 2000, the sparse and degenerated design is caused by regulation imposed by stock exchanges on the tradable maturities from call options.

For numerical tractability, see Fengler et al. (2007), observations v_{jt} are transformed into log-implied volatilities $\tilde{Y}_{jt} = \log v_{jt}$ and based on $\{(\tilde{Y}_{jt}, X_{jt})\}$, we use DSFM to model

$$\tilde{Y}_{jt} = Z_t^\top m(X_{jt}) + \varepsilon_{jt}. \tag{3.6}$$

The implied volatility surface at t is estimated by $\hat{\mathcal{V}}_t = \exp(\hat{Z}_t^\top \hat{\Gamma} \psi)$, recall (2.2). The RN surface is estimated using (3.5) by $\hat{\mathcal{H}}_t = \mathcal{H}(s_t, r, K, \tau, \hat{\mathcal{V}}_t)$. The dynamics of the unobservable sequence of RN surfaces $\{\mathcal{H}_t\}$ implied in the observations may be investigated by analyzing the lower dimensional $\{\hat{Z}_t\}$.




Fig. 2 Implied volatilities (*left*) and data design (*right*), ODAx on 2 May 2000

Table 1 Descriptive statistics, number of intraday observations $J_t, t = 1, \dots, 253$

Mean	Std. dev.	Max	Min
2845.92	1589.90	11298	616

4 Application

In this section, the implied volatility and risk neutral surfaces are estimated with DSFM from intraday prices of calls on the DAX index, i.e., S_t represents the value of the DAX index at time t . The dataset contains prices observed from 1 Jan. 2001 to 1 Jan. 2002 corresponding to $T = 253$ trading days. The descriptive statistics of the number of intraday observations J_t are in Table 1, the total number of intraday observations across days is $\sum_{t=1}^T J_t = 720017$.

Tensor B-splines, quadratic in τ and cubic in κ directions placed on 8×6 knots, are used for the series estimators of m . The number of basis functions is chosen based on

$$EV(L) = 1 - \frac{\sum_{t=1}^T \sum_{j=1}^{J_t} \{\tilde{Y}_{jt} - \hat{Z}_t^\top \hat{m}(X_{jt})\}^2}{\sum_{t=1}^T \sum_{j=1}^{J_t} (\tilde{Y}_{jt} - \bar{Y})^2},$$

where $\bar{Y} = (\sum_{t=1}^T \sum_{j=1}^{J_t} \tilde{Y}_{jt}) / \sum_{t=1}^T J_t$. The value $EV(L)$ may be interpreted as the ratio of variation explained by the model to total variation. As established by numerous simulations in Park et al. (2009), the order of the splines and number of knots have negligible influence on $EV(L)$.

4.1 Simulation

The choice of the number of basis functions based on the explained variation criteria is validated by a small simulation study. Datasets $\{(Y_{jt}, X_{jt})\}$ are generated following

$$\begin{aligned}
 Y_{jt} &= \sum_{l=0}^{L^*} Z_{lt} m_l(X_{jt}) + \varepsilon_{jt}, \quad j = 1, \dots, J, \quad t = 1, \dots, T, \\
 \varepsilon_{jt} &\sim N(0, \sigma_\varepsilon^2), \\
 X_{jt} &\sim U([0, 1]^2),
 \end{aligned}
 \tag{4.1}$$

where ε_{jt} and X_{jt} are i.i.d. For $\zeta_t = (Z_{1t}, \dots, Z_{L^*t})^\top$, with 0_d denoting the $(d \times 1)$ vector of zeros and I_d the d identity matrix we define

$$\begin{aligned}
 Z_t &= (1, \zeta_t)^\top, \\
 \zeta_t &= A_{L^*} \zeta_{t-1} + u_t, \\
 u_t &\sim N(0_{L^*}, \sigma_u^2 I_{L^*}),
 \end{aligned}$$

where u_t is i.i.d. and A_{L^*} is a square matrix containing the first L^* rows and L^* columns from A ,

$$A = \begin{pmatrix} 0.95 & -0.2 & 0 & 0.1 \\ 0 & 0.8 & 0.1 & 0.2 \\ 0.1 & 0 & 0.6 & -0.1 \\ 0 & 0.1 & -0.2 & 0.5 \end{pmatrix}.$$

The basis functions are defined as

$$\begin{aligned}
 m_0(\kappa, \tau) &= 1, \\
 m_1(\kappa, \tau) &= 3.46(\kappa - 0.5), \\
 m_2(\kappa, \tau) &= 9.45\{(\kappa - 0.5)^2 + (\tau - 0.5)^2\} - 1.6, \\
 m_3(\kappa, \tau) &= 1.41 \sin(2\pi\tau), \\
 m_4(\kappa, \tau) &= 1.41 \cos(2\pi\kappa),
 \end{aligned}$$

and are close to orthogonal, enhancing similar choice from Park et al. (2009). The value L^* denotes the true number of dynamic basis functions.

Setting $T = 500$, $J = 100$, $\sigma_\varepsilon = 0.05$, and $\sigma_u = 0.1$, $i = 1, \dots, 100$ samples following (4.1) are generated with $L^* = 2, 3$ and 4. Each of them is estimated by DSFM with $L = 1, \dots, 6$, and the corresponding $EV_i(L)$ is computed. The average explained variation under the true L^* , defined as $\mathcal{EV}(L; L^*) = \frac{1}{100} \sum_i EV_i(L)$, is also calculated.

Table 2 shows $\mathcal{EV}(L; L^*)$ and indicates that the increase in the average explained variation between estimation with L^* and $L^* + 1$ dynamic basis functions, $\mathcal{EV}(L^* + 1; L^*) - \mathcal{EV}(L^*; L^*)$, is close to zero across values of L^* . Therefore,

Table 2 Average explained variation $\mathcal{E}\mathcal{V}(L; L^*)$ based on 100 samples from (4.1), across number of dynamic basis functions used in the estimation L and the true L^*

$\mathcal{E}\mathcal{V}(L; L^*)$		L^*		
		2	3	4
L	1	0.86	0.75	0.71
	2	0.99	0.90	0.89
	3	0.99	0.99	0.97
	4	0.99	0.99	0.99
	5	0.99	0.99	0.99

Table 3 Number of basis functions and explained variation

L	1	2	3	4	5
$EV(L)$	0.77	0.97	0.98	0.98	0.98

for DSFM estimation, we select the smallest L such that $EV(L - 1) < EV(L) \approx EV(L + 1)$.

4.2 Results

The implied volatility and RN surfaces are estimated with DSFM as in (3.6) with $L = 3$. Table 3 shows that the addition of the fourth or fifth dynamic basis function results in negligible increase in $EV(L)$.

Following Fengler et al. (2007) and Park et al. (2009), the estimated \widehat{Z}_t and \widehat{m} are respectively transformed and orthonormalized so that $\{\widehat{Z}_t^\top \widehat{m}_l\}$ has a larger contribution than $\{\widehat{Z}_{(l+1)t}^\top \widehat{m}_{l+1}\}$, $l = 1, \dots, L - 1$, to the total variation $\sum_{t=1}^T \int \widehat{Z}_t^\top \widehat{m}$. This transformation aims to improve the interpretation of the basis functions in the analysis of the dynamics of implied volatility surfaces. In the analysis of risk neutral surfaces dynamics, however, it does not present a clear advantage. The covariance structures from $\{\widehat{Z}_t\}$ and $\{Z_t\}$ are then asymptotically equivalent up to orthogonal transformations.

Figures 3 and 4 depict the estimated loading factors series $\{\widehat{Z}_t\}$ and basis functions \widehat{m}_l . The upward and downward peaks observed in \widehat{Z}_{2t} occur on days 6 Feb. 2001 and 5 Nov. 2001 and are caused respectively by extremely unbalanced design and low price levels. The first day has $J_t = 1697$ observations concentrated on short maturities, while the latter has $J_t = 3268$ with very low prices at high maturities.

From (3.5), we obtain a sequence of RN surfaces $\{\widehat{\mathcal{H}}_t\}$, $t = 1, \dots, 253$. We define $\widehat{\mathcal{H}}_t(\kappa, \tau)$ as $\mathcal{H}(\kappa, \tau; s_t, r, \widehat{V}_t)$ where $\kappa = e^{r\tau} K / s_t$. Figure 5 shows $\widehat{\mathcal{H}}_t(\kappa, \tau)$ across moneyness κ and maturity τ at t corresponding to 10 Jul. 2001.

In a first step, we investigate the covariance structure of $\{\widehat{Z}_t\}$ by means of VAR analysis. Table 4 presents the parameters from the VAR(2) model fitted on $\{\widehat{Z}_t\}$. The order 2 is selected based on Akaike (AIC), Schwarz (SC) and Hannan–Quinn (HQ) criteria, see Table 5. Moreover, the VAR(2) model is stationary as the roots of the characteristic polynomial lie inside of the unit circle.

A natural issue is to analyze the dependences between $\{Z_t\}$ and the shape of the RN surfaces $\{\widehat{\mathcal{H}}_t\}$. In order to investigate this relation, we compute the skewness

Fig. 3 Estimated $\{\widehat{Z}_{lt}\}$, $l = 1, 2, 3$ (top to bottom)





Fig. 4 Estimated basis functions $\widehat{m}_l, l = 0, \dots, 3$, clockwise

γ and excess kurtosis η of $\widehat{q}_{t,T}(\cdot|s_t)$ across t for a maturity τ where $\widehat{q}_{t,T}(\cdot|s_t) = \widehat{\mathcal{H}}_t(\cdot, \tau)$. Figure 6 displays the skewness $\{\gamma_t\}$ and excess kurtosis $\{\eta_t\}$ associated with $\widehat{q}_{t,T}$ for maturity $\tau = 18$ days together with $\{\widehat{Z}_{1t}\}$ and $\{\widehat{Z}_{3t}\}$, motivating the investigation of their joint autocovariance structure.

The dynamic structure of the pairs $\{(\widehat{Z}_{1t}, \eta_t)\}$ and $\{(\widehat{Z}_{3t}, \gamma_t)\}$ for $\tau = 18$ is modeled by VAR(2) models. The choice of the VAR order is again based on AIC, SC, and HQ selection criteria. Portmanteau and LM tests on VAR residuals reject autocorrelations up to lag 12 and the roots of the characteristic polynomial lie inside of the unit circle.




Fig. 5 Estimated RN surface, $\widehat{\mathcal{H}}_t$ at t corresponding to 10 Jul. 2001

Table 4 Estimated parameters for the VAR(2) model on $\{\widehat{Z}_t\}$

	VAR(2)						
	Const	$\widehat{Z}_{1,t-1}$	$\widehat{Z}_{1,t-2}$	$\widehat{Z}_{2,t-1}$	$\widehat{Z}_{2,t-2}$	$\widehat{Z}_{3,t-1}$	$\widehat{Z}_{3,t-2}$
\widehat{Z}_{1t}	0.01	1.09	-0.16	0.10	-0.36	0.32	-0.23
\widehat{Z}_{2t}	0.01	-0.27	0.26	0.31	0.12	-1.14	0.33
\widehat{Z}_{3t}	0.01	-0.08	0.62	-0.05	-0.04	0.41	0.35

Table 5 Lag selection criteria for VAR models on $\{\widehat{Z}_t\}$. The asterisks denote the smallest value for each criterion

Order	AIC	SC	HQ
1	-11.03	-10.99	-11.01
2	-15.71	-15.54*	-15.64*
3	-15.77*	-15.46	-15.64
4	-15.76	-15.32	-15.58
5	-15.72	-15.16	-15.45

Modeling the dynamics of risk neutral densities using DSFM allows quantifying the mechanisms governing risk perceptions from agents acting in a market. Insights are obtained in two directions, concerning the autocovariance structure of $\{\widehat{Z}_t\}$, i.e., the time behavior of the RN surfaces and their cross-correlation with the skewness and excess kurtosis from the estimated risk neutral densities, i.e., the relation between the dynamics and shape of the obtained RN surfaces. As seen in Tables 6 and 7 the excess kurtosis and skewness from $\widehat{q}_{t,T}$ at maturity $\tau = 18$ are determined by the corresponding lagged values of \widehat{Z}_t .




Fig. 6 Left: RN excess kurtosis $\{\eta_t\}$, $\tau = 18$ (top), $\{\widehat{Z}_{1t}\}$ (bottom). Right: RN skewness $\{\gamma_t\}$, $\tau = 18$ (top), $\{\widehat{Z}_{2t}\}$ (bottom)

Table 6 Estimated parameters for the VAR(2) model on $\{\widehat{Z}_{1t}, \eta_t\}$

	VAR(2)				
	Const	$\widehat{Z}_{1,t-1}$	$\widehat{Z}_{1,t-2}$	η_{t-1}	η_{t-2}
\widehat{Z}_{1t}	0.04	0.86	0.08	0.01	0.00
η_t	-0.51	2.63	-1.75	0.67	0.19

Table 7 Estimated parameters for the VAR(2) model on $\{\widehat{Z}_{3t}, \gamma_t\}$

	VAR(2)				
	Const	$\widehat{Z}_{3,t-1}$	$\widehat{Z}_{3,t-2}$	γ_{t-1}	γ_{t-2}
\widehat{Z}_{3t}	0.00	0.20	0.27	0.01	-0.02
γ_t	0.00	-1.69	0.68	0.81	0.24

The presented methodology allows the investigation of the dynamics from risk neutral skewness and excess kurtosis based on statistical inference on $\{\widehat{Z}_t\}$. A natural further step is to perform econometric analysis on the cointegration between the lower dimensional time series and macroeconomic and financial indicators. This could provide deeper insights into the relation between risk assessments from investors acting in a market and the flow of economic information at which they are exposed.

Acknowledgements Financial support from the Deutsche Forschungsgemeinschaft via SFB 649 “Economic Risk” is gratefully acknowledged. The authors also thank the editor, an associate editor and two referees for their helpful comments.

Appendix: Assumptions

The results from Theorems 2.1 and 2.2, see Park et al. (2009), rely on the following assumptions:

- (A1) The variables X_{11}, \dots, X_{JT} , $\varepsilon_{11}, \dots, \varepsilon_{JT}$ and Z_1, \dots, Z_T are independent. The process Z_t is allowed to be nonrandom.
- (A2) For $t = 1, \dots, T$, the variables X_{1t}, \dots, X_{Jt} are identically distributed, have support $[0, 1]^d$ and a density f_t that is bounded from below and above on $[0, 1]^d$, uniformly over $t = 1, \dots, T$.

(A3) We assume that $E[\varepsilon_{jt}] = 0$ for $t = 1, \dots, T$ and $j = 1, \dots, J$, and

$$\sup_{t=1, \dots, T, j=1, \dots, J} E \exp[c\varepsilon_{jt}^2] < \infty$$

for $c > 0$ small enough.

(A4) The functions ψ_k may depend on the increasing indices T and J and are normed so that $\int_{[0,1]^d} \psi_k^2(x) dx = 1$ for $k = 1, \dots, K$. Furthermore, $\sup_{x \in [0,1]^d} \|\psi(x)\| = \mathcal{O}(K^{1/2})$.

(A5) The components m_0, \dots, m_L can be approximated by ψ_1, \dots, ψ_K , i.e.,

$$\delta_K = \sup_{x \in [0,1]^d} \inf_{\Gamma \in \mathcal{G}} |m(x) - \Gamma \psi(x)| \rightarrow 0 \tag{A.1}$$

for $l = 0, \dots, L$ and $K \rightarrow \infty$. We denote by Γ^* the matrix that fulfills

$$\sup_{x \in [0,1]^d} |m(x) - \Gamma^* \psi(x)| \leq 2\delta_K.$$

(A6) There exist constants $0 < C_L < C_U < \infty$ such that all eigenvalues of the random matrix $T^{-1} \sum_{t=1}^T Z_t Z_t^\top$ lie in the interval $[C_L, C_U]$ with probability tending to one.

(A7) The minimization (2.2) runs over all values of (Γ, z) with

$$\sup_{x \in [0,1]^d} \max_{1 \leq t \leq T} \|Z_t^\top \Gamma \psi(x)\| \leq M_T,$$

where M_T fulfills $\max_{1 \leq t \leq T} \|Z_t\| \leq M_T/C_m$ (with probability tending to one) for a constant $C_m > \sup_{x \in [0,1]^d} \|m(x)\|$.

(A8) It holds that

$$\xi^2 = (K + T)M_T^2 \log(JTM_T)(JT)^{-1} \rightarrow 0, \tag{A.2}$$

where the dimension L is fixed.

(A9) Z_t is a martingale difference with $E[Z_t | Z_1, \dots, Z_{t-1}] = 0$ and for some $C > 0$ $E[\|Z_t\|^2 | Z_1, \dots, Z_{t-1}] < C$ (a.s.). The matrix $E[Z_t Z_t^\top]$ has full rank. The process Z_t is independent of X_{11}, \dots, X_{TJ} and $\varepsilon_{11}, \dots, \varepsilon_{TJ}$.

(A10) The functions m_0, \dots, m_L are linearly independent. In particular, no function is equal to 0.

(A11) It holds that $(K^{1/2}M_T + T^{1/4})(\xi + \delta_K) = \mathcal{O}(1)$.

References

Ait-Sahalia, Y., Lo, A.: Nonparametric estimation of state-price densities implicit in financial asset prices. *J. Finance* **53**, 499–547 (1998)

Benko, M., Kneip, A., Härdle, W.: Common functional principal components. *Ann. Stat.* **37**(1), 1–34 (2009)

Breeden, D., Litzenberger, R.: Prices of state-contingent claims implicit in options prices. *J. Bus.* **51**, 621–651 (1978)

- Brüggemann, R., Härdle, W., Mungo, J., Trenkler, C.: VAR modeling for dynamic loadings driving volatility strings. *J. Financ. Econ.* **6**, 361–381 (2008)
- Cont, R., da Fonseca, J.: The dynamics of implied volatility surfaces. *Quant. Finance* **2**, 45–60 (2002)
- Dauxois, J., Pousse, A., Romain, Y.: Asymptotic theory for the principal component analysis of a vector random function: some applications to statistical inference. *J. Multivar. Anal.* **12**, 136–154 (1982)
- Fengler, M.: *Semiparametric Modeling of Implied Volatility*. Springer, Heidelberg (2005)
- Fengler, M., Härdle, W., Mammen, E.: A semiparametric factor model for implied volatility surface dynamics. *J. Financ. Econ.* **5**, 189–218 (2007)
- Gasser, T., Kneip, A.: Searching for structure in curve samples. *J. Am. Stat. Assoc.* **90**(432), 1179–1188 (1995)
- Hafner, R.: *Stochastic Implied Volatility*. Springer, Heidelberg (2004)
- Hall, P., Müller, H., Wang, J.: Properties of principal component methods for functional and longitudinal data analysis. *Ann. Stat.* **34**(3), 1493–1517 (2006)
- Hernández-Hernández, D., Schied, A.: A control approach to robust maximization with logarithmic utility and time-consistent penalties. *Stoch. Process. Appl.* **117**(8), 980–1000 (2007)
- Kneip, A.: Nonparametric estimation of common regressors for similar curve data. *Ann. Stat.* **22**(3), 1386–1427 (1994)
- Kneip, A., Gasser, T.: Statistical tools to analyse data representing a sample of curves. *Ann. Stat.* **20**(3), 1266–1305 (1992)
- Park, B., Mammen, E., Härdle, W., Borak, S.: Time series modelling with semiparametric factor dynamics. *J. Am. Stat. Assoc.* **104**(485), 284–298 (2009)
- Ramsay, J.O., Dalzell, C.T.: Some tools for functional data analysis. *J. R. Stat. Soc. B* **53**(3), 539–572 (1991)
- Rice, J., Silverman, B.W.: Estimating the mean and covariance structure nonparametrically when the data are curves. *J. R. Stat. Soc. B* **53**, 233–243 (1991)

Inhomogeneous Dependence Modeling With Time-Varying Copulae

Enzo GIACOMINI

Center for Applied Statistics and Economics, Humboldt-Universität zu Berlin, 10178 Berlin, Germany; and Weierstrass Institute for Applied Analysis and Stochastics, 10117 Berlin, Germany (*giacomini@wiwi.hu-berlin.de*)

Wolfgang HÄRDLE

AU1 Weierstrass Institute for Applied Analysis and Stochastics, 10117 Berlin, Germany

Vladimir SPOKOINY

Center for Applied Statistics and Economics, Humboldt-Universität zu Berlin, 10178 Berlin, Germany; and Weierstrass Institute for Applied Analysis and Stochastics, 10117 Berlin, Germany

Measuring dependence in multivariate time series is tantamount to modeling its dynamic structure in space and time. In risk management, the nonnormal behavior of most financial time series calls for non-Gaussian dependences. The correct modeling of non-Gaussian dependences is, therefore, a key issue in the analysis of multivariate time series. In this article we use copula functions with adaptively estimated time-varying parameters for modeling the distribution of returns. Furthermore, we apply copulae to the estimation of value-at-risk of portfolios and show their better performance over the *RiskMetrics* approach.

KEY WORDS: Adaptive estimation; Nonparametric estimation; Value-at-risk.

1. INTRODUCTION

Time series of financial data are high dimensional and typically have a non-Gaussian behavior. The standard modeling approach based on properties of the multivariate normal distribution therefore often fails to reproduce the stylized facts (i.e., fat tails, asymmetry) observed in returns from financial assets.

A correct understanding of the time-varying multivariate (conditional) distribution of returns is vital to many standard applications in finance such as portfolio selection, asset pricing, and value-at-risk (var) calculation. Empirical evidence from asymmetric return distributions have been reported in the recent literature. Longin and Solnik (2001) investigate the distribution of joint extremes from international equity returns and reject multivariate normality in their lower orthant; Ang and Chen (2002) test for conditional correlation asymmetries in U.S. equity data, rejecting multivariate normality at daily, weekly, and monthly frequencies; and Hu (2006) models the distribution of index returns with mixtures of copulae, finding asymmetries in the dependence structure across markets. For a concise survey on stylized empirical facts from financial returns see Cont (2001) and Granger (2003).

Modeling distributions with copulae has drawn attention from many researchers because it avoids the “procrustean bed” of normality assumptions, producing better fits of the empirical characteristics of financial returns. A natural extension is to apply copulae in a dynamic framework with conditional distributions modeled by copulae with time-varying parameters. The question, though, is how to steer the time-varying copulae parameters. This question is the focus of this article.

A possible approach is to estimate the parameter from structurally invariant periods. There is a broad field of econometric literature on structural breaks. Tests for unit root in macroeconomic series against stationarity with a structural

break at a known change point have been investigated by Perron (1989), and for an unknown change point by Zivot and Andrews (1992), Stock (1994) and Hansen (2001); Andrews (1993) tests for parameter instability in nonlinear models; Andrews and Ploberger (1994) construct asymptotic optimal tests for multiple structural breaks. In a different set up, Quintos, Fan, and Philips (2001) test for a constant tail index coefficient in Asian equity data against a break at an unknown point.

Time-varying copulae and structural breaks are combined in Patton (2006). The dependence structure across exchange rates is modeled with time-varying copulae with a parameter specified to evolve as an ARMA-type process. Tests for a structural break in the ARMA coefficients at a known change point have been performed, and strong evidence of a break was found. In a similar fashion, Rodriguez (2007) models the dependence across sets of Asian and Latin American stock indexes using time-varying copula where the parameter follows regime-switching dynamics. Common to these articles is that they use a fixed (parametric) structure for the pattern of changes in the copula parameter. **AU2**

In this article we follow a semiparametric approach, because we are not specifying the parameter changing scheme. Rather, we select locally the time-varying copula parameter. The choice is performed via an adaptive estimation under the assumption of local homogeneity: For every time point there exists an interval of time homogeneity in which the copula parameter can be well approximated by a constant. This interval is recovered from the data using local change point analysis. This does not imply that the model follows a change

point structure. The adaptive estimation also applies when the parameter varies smoothly from one value to another (see Spokoiny 2008).

[F1] Figure 1 shows the time-varying copula parameter determined by our procedure for a portfolio composed of daily prices of six German equities and the “global” copula parameter, shown by a constant horizontal line. The absence of parametric specification for time variations in the dependence structure (its dynamics is obtained adaptively from the data) allows for flexibility in estimating dependence shifts across time.

The obtained time-varying dependence structure can be used in financial engineering applications, the most prominent being the calculation of the var of a portfolio. Using copulae with adaptively estimated dependence parameters we estimate the var from DAX portfolios over time. As a benchmark procedure we choose *RiskMetrics*, a widely used methodology based on conditional normal distributions with a GARCH specification for the covariance matrix. Backtesting underlines the improved performance of the proposed adaptive time-varying copulae fitting.

[AU3]

This article is organized as follows: Section 2 presents the basic copulae definitions, Section 3 discusses the var and its estimation procedure. The adaptive copula estimation is exposed in Section 4 and is applied to simulated data in Section 5. In Section 6, the var from DAX portfolios is estimated based on adaptive time-varying copulae. The estimation performance is compared with the *RiskMetrics* approach by means of backtesting. Section 7 concludes.

2. COPULAE

Copulae merge marginally into joint distributions, providing a natural way for measuring the dependence structure between random variables. Copulae are present in the literature since Sklar (1959), although related concepts originate in Hoeffding (1940) and Fréchet (1951), and have been widely studied in the statistical literature (see Joe 1997, Nelsen 1998, and Mari and Kotz 2001). Applications of copulae in finance, insurance, and econometrics have been investigated in Embrechts, McNeil, and Straumann (2002); Embrechts, Hoeing, and Juri (2003a); Franke, Härdle, and Hafner (2004); and Patton (2004) among others. Cherubini, Luciano, and Vecchiato (2004) and McNeil, Frey, and Embrechts (2005) provide an overview of copulae for practical problems in finance and insurance.

Assuming absolutely continuous distributions and continuous marginals throughout this article, we have from Sklar’s

theorem that for a d -dimensional distribution function F with marginal cdf’s F_1, \dots, F_d there exists a unique copula $C : [0, 1]^d \rightarrow [0, 1]$ satisfying

$$F(x_1, \dots, x_d) = C\{F_1(x_1), \dots, F_d(x_d)\} \quad (2.1)$$

for every $\mathbf{x} = (x_1, \dots, x_d)^T \in \mathbb{R}^d$. Conversely, for a random vector $\mathbf{X} = (X_1, \dots, X_d)^T$ with cdf $F_{\mathbf{X}}$, the copula of \mathbf{X} may be written as $C_{\mathbf{X}}(u_1, \dots, u_d) = F_{\mathbf{X}}\{F_1^{-1}(u_1), \dots, F_d^{-1}(u_d)\}$, where $u_j = F_j(x_j)$, F_j is the cdf of X_j , and $F_j^{-1}(\alpha) = \inf\{x_j : F_j(x_j) \geq \alpha\}$ its generalized inverse, $j = 1, \dots, d$. A prominent copula is the Gaussian

$$C_{\Psi}^{Ga}(u_1, \dots, u_d) = F_Y\{\Phi^{-1}(u_1), \dots, \Phi^{-1}(u_d)\} \quad (2.2)$$

where $\Phi(s)$, $s \in \mathbb{R}$ stands for the one-dimensional standard normal cdf, F_Y is the cdf of $\mathbf{Y} = (Y_1, \dots, Y_d)^T \sim N_d(\mathbf{0}, \Psi)$, $\mathbf{0}$ is the $(d \times 1)$ vector of zeros, and Ψ is a correlation matrix. The Gaussian copula represents the dependence structure of the multivariate normal distribution. In contrast, the Clayton copula given by

$$C_{\theta}(u_1, \dots, u_d) = \left\{ \left(\sum_{j=1}^d u_j^{-\theta} \right) - d + 1 \right\}^{-\theta^{-1}} \quad (2.3)$$

for $\theta > 0$, expresses asymmetric dependence structures.

The dependence at upper and lower orthants of a copula C may be expressed by the upper and lower tail dependence coefficients $\lambda_U = \lim_{u \rightarrow 0} \hat{C}(u, \dots, u)/u$ and $\lambda_L = \lim_{u \rightarrow 0} C(u, \dots, u)/u$, where $u \in (0, 1]$ and \hat{C} is the survival copula of C (see Joe 1997 and Embrechts, Lindskog, and McNeil 2003b). Although Gaussian copulae are asymptotically independent at the tails ($\lambda_L = \lambda_U = 0$), the d -dimensional Clayton copulae exhibit lower tail dependence ($\lambda_L = d^{-1/\theta}$) but are asymptotically independent at the upper tail ($\lambda_U = 0$). Joe (1997) provides a summary of diverse copula families and detailed description of their properties.

For estimating the copula parameter, consider a sample $\{\mathbf{x}_t\}_{t=1}^T$ of realizations from \mathbf{X} where the copula of \mathbf{X} belongs to a parametric family $C = \{C_{\theta}, \theta \in \Theta\}$. Using Equation (2.1), the log-likelihood reads as $L(\theta; \mathbf{x}_1, \dots, \mathbf{x}_T) = \sum_{t=1}^T [\log c(F_1(x_{t,1}), \dots, F_d(x_{t,d}); \theta) + \sum_{j=1}^d \log f_j(x_{t,j})]$, where $c(u_1, \dots, u_d) = \partial^d C(u_1, \dots, u_d) / \partial u_1 \dots \partial u_d$ is the density of the copula C and f_j is the probability density function of F_j . The canonical maximum likelihood estimator $\hat{\theta}$ maximizes the pseudo log-likelihood with empirical marginal cdf’s $\hat{L}(\theta) = \sum_{t=1}^T \log c\{\hat{F}_1(x_{t,1}), \dots, \hat{F}_d(x_{t,d}); \theta$, where




Figure 1. Time-varying dependence. Time-varying dependence parameter and global parameter (horizontal line) estimated with Clayton copula, stock returns from Allianz, Münchener Rückversicherung, BASF, Bayer, DaimlerChrysler, and Volkswagen.

$$\widehat{F}_j(s) = \frac{1}{T+1} \sum_{k=1}^T 1_{\{x_{kj} \leq s\}} \quad (2.4)$$

for $j = 1, \dots, d$. Note that \widehat{F}_j differs from the usual empirical cdf by the denominator $T + 1$. This ensures that $\{\widehat{F}_1(x_{t,1}), \dots, \widehat{F}_d(x_{t,d})\}^\top \in (0, 1)^d$ and avoids infinite values the copula density may take on the boundary of the unit cube (see McNeil, Frey, and Embrechts 2005). Joe (1997); Cherubini, Luciano, and Vecchiato (2004); and Chen and Fan (2006) provide a detailed exposition of inference methods for copulae.

3. VALUE-AT-RISK AND COPULAE

The dependence (over time) between asset returns is especially important in risk management, because the profit and loss (P&L) function determines the var. More precisely, the var of a portfolio is determined by the multivariate distribution of risk factor increments. If $\mathbf{w} = (w_1, \dots, w_d)^\top \in \mathbb{R}^d$ denotes a portfolio of positions on d assets and $\mathbf{S}_t = (S_{t,1}, \dots, S_{t,d})^\top$ a nonnegative random vector representing the prices of the assets at time t , the value V_t of the portfolio \mathbf{w} is given by $V_t = \sum_{j=1}^d w_j S_{t,j}$. The random variable

$$L_t = (V_t - V_{t-1}), \quad (3.1)$$

called the profit and loss (P&L) function, expresses the change in the portfolio value between two subsequent time points. Defining the log-returns $\mathbf{X}_t = (X_{t,1}, \dots, X_{t,d})^\top$, where $X_{t,j} = \log S_{t,j} - \log S_{t-1,j}$ and $\log S_{0,j} = 0, j = 1, \dots, d$, Equation (3.1) can be written as

$$L_t = \sum_{j=1}^d w_j S_{t-1,j} \{\exp(X_{t,j}) - 1\}. \quad (3.2)$$

The cdf of L_t is given by $F_{t,L_t}(x) = P_t(L_t \leq x)$. The var at level α from a portfolio \mathbf{w} is defined as the α quantile from F_{t,L_t} :

$$\text{var}_t(\alpha) = F_{t,L_t}^{-1}(\alpha). \quad (3.3)$$

It follows from Equation (3.2) that F_{t,L_t} depends on the specification of the d -dimensional distribution of the risk factors \mathbf{X}_t . Thus, modeling their distribution over time is essential for obtaining the quantiles (Eq. 3.3).

The *RiskMetrics* technique, a widely used methodology for var estimation, assumes that risk factors \mathbf{X}_t follow a conditional multivariate normal distribution $\mathcal{L}(\mathbf{X}_t | \mathcal{F}_{t-1}) = N(\mathbf{0}, \boldsymbol{\Sigma}_t)$, where $\mathcal{F}_{t-1} = \sigma(\mathbf{X}_1, \dots, \mathbf{X}_{t-1})$ is the σ field generated by the first $t - 1$ observations, and estimates the covariance matrix $\boldsymbol{\Sigma}_t$ for one period return as

$$\widehat{\boldsymbol{\Sigma}}_t = \lambda \widehat{\boldsymbol{\Sigma}}_{t-1} + (1 - \lambda) \mathbf{X}_{t-1} \mathbf{X}_{t-1}^\top, \quad (3.4)$$




Figure 2. Local change point procedure. Choice of intervals I_k and I_{k-1} .

where the parameter λ is the so-called decay factor. $\lambda = 0.94$ provides the best backtesting results for daily returns according to Morgan (1996). Using the copulae-based approach, one first corrects the contemporaneous mean and volatility in the log-returns process:

$$X_{t,j} = \mu_{t,j} + \sigma_{t,j} \varepsilon_{t,j}, \quad (3.5)$$

where $\mu_{t,j} = E[X_{t,j} | \mathcal{F}_{t-1}]$ is the conditional mean and $\sigma_{t,j}^2 = E[(X_{t,j} - \mu_{t,j})^2 | \mathcal{F}_{t-1}]$ is the conditional variance of $X_{t,j}$. The standardized innovations $\boldsymbol{\varepsilon}_t = (\varepsilon_{t,1}, \dots, \varepsilon_{t,d})^\top$ have joint cdf $F_{\boldsymbol{\varepsilon}_t}$ given by

$$F_{\boldsymbol{\varepsilon}_t}(x_1, \dots, x_d) = C_\theta \{F_{t,1}(x_1), \dots, F_{t,d}(x_d)\}, \quad (3.6)$$

where $F_{t,j}$ is the cdf of $\varepsilon_{t,j}$ and C_θ is a copula belonging to a parametric family $C = C_\theta, \theta \in \Theta$. For details on the previous model specification, see Chen and Fan (2006) and Chen, Fan, and Tsyrennikov (2006). For the Gaussian copula with Gaussian marginals, we recover the conditional Gaussian *RiskMetrics* framework.

To obtain the var in this setup, the dependence parameter and cdf's from residuals are estimated from a sample of log-returns and are used to generate P&L Monte Carlo samples. Their quantiles at different levels are the estimators for the var (see Embrechts, McNeil, and Straumann 2002).

The whole procedure can be summarized as follows (see Härdle, Kleinow, and Stahl 2002; and Giacomini and Härdle 2005): For a portfolio $\mathbf{w} \in \mathbb{R}^d$ and a sample $\{x_{t,j}\}_{t=1}^T, j = 1, \dots, d$ of log-returns, the var at level α is estimated according to the following steps:

1. Determination of innovations $\{\hat{\varepsilon}_t\}_{t=1}^T$ by, for example, “deGARCHing”
2. Specification and estimation of marginal cdf's $F_j(\hat{\varepsilon}_j)$
3. Specification of a parametric copula family C and estimation of the dependence parameter θ
4. Generation of Monte Carlo sample of innovations ε and losses L
5. Estimation of $\widehat{\text{var}}(\alpha)$, the empirical α quantile of F_L

4. MODELING WITH TIME-VARYING COPULAE

Similar to the *RiskMetrics* procedure, one can perform a moving (fixed-length) window estimation of the copula parameter. This procedure, though, does not fine-tune local changes in dependencies. In fact, the cdf $F_{\boldsymbol{\varepsilon}_t}$ from Equation (3.6) is modeled as $F_{t,\boldsymbol{\varepsilon}_t} = C_\theta \{F_{t,1}(\cdot), \dots, F_{t,d}(\cdot)\}$ with probability measure P_θ . The moving window of fixed width will estimate a θ_t for each t , but it has clear limitations. The choice of a small window results in a high pass filtering and, hence, in a very unstable estimate with huge variability. The choice of a large window leads to a poor sensitivity of the estimation procedure




Figure 3. Homogeneity test. Testing interval I , tested interval I , and subintervals J and J^c for a point $\tau \in I$.

and to a high delay in the reaction to changes in dependence measured by the parameter θ_r .

To choose an interval of homogeneity, we use a local parametric fitting approach as introduced by Polzehl and Spokoiny (2006), Belomestny and Spokoiny (2007) and Spokoiny (2008). The basic idea is to select for each time point t_0 an interval $I_{t_0} = [t_0 - m_{t_0}, t_0]$ of length m_{t_0} in such a way that the time-varying copula parameter θ_t can be well approximated by a constant value θ . The question is, of course, how to select m_{t_0} in an online situation from historical data. The aim should be to select I_{t_0} as close as possible to the so-called ‘‘oracle’’ choice interval. The oracle choice is defined as the largest interval $I = [t_0 - m_{t_0}^*, t_0]$, for which the small modeling bias condition

$$\Delta_I(\theta) = \sum_{t \in I} \mathcal{K}(P_{\theta_t}, P_{\theta}) \leq \Delta \quad (4.1)$$

for some $\Delta \geq 0$ holds. Here, θ is constant and $\mathcal{K}(P_{\vartheta}, P_{\vartheta'}) = E_{\vartheta} \log\{p(y, \vartheta)/p(y, \vartheta')\}$ denotes the Kullback-Leibler divergence. In such an oracle choice interval, the parameter $\theta_{t_0} = \theta_t|_{t=t_0}$ can be ‘‘optimally’’ estimated from $I = [t_0 - m_{t_0}^*, t_0]$. The error and risk bounds are calculated in Spokoiny (2008). It is important to mention that the concept of local parametric approximation allows one to treat in a unified way the case of ‘‘switching regime’’ models with spontaneous changes of parameters and the ‘‘smooth transition’’ case when the parameter varies smoothly in time.

The oracle choice of the interval of homogeneity depends on the unknown time-varying copula parameter θ_r . The next section presents an adaptive (data-driven) procedure that mimics the oracle in the sense that it delivers the same accuracy of estimation as the oracle one. The trick is to find the largest interval in which the hypothesis of a local constant copula

parameter is supported. The local change point (LCP) detection procedure originates from Mercurio and Spokoiny (2004) and sequentially tests the hypothesis: θ_t is constant (i.e., $\theta_t = \theta$) within some interval I (local parametric assumption).

The LCP procedure for a given point t_0 starts with a family of nested intervals $I_0 \subset I_1 \subset I_2 \subset \dots \subset I_K = I_{K+1}$ of the form $I_k = [t_0 - m_k, t_0]$. The sequence m_k determines the length of these interval ‘‘candidates’’ (see Section 4.2). Every interval I_k leads to an estimate $\hat{\theta}_k$ of the copula parameter θ_{t_0} . The procedure selects one interval \hat{I} out of the given family and, therefore, the corresponding estimate $\hat{\theta} = \hat{\theta}_{\hat{I}}$.

The idea of the procedure is to screen each interval $\mathcal{J}_k = [t_0 - m_k, t_0 - m_{k-1}]$ sequentially and check each point $\tau \in \mathcal{J}_k$ as a possible change point location (see Section 4.1 for more details). The family of intervals I_k and \mathcal{J}_k are illustrated in Figure 2. The interval I_k is accepted if no change point is detected within $\mathcal{J}_1, \dots, \mathcal{J}_k$. If the hypothesis of homogeneity is rejected for an interval candidate I_k , the procedure stops and selects the latest accepted interval. The formal description reads as follows:

Start the procedure with $k = 1$ and test the hypothesis $H_{0,k}$ of no structural changes within \mathcal{J}_k using the larger testing interval I_{k+1} . If no change points were found in \mathcal{J}_k , then I_k is accepted. Take the next interval \mathcal{J}_{k+1} and repeat the previous step until homogeneity is rejected or the largest possible interval $I_K = [t_0 - m_K, t_0]$ is accepted. If $H_{0,k}$ is rejected for \mathcal{J}_k , the estimated interval of homogeneity is the last accepted interval $\hat{I} = I_{k-1}$. If the largest possible interval I_K is accepted, we take $\hat{I} = I_K$. We estimate the copula dependence parameter θ at time instant t_0 from observations in \hat{I} , assuming the homogeneous model within \hat{I} (i.e., we define $\hat{\theta}_{t_0} = \hat{\theta}_{\hat{I}}$). We also denote by \hat{I}_k the largest accepted interval after k steps of

Table 1. Critical values $\mathfrak{z}_k(\rho, \theta^*)$

k	$\theta^* = 0.5$			$\theta^* = 1.0$			$\theta^* = 1.5$		
	$\rho = 0.2$	$\rho = 0.5$	$\rho = 1.0$	$\rho = 0.2$	$\rho = 0.5$	$\rho = 1.0$	$\rho = 0.2$	$\rho = 0.5$	$\rho = 1.0$
1	3.64	3.29	2.88	3.69	3.29	2.84	3.95	3.49	2.96
2	3.61	3.14	2.56	3.43	2.91	2.35	3.69	3.02	2.78
3	3.31	2.86	2.29	3.32	2.76	2.21	3.34	2.80	2.09
4	3.19	2.69	2.07	3.04	2.57	1.80	3.14	2.55	1.86
5	3.05	2.53	1.89	2.92	2.22	1.53	2.95	2.65	1.49
6	2.87	2.26	1.48	2.92	2.17	1.19	2.83	2.04	0.94
7	2.51	1.88	1.02	2.64	1.82	0.56	2.62	1.79	0.31
8	2.49	1.72	0.35	2.33	1.39	0.00	2.35	1.33	0.00
9	2.18	1.23	0.00	2.03	0.81	0.00	2.10	0.60	0.00
10	0.92	0.00	0.00	0.82	0.00	0.00	0.79	0.00	0.00

NOTE: Critical values are obtained according to Equation (4.2), based on 5,000 simulations. Clayton copula, $m_0 = 20$ and $c = 1.25$.




Figure 4. LCP and sudden jump in copula parameter. Pointwise median (full), and 0.25 and 0.75 quantiles (dotted) from $\hat{\theta}_t$. True parameter θ_t (dashed) with $\vartheta_a = 0.10$, $\vartheta_b = 0.50, 0.75$, and 1.00 (left, top to bottom); and $\vartheta_b = 0.10$, $\vartheta_a = 0.50, 0.75$, and 1.00 (right, top to bottom). Based on 100 simulations from Clayton copula, estimated with LCP, $m_0 = 20$, $c = 1.25$, and $\rho = 0.5$.

the algorithm and, by $\hat{\theta}_k$ the corresponding estimate of the copula parameter.

It is worth mentioning that the objective of the described estimation algorithm is not to detect the points of change for the copula parameter, but rather to determine the current dependence structure from historical data by selecting an interval of time homogeneity. This distinguishes our approach from other procedures for estimating a time-varying parameter by change point detection. A visible advantage of our approach is that it equally applies to the case of spontaneous changes in the dependence structure and in the case of smooth transition in the copula parameter. The obtained dependence structure can be used for different purposes in financial engineering, the most prominent being the calculation of the var (see also Section 6).

The theoretical results from Spokoiny and Chen (2007) and Spokoiny (2008) indicate that the proposed procedure provides the rate optimal estimation of the underlying parameter when this varies smoothly with time. It has also been shown that the procedure is very sensitive to structural breaks and provides the minimal possible delay in detection of changes, where the delay depends on the size of change in terms of Kullback-Leibler divergence.

4.1 Test of Homogeneity Against a Change Point Alternative

In the homogeneity test against a change point alternative we want to check every point of an interval I (recall Fig. 2), here called the “tested interval,” on a possible change in the dependence structure at this moment. To perform this check, we assume a larger testing interval J of form $J = [t_0 - m, t_0]$, so that I is an internal subset within J . The null hypothesis H_0 means that $\forall t \in I, \theta_t = \theta$ (i.e., the observations in I follow the

model with dependence parameter θ). The alternative hypothesis H_1 claims that $\exists \tau \in I$ such that $\theta_t = \theta_1$ for $t \in J = [\tau, t_0]$ and $\theta_t = \theta_2 \neq \theta_1$ for $t \in J^c = [t_0 - m, \tau]$ (i.e., the parameter θ changes spontaneously in some point $\tau \in I$). Figure 3 depicts I, J , and the subintervals J and J^c determined by the point $\tau \in I$.

Let $L_J(\theta)$ be the log-likelihood and $\tilde{\theta}_J$ the maximum likelihood estimate for the interval J . The log-likelihood functions corresponding to H_0 and H_1 are $L_J(\theta)$ and $L_J(\theta_1) + L_{J^c}(\theta_2)$, respectively. The likelihood ratio test for the single change point with known fixed location τ can be written as

Table 2. Detection delay statistics

$(\vartheta_a, \vartheta_b)$	r	Mean	SD	Max	Min
(0.50, 0.10)	0.25	9.06	7.28	56	0
	0.50	13.64	9.80	60	0
	0.75	21.87	14.52	89	3
(0.75, 0.10)	0.25	5.16	4.24	21	0
	0.50	8.85	5.55	25	0
	0.75	16.72	10.37	64	3
(1.00, 0.10)	0.25	4.47	2.94	12	0
	0.50	7.94	4.28	22	0
	0.75	14.79	7.38	62	5
(0.10, 0.50)	0.25	8.94	6.65	36	0
	0.50	14.21	9.06	53	0
	0.75	21.43	12.15	68	0
(0.10, 0.75)	0.25	9.00	4.80	25	0
	0.50	14.30	5.96	40	3
	0.75	21.00	10.97	75	6
(0.10, 1.00)	0.25	7.39	3.67	19	0
	0.50	13.10	4.13	22	2
	0.75	20.13	7.34	55	10

NOTE: The detection delays δ are calculated as in Equation (5.1), with the statistics based on 100 simulations. Clayton copula, $m_0 = 20$, $c = 1.25$, and $\rho = .5$. SD, standard deviation.




Figure 5. Divergences for upward and downward jumps. Kullback-Leibler divergences $\mathcal{K}(0.10, \vartheta)$ (full) and $\mathcal{K}(\vartheta, 0.10)$ (dashed) for Clayton copula.

$$T_{I,\tau} = \max_{\theta_1, \theta_2} \{L_J(\theta_1) + L_{J^c}(\theta_2)\} - \max_{\theta} L_I(\theta) \\ = L_J(\tilde{\theta}_J) + L_{J^c}(\tilde{\theta}_{J^c}) - L_I(\tilde{\theta}_I).$$

The test statistic for an unknown change point location is defined as $T_I = \max_{\tau \in I} T_{I,\tau}$. The change point test compares this test statistic with a critical value λ_I , which may depend on the interval I . One rejects the hypothesis of homogeneity if $T_I > \lambda_I$.

4.2 Parameters of the LCP Procedure

To apply the LCP testing procedure for local homogeneity, we have to specify some parameters. This includes selecting interval candidates I_k or, equivalently, of the tested intervals J_k and choosing respective critical values λ_k . One possible parameter set that has been used successfully in simulations is presented in the following section.

4.2.1 Selection of interval candidates J_k and internal points I_k . It is useful to take the set of numbers m_k defining the length of I_k and J_k in the form of a geometric grid. We fix the

value m_0 and define $m_k = \lfloor m_0 c^k \rfloor$ for $k = 1, 2, \dots, K$ and $c > 1$ where $\lfloor x \rfloor$ means the integer part of x . We set $I_k = [t_0 - m_k, t_0]$ and $J_k = [t_0 - m_k, t_0 - m_{k-1}]$ for $k = 1, 2, \dots, K$ (see Fig. 2).

4.2.2 Choice of the critical values λ_k . The algorithm is in fact a multiple testing procedure. Mercurio and Spokoiny (2004) suggested selecting the critical value z_k to provide the overall first type error probability of rejecting the hypothesis of homogeneity in the homogeneous situation. Here we follow another proposal from Spokoiny and Chen (2007), which focuses on estimation losses caused by the “false alarm”—in our case obtaining a homogeneity interval that is too small—rather than on its probability.

In the homogeneous situation with $\theta_t \equiv \theta^*$ for all $t \in I_{k+1}$, the desirable behavior of the procedure is that after the first k steps the selected interval \hat{I}_k coincides with I_k and the corresponding estimate $\hat{\theta}_k$ coincides with $\tilde{\theta}_k$, which means there is no false alarm. On the contrary, in the case of a false alarm, the selected interval \hat{I}_k is smaller than I_k and, hence, the corresponding estimate $\hat{\theta}_k$ has larger variability than $\tilde{\theta}_k$. This means that the false alarm during the early steps of the procedure is more critical than during the final steps, because it may lead to selecting an estimate with very high variance. The difference between $\hat{\theta}_k$ and $\tilde{\theta}_k$ can naturally be measured by the value $L_{I_k}(\tilde{\theta}_k, \hat{\theta}_k) = L_{I_k}(\tilde{\theta}_k) - L_{I_k}(\hat{\theta}_k)$ normalized by the risk of the nonadaptive estimate $\tilde{\theta}_k$, $\mathfrak{R}(\theta^*) = \max_{k \geq 1} E_{\theta^*} |L_{I_k}(\tilde{\theta}_k, \theta^*)|^{1/2}$. The conditions we impose read as

$$E_{\theta^*} |L_{I_k}(\tilde{\theta}_k, \hat{\theta}_k)|^{1/2} \leq \rho \mathfrak{R}(\theta^*), \quad k = 1, \dots, K, \quad \theta^* \in \Theta. \quad (4.2)$$

The critical values λ_k are selected as minimal values providing these constraints. In total we have K conditions to select K critical values $\lambda_1, \dots, \lambda_K$. The values λ_k can be selected sequentially by Monte Carlo simulation, where one simulates under $H_0 : \theta_t = \theta^*, \forall t \in I_k$. The parameter ρ defines how conservative the procedure is. A small ρ value leads to larger critical values and hence to a conservative and nonsensitive procedure, whereas an increase in ρ results in more sensitivity at cost of stability. For details, see Spokoiny and Chen (2007) or Spokoiny (2008).




Figure 6. Mean detection delay and parameter jumps. Mean detection delays (dots) at rule $r = 0.75, 0.50$, and 0.25 from top to bottom. Left: $\vartheta_b = 0.10$ (upward jump). Right: $\vartheta_a = 0.10$ (downward jump), based on 100 simulations from Clayton copula, $m_0 = 20$, $c = 1.25$, and $\rho = 0.5$.




Figure 7. LCP and smooth change in copula parameter. Pointwise median (full), 0.25 and 0.75 quantiles (dotted) from $\hat{\theta}_t$ and true parameter θ_t (dashed) with $\vartheta_a = 0.10$ and $\vartheta_b = 1.00$ (left), and $\vartheta_a = 1.00$ and $\vartheta_b = 0.10$ (right). Based on 100 simulations from Clayton copula, estimated with LCP, $m_0 = 20$, $c = 1.25$, and $\rho = 0.5$.

5. SIMULATED EXAMPLES

In this section we apply the LCP procedure on simulated data with a dependence structure given by the Clayton copula. We generate sets of six-dimensional data with a sudden jump in the dependence parameter given by

$$\theta_t = \begin{cases} \vartheta_a & \text{if } -390 \leq t \leq 10 \\ \vartheta_b & \text{if } 10 < t \leq 210 \end{cases}$$

for different values of $(\vartheta_a, \vartheta_b)$: One of them is fixed at .1 (close to independence) and the other is set to larger values.

The LCP procedure is implemented with the family of interval candidates in form of a geometric grid defined by $m_0 = 20$ and $c = 1.25$. The critical values, selected according to Equation (4.2) for different ρ and θ^* , are displayed in Table 1. The choice of θ^* has negligible influence in the critical values for fixed ρ , therefore we use $\mathfrak{z}_1, \dots, \mathfrak{z}_K$ obtained with $\theta^* = 1.0$. Based on our experience, see Spokoiny and Chen (2007) and Spokoiny (2008), the default choice for ρ is 0.5.

Figure 4 shows the pointwise median and quantiles of the estimated parameter $\hat{\theta}_t$ for distinct values of $(\vartheta_a, \vartheta_b)$ based on 100 simulations. The detection delay δ at rule $r \in [0, 1]$ to jump of size $\gamma = \theta_t - \theta_{t-1}$ at t is expressed by

$$\delta(t, \gamma, r) = \min\{u \geq t : \hat{\theta}_u = \theta_{t-1} + r\gamma\} - t \quad (5.1)$$

and represents the number of steps necessary for the estimated parameter to reach the r fraction of a jump in the true parameter.

Detection delays are proportional to the probability of error of type II (i.e., the probability of accepting homogeneity in case of a jump). Thus, tests with higher power correspond to lower delays δ . Moreover, because the Kullback-Leibler divergences for upward and downward jumps are proportional to the power of the respective homogeneity tests, larger divergences result in faster jump detections.

The descriptive statistics for detection delays to jumps at $t = 11$ for different values of $(\vartheta_a, \vartheta_b)$ are in Table 2. The mean detection delay decreases with $\gamma = \vartheta_b - \vartheta_a$ and are higher for downward jumps than for upward jumps. Figure 5 shows that for Clayton copulae the Kullback-Leibler divergence is higher for upward jumps than for downward jumps. Figure 6 displays the mean detection delays against jump size for upward and downward jumps.

The LCP procedure is also applied on simulated data with smooth transition in the dependence parameter given by

$$\theta_t = \begin{cases} \vartheta_a & \text{if } -350 \leq t \leq 50 \\ \vartheta_a + \frac{t-50}{100}(\vartheta_b - \vartheta_a) & \text{if } 50 < t \leq 150 \\ \vartheta_b & \text{if } 150 < t \leq 350. \end{cases}$$

Figure 7 depicts the pointwise median and quantiles of the estimated parameter $\hat{\theta}_t$ and the true parameter θ_t for $(\vartheta_a, \vartheta_b)$ set to $(0.10, 1.00)$ and $(1.00, 0.10)$.

6. EMPIRICAL RESULTS

In this section the var from German stock portfolios is estimated based on time-varying copulae and *RiskMetrics* approaches. The time-varying copula parameters are selected by local change point (LCP) and moving window procedures. Backtesting is used to evaluate the performances of the three methods in var estimation.

Two groups of six stocks listed on DAX are used to compose the portfolios. Stocks from group 1 belong to three different industries: automotive (Volkswagen and DaimlerChrysler), insurance (Allianz and Münchener Rückversicherung), and chemical (Bayer and BASF). Group 2 is composed of stocks from six industries: electrical (Siemens), energy (E.ON), metallurgical (ThyssenKrupp), airlines (Lufthansa), pharmaceutical (Schering), and chemical (Henkel). The portfolio values are calculated using 1,270 observations, from January 1, 2000 to December 31, 2004, of the daily stock prices (data available at <http://sfb649.wiwi.hu-berlin.de/fedc>).

The selected copula belongs to the Clayton family (Eq. 2.3). Clayton copulae have a natural interpretation and are well advocated in risk management applications. In line with the stylized facts for financial returns, Clayton copulae are asymmetric and present lower tail dependence, modeling joint

Table 3. p Values from tests on residuals $\hat{\varepsilon}_{t,j}$

j	Ljung-Box		ARCH	
	Group 1	Group 2	Group 1	Group 2
1	0.33	0.52	0.15	0.04
2	0.13	0.35	0.15	0.98
3	0.21	0.08	0.34	0.72
4	0.99	0.05	0.10	0.18
5	0.90	0.07	0.91	0.77
6	0.28	0.81	0.28	0.94




Figure 8. Time-varying dependence, group 1. Copula parameter $\hat{\theta}_t$ estimated with LCP method, Clayton copula, $m_0 = 20$, $c = 1.25$, and $\rho = 0.5$.

extreme events at lower orthants with higher probability than Gaussian copulae for the same correlation, see McNeil, Frey, and Embrechts (2005). This fact is essential for var calculations and is illustrated by the ratio between Equations (2.2) and (2.3) for off-diagonal elements of Ψ set to 0.25 and $\theta = 0.5$. For the quantiles $u_i = 0.05$, $i = 1, \dots, 6$ the ratio $C_{\Psi}^{Ga}(u_1, \dots, u_6) / C_{\theta}(u_1, \dots, u_6)$ equals 2.3×10^{-2} , whereas for the 0.01 quantiles it equals 1.3×10^{-3} .

The var estimation follows the steps described in Section 3. Using the *RiskMetrics* approach, the log-returns X_t are assumed conditionally normal distributed with zero mean and covariance matrix following a GARCH specification with fixed decay factor $\lambda = 0.94$ as in Equation (3.4).

In the time-varying copulae estimation, the log-returns are modeled as in Equation (3.5), where the innovations ε_t have cdf $F_{t,\varepsilon_i}(x_1, \dots, x_d) = C_{\theta_i}\{F_{t,1}(x_1), \dots, F_{t,d}(x_d)\}$ and C_{θ} is the Clayton copula. The univariate log-returns $X_{t,j}$ corresponding to stock j are devolatilized according to *RiskMetrics* (i.e., with zero conditional means and conditional variances $\sigma_{t,j}^2$ estimated by the univariate version of Equation (3.4) with a decay factor equal to 0.94). We note that this choice sets the same specification for the dynamics of the univariate returns across all methods (*RiskMetrics*, moving windows, and LCP), making their performances in var estimation comparable. Moreover, as the means from daily returns are clearly dominated by the variances and are approximately independent on the available information sets (see Jorion 1995; Fleming, Kirby, and Ostdiek 2001; and Christoffersen and Diebold 2006), their specification is very unlikely to cause a perceptible bias in the estimated variances and dependence parameters. Therefore, the zero mean assumption is, as pointed out by Kim, Malz, and Mina (1999), as good as any other choice. Daily returns are also modeled with zero conditional means in Fan and Gu (2003) and Härdle, Herwartz, and Spokoiny (2003) among others.

The GARCH specification (Eq. 3.4) with $\lambda = .94$ optimizes variance forecasts across a large number of assets (Morgan 1996), and is widely used in the financial industry. Different choices for the decay factor (like 0.85 or 0.98) result in negligible changes (about 3%) in the estimated dependence parameter.

The p values from the Ljung-Box test for serial correlation and from ARCH test for heteroscedasticity effects in the obtained residuals $\hat{\varepsilon}_{t,j}$ are in Table 3. Normality is rejected by Jarque-Bera test, with p values approximately 0.00 for all residuals in both groups. The empirical cdf's of residuals as defined in Equation (2.4) are used for the copula estimation.

With the moving windows approach, the size of the estimating window is fixed as 250 days corresponding to 1 business year (the same size is used in, for example, Fan and Gu (2003)). For the LCP procedure, following Section 4.2, we set the family of interval candidates as a geometric grid with $m_0 = 20$, $c = 1.25$, and $\rho = 0.5$. We have chosen these parameters from our experience in simulations (for details on robustness of the reported results with respect to the choice of m_0 and c , refer to Spokoiny (2008)).

The performance of the var estimation is evaluated based on backtesting. At each time t , the estimated var at level α for a portfolio w is compared with the realization l_t of the corresponding P&L function (see Eq. 3.2), with an exceedance occurring for each l_t less than $\widehat{\text{var}}_t(\alpha)$. The ratio of the number of exceedances to the number of observations gives the exceedance ratio

$$\hat{\alpha}_w(\alpha) = \frac{1}{T} \sum_{t=1}^T \mathbf{1}_{\{l_t < \widehat{\text{var}}_t(\alpha)\}}$$

Because the first 250 observations are used for estimation, $T = 1,020$. The difference between $\hat{\alpha}$ and the desired level α is expressed by the relative exceedance error




Figure 9. Time-varying dependence, group 2. Copula parameter $\hat{\theta}_t$ estimated with LCP method, Clayton copula, $m_0 = 20$, $c = 1.25$, and $\rho = 0.5$.




Figure 10. Estimated var across methods, group 1. P&L realizations l_t (dots), $\widehat{\text{var}}_t(\alpha)$ (line), and exceedance times (crosses). Estimated with LCP (top), moving windows (middle), and *RiskMetrics* (bottom) for equally weighted portfolio w^* at level $\alpha = 0.05$.

$$e_w = (\hat{\alpha} - \alpha)/\alpha.$$

We compute exceedance ratios and relative exceedance errors to levels $\alpha = 0.05$ and 0.01 for a set $W = \{w^*, w_n; n = 1, \dots, 100\}$ of portfolios, where each $w_n = (w_{n,1}, \dots, w_{n,6})^T$ is a realization of a random vector uniformly distributed on $S = \{(x_1, \dots, x_6) \in \mathbb{R}^6 : \sum_{i=1}^6 x_i = 1, x_i \geq .1\}$, and $w^* = 1/6 \mathbf{I}_6$, with \mathbf{I}_d denoting the $(d \times 1)$ vector of ones, is the equally weighted portfolio. The degree of diversification of a portfolio can be measured based on the majorization preordering on S (see Marshall and Olkin 1979). In other words, a portfolio w_a is more diversified than portfolio w_b if $w_a \prec w_b$. Under the majorization preordering the vector w^* satisfies $w^* \preceq w$ for all $w \in S$; therefore, the equally weighted portfolio is the most diversified portfolio from W , see Ibragimov and Walden (2007).

The average relative exceedance error over portfolios and the corresponding standard deviation

$$A_W = \frac{1}{|W|} \sum_{w \in W} e_w$$

$$D_W = \left\{ \frac{1}{|W|} \sum_{w \in W} (e_w - A_W)^2 \right\}^{\frac{1}{2}}$$

are used to evaluate the performances of the time-varying copulae and *RiskMetrics* methods in var estimation.

The dependence parameter estimated with LCP for stocks from groups 1 and 2 are shown in Figures 8 and 9. The different industry concentrations in each group are reflected in the higher parameter values obtained for group 1. The P&L and the var at level 0.05 estimated with LCP, moving windows, and




Figure 11. Estimated var across methods, group 2. P&L realizations l_t (dots), $\widehat{\text{var}}_t(\alpha)$ (line), and exceedance times (crosses). Estimated with LCP (top), moving windows (middle), and *RiskMetrics* (bottom) for equally weighted portfolio w^* at level $\alpha = 0.05$.

Table 4. Exceedance ratios and errors, group 1

	RiskMetrics		Moving windows		LCP	
	$\alpha = 5.00$	$\alpha = 1.00$	$\alpha = 5.00$	$\alpha = 1.00$	$\alpha = 5.00$	$\alpha = 1.00$
	$\hat{\alpha}_{w^*}$	6.11	1.48	5.62	0.59	5.52
$\hat{\alpha}_{w_1}$	5.91	1.38	5.42	0.49	5.42	0.69
$\hat{\alpha}_{w_2}$	6.40	1.28	5.91	0.49	5.71	0.59
A_W	0.23	0.45	0.11	-0.49	0.11	-0.36
D_W	0.04	0.14	0.06	0.08	0.06	0.10

NOTE: Exceedance ratios for portfolios w^* , w_1 , and w_2 , and average and standard deviation from relative exceedance errors. Across levels and methods, ratios and levels are expressed as a percentage.

RiskMetrics methods for the equally weighted portfolio w^* are in Figures 10 (group 1) and 11 (group 2). Exceedance ratios for portfolios w^* , w_1 , and w_2 ; average relative exceedance errors; and corresponding standard deviations across methods and levels are shown in Tables 4 (group 1) and 5 (group 2).

Based on the exceedance errors, the LCP procedure outperforms the moving windows (second best) and *RiskMetrics* methods in var estimation in group 1. At level 0.05, the average error associated with copula methods is about half the error from *RiskMetrics* estimation for nearly the same standard deviation. At level 0.01, the LCP average error is the smallest in absolute value, and copula methods present less standard deviations. At this level, copula methods overestimate var, and *RiskMetrics* underestimates it. Although overestimation of var means that a financial institution would be requested to keep more capital aside than necessary to guarantee the desired confidence level, underestimation means that less capital is reserved and the desired level is not guaranteed. Therefore, from the regulatory point of view, overestimation is preferred to underestimation. In the less concentrated group 2, LCP outperforms moving windows and *RiskMetrics* at the level 0.05, presenting the smallest average error in magnitude for nearly the same value of D_W . At level 0.01, copula methods overestimate and *RiskMetrics* underestimates the var by about 60%.

It is interesting to note the effect of portfolio diversification on the exceedance errors for group 1 and level 0.01. The errors decrease with increasing portfolio diversification for copulae methods but become larger under the *RiskMetrics* estimation. For other groups and levels, the diversification effects are not clear. Refer to Ibragimov (2007) and Ibragimov and Walden

(2007) for details on the effects of portfolio diversification under heavy-tailed distributions in risk management.

7. CONCLUSION

In this article we modeled the dependence structure from German equity returns using time-varying copulae with adaptively estimated parameters. In contrast to Patton (2006) and Rodriguez (2007), we neither specified the dynamics nor assumed regime switching models for the copula parameter. The parameter choice was performed under the local homogeneity assumption with homogeneity intervals recovered from the data through local change point analysis.

We used time-varying Clayton copulae, which are asymmetric and present lower tail dependence, to estimate the var from portfolios of two groups of German securities, presenting different levels of industry concentration. *RiskMetrics*, a widely used methodology based on multivariate normal distributions, was chosen as a benchmark for comparison. Based on back-testing, the adaptive copula achieved the best var estimation performance in both groups, with average exceedance errors mostly small in magnitude and corresponding to sufficient capital reserve for covering losses at the desired levels.

The better var estimates provided by Clayton copulae indicate that the dependence structure from German equities may contain nonlinearities and asymmetries, such as stronger dependence at lower tails than at upper tails, that cannot be captured by the multivariate normal distribution. This asymmetry translates into extremely negative returns being more correlated than extremely positive returns. Thus, our results for the German equities resemble those from Longin and Solnik (2001), Ang and Chen (2002) and Patton (2006) for international markets, U.S. equities, and Deutsch mark/Japanese yen exchange rates, where empirical evidence for asymmetric dependences with increasing correlations in market downturns were found.

Furthermore, in the non-Gaussian framework, with nonlinearities and asymmetries taken into consideration through the use of Clayton copulae, the adaptive estimation produces better var fits than the moving window estimation. The high sensitive adaptive procedure can capture local changes in the dependence parameter that are not detected by the estimation with a scrolling window of fixed size.

The main advantage of using time-varying copulae to model dependence dynamics is that the normality assumption is not needed. With the proposed adaptively estimated time-varying copulae, neither normality assumption nor specification for the dependence dynamics are necessary. Hence, the method provides more flexibility in modeling dependences between markets and economies over time.

ACKNOWLEDGMENTS

Financial support from the *Deutsche Forschungsgemeinschaft* via *SFB 649 "Ökonomisches Risiko,"* Humboldt-Universität zu Berlin is gratefully acknowledged. The authors also thank the editor, an associate editor, and two referees for their helpful comments.

[Received October 2006. Revised November 2007.]

Table 5. Exceedance ratios and errors, group 2

	RiskMetrics		Moving windows		LCP	
	$\alpha = 5.00$	$\alpha = 1.00$	$\alpha = 5.00$	$\alpha = 1.00$	$\alpha = 5.00$	$\alpha = 1.00$
	$\hat{\alpha}_{w^*}$	5.42	1.58	4.53	0.39	4.53
$\hat{\alpha}_{w_1}$	5.81	1.77	5.02	0.39	5.02	0.39
$\hat{\alpha}_{w_2}$	5.62	1.58	5.12	0.39	5.22	0.30
A_W	0.16	0.57	-0.10	-0.65	-0.09	-0.65
D_W	0.04	0.16	0.06	0.09	0.06	0.08

NOTE: Exceedance ratios for portfolios w^* , w_1 , and w_2 , and average and standard deviation from relative exceedance errors. Across levels and methods, ratios and levels are expressed as a percentage.

AU8

F10, T1

T4, T5

AU9

REFERENCES

- Andrews, D. W. K. (1993), "Tests for Parameter Instability and Structural Change With Unknown Change Point," *Econometrica*, 61, 821–856.
- Andrews, D. W. K., and Ploberger, W. (1994), "Optimal Tests When a Nuisance Parameter Is Present Only Under the Alternative," *Econometrica*, 62, 1383–1414.
- Ang, A., and Chen, J. (2002), "Asymmetric Correlations of Equity Portfolios," *Journal of Financial Economics*, 63, 443–494.
- Belomestny, D., and Spokoiny, V. (2007), "Spatial Aggregation of Local Likelihood Estimates With Applications to Classification," *The Annals of Statistics*, 35, 2287–2311.
- Chen, X., and Fan, Y. (2006), "Estimation and Model Selection of Semiparametric Copula-Based Multivariate Dynamic Models Under Copula Misspecification," *Journal of Econometrics*, 135, 125–154.
- Chen, X., Fan, Y., and Tsyrennikov, V. (2006), "Efficient Estimation of Semiparametric Multivariate Copula Models," *Journal of the American Statistical Association*, 101, 1228–1240.
- Cherubini, U., Luciano, E., and Vecchiato, W. (2004), *Copula Methods in Finance*, Chichester: Wiley.
- Christoffersen, P., and Diebold, F. (2006), "Financial Asset Returns, Direction-of-Change Forecasting, and Volatility Dynamics," *Management Science*, 52, 1273–1287.
- Cont, R. (2001), "Empirical Properties of Asset Returns: Stylized Facts and Statistical Issues," *Quantitative Finance*, 1, 223–236.
- Embrechts, P., Hoeing, A., and Juri, A. (2003a), "Using Copulae to Bound the Value-at-Risk for Functions of Dependent Risks," *Finance and Stochastics*, 7, 145–167.
- Embrechts, P., Lindskog, F., and McNeil, A. (2003b), "Modelling Dependence with Copulas and Applications to Risk Management," in *Handbook of Heavy Tailed Distributions in Finance*, ed. S. Rachev, Amsterdam: North-Holland, pp. 329–384.
- Embrechts, P., McNeil, A., and Straumann, D. (2002), "Correlation and Dependence in Risk Management: Properties and Pitfalls," in *Risk Management: Value at Risk and Beyond*, ed. M. Dempster, Cambridge, UK: Cambridge University Press.
- Fan, J., and Gu, J. (2003), "Semiparametric Estimation of Value-at-Risk," *The Econometrics Journal*, 6, 261–290.
- Fleming, J., Kirby, C., and Ostdiek, B. (2001), "The Economic Value of Volatility Timing," *The Journal of Finance*, 56, 239–354.
- Franke, J., Härdle, W., and Hafner, C. (2004), *Statistics of Financial Markets*, Heidelberg: Springer-Verlag.
- Fréchet, M. (1951), "Sur les Tableaux de Correlation Dont les Marges Sont Données," *Annales de l'Université de Lyon, Sciences Mathématiques et Astronomie*, 14, 5–77.
- Giacomini, E., and Härdle, W. (2005), "Value-at-Risk Calculations With Time Varying Copulae," in *Bulletin of the International Statistical Institute, Proceedings of the 55th Session*.
- Granger, C. (2003), "Time Series Concept for Conditional Distributions," *Oxford Bulletin of Economics and Statistics*, 65, 689–701.
- Hansen, B. E. (2001), "The New Econometrics of Structural Change: Dating Breaks in U.S. Labor Productivity," *The Journal of Economic Perspectives*, 15, 117–128.
- Härdle, W., Herwartz, H., and Spokoiny, V. (2003), "Time Inhomogeneous Multiple Volatility Modelling," *Journal of Financial Econometrics*, 1, 55–95.
- Härdle, W., Kleinow, T., and Stahl, G. (2002), *Applied Quantitative Finance*, Springer-Verlag, Heidelberg.
- Hoëffding, W. (1940), "Maßstabvariante Korrelationstheorie," *Schriften des mathematischen Seminars und des Instituts für angewandte Mathematik der Universität Berlin*, 5, 181–233.
- Hu, L. (2006), "Dependence Patterns Across Financial Markets: A Mixed Copula Approach," *Applied Financial Economics*, 16, 717–729.
- Ibragimov, R. (2007), "Efficiency of Linear Estimators Under Heavy-Tailedness: Convolutions of α -Symmetric Distributions," *Econometric Theory*, 23, 501–517.
- Ibragimov, R., and Walden, J. (2007), "The Limits of Diversification When Losses May be Large," *Journal of Banking and Finance*, 31, 2551–2569.
- Joe, H. (1997), *Multivariate Models and Dependence Concepts*, London: Chapman & Hall.
- Jorion, P. (1995), "Predicting Volatility in the Foreign Exchange Market," *The Journal of Finance*, 50, 507–528.
- Morgan, J. P. (1996), *RiskMetrics Technical Document*, New York: RiskMetrics Group.
- Kim, J., Malz, A. M., and Mina, J. (1999), *Long Run Technical Document*, New York: RiskMetrics Group.
- Longin, F., and Solnik, B. (2001), "Extreme Correlation on International Equity Markets," *The Journal of Finance*, 56, 649–676.
- Mari, D., and Kotz, S. (2001), *Correlation and Dependence*, London: Imperial College Press.
- Marshall, A., and Olkin, I. (1979), *Inequalities: Theory of Majorizations and Its Applications*, New York: Academic Press.
- McNeil, A. J., Frey, R., and Embrechts, P. (2005), *Quantitative Risk Management: Concepts, Techniques and Tools*, Princeton, NJ: Princeton University Press.
- Mercurio, D., and Spokoiny, V. (2004), "Estimation of Time Dependent Volatility via Local Change Point Analysis With Applications to Value-at-Risk," *Annals of Statistics*, 32, 577–602.
- Nelsen, R. (1998), *An Introduction to Copulas*, New York: Springer-Verlag.
- Patton, A. (2004), "On the Out-of-Sample Importance of Skewness and Asymmetric Dependence for Asset Allocation," *Journal of Financial Econometrics*, 2, 130–168.
- (2006), "Modelling Asymmetric Exchange Rate Dependence," *International Economic Review*, 47, 527–556.
- Perron, P. (1989), "The Great Crash, the Oil Price Shock and the Unit Root Hypothesis," *Econometrica*, 57, 1361–1401.
- Polzehl, J., and Spokoiny, V. (2006), "Propagation–Separation Approach for Likelihood Estimation," *Probability Theory and Related Fields*, 135, 335–362.
- Quintos, C., Fan, Z., and Philips, P. C. B. (2001), "Structural Change Tests in Tail Behaviour and the Asian Crisis," *The Review of Economic Studies*, 68, 633–663.
- Rodriguez, J. C. (2007), "Measuring Financial Contagion: A Copula Approach," *Journal of Empirical Finance*, 14, 401–423.
- Sklar, A. (1959), "Fonctions de Répartition à n Dimensions et Leurs Marges," *Publications de l'Institut de Statistique de l'Université de Paris*, 8, 229–231.
- Spokoiny, V. (2008), *Local Parametric Methods in Nonparametric Estimation*, Berlin, Heidelberg: Springer-Verlag.
- Spokoiny, V., and Chen, Y. (2007), *Multiscale Local Change Point Detection with Applications to Value-at-Risk*, Preprint 904, Berlin: Weierstrass Institute Berlin.
- Stock, J.H. (1994), "Unit Roots, Structural Breaks and Trends," in *Handbook of Econometrics*, Vol. 4, ed. R. F. Engle and D. McFadden, Amsterdam: North-Holland, pp. 2739–2841.
- Zivot, E., and Andrews, D. W. K. (1992), "Further Evidence on the Great Crash, the Oil Price Shock and the Unit Root Hypothesis," *Journal of Business & Economic Statistics*, 10, 251–270.



Dynamics of state price densities

Wolfgang Härdle^a, Zdeněk Hlávka^{b,*}

^a CASE—Center for Applied Statistics and Economics, Wirtschaftswissenschaftliche Fakultät, Humboldt-Universität zu Berlin, Spandauer Str. 1, 10178 Berlin, Germany

^b Charles University in Prague, Department of Statistics, Sokolovská 83, 18675 Praha, Czech Republic

ARTICLE INFO

Article history:

Received 11 January 2009

Accepted 12 January 2009

Available online 15 January 2009

JEL classification:

C13

C14

G13

Keywords:

Option pricing

State price density

Nonlinear least squares

Constrained estimation

ABSTRACT

State price densities (SPDs) are an important element in applied quantitative finance. In a Black–Scholes world they are lognormal distributions, but in practice volatility changes and the distribution deviates from log-normality. In order to study the degree of this deviation, we estimate SPDs using EUREX option data on the DAX index via a nonparametric estimator of the second derivative of the (European) call pricing function. The estimator is constrained so as to satisfy no-arbitrage constraints and corrects for the intraday covariance structure in option prices. In contrast to existing methods, we do not use any parametric or smoothness assumptions.

© 2009 Elsevier B.V. All rights reserved.

1. Introduction

The dynamics of option prices carries information on changes in state price densities (SPDs). The SPD contains important information on the behavior and expectations of the market and is used for pricing and hedging. The most important application of an SPD is that it allows one to price options with complicated payoff functions simply by (numerical) integration of the payoff with respect to this density.

Prices $C_t(K, T)$ of European options with strike price K observed at time t and expiring at time T allow one to deduce the state price density $f(\cdot)$ using the relationship (Breedon and Litzenberger, 1978)

$$f(K) = \exp\{r(T-t)\} \frac{\partial^2 C_t(K, T)}{\partial K^2}. \quad (1)$$

Eq. (1) can be used to estimate the SPD $f(K)$ from the observed option prices. An extensive overview of parametric and other estimation techniques can be found, for example, in Jackwerth (1999). An application to option pricing is given in Buehler (2006).

Kernel smoothers were in this framework proposed and successfully applied by, for example, Ait-Sahalia and Lo (1998), Ait-Sahalia and Lo (2000), Ait-Sahalia et al. (2000), or Huynh et al. (2002). Ait-Sahalia and Duarte (2003) proposed a method for

nonparametric estimation of the SPD under constraints like positivity, convexity, and boundedness of the first derivative. Bondarenko (2003) calculates arbitrage-free SPD estimates using positive convolution approximation (PCA) methodology and demonstrates its properties in a Monte Carlo study based on closing prices of the S&P 500 options. Another sophisticated approach based on smoothing splines allowing one to include these constraints is described and applied on simulated data in Yatchew and Härdle (2006). In the majority of these papers, the focus was more on the smoothing techniques rather than on a no-arbitrage argument, although a crucial element of local volatility models is the absence of arbitrage (Dupire, 1994). Highly numerically efficient pricing algorithms, for example, by Andersen and Brotherton-Ratcliffe (1997), rely heavily on no-arbitrage properties. Kahalé (2004) proposed a procedure that requires solving a set of nonlinear equations with no guarantee of a unique solution. Moreover, for that algorithm the data feed is already (unrealistically) expected to be arbitrage free (Fengler, 2005; Fengler et al., 2007). In addition, the covariance structure of the quoted option prices (Renault, 1997) is rarely incorporated into the estimation procedure.

In Table 1, we give an overview of selected properties of different estimation techniques. The parametric approach may be used to estimate parameters of a probability density lying in some preselected family. The parametric models may be further extended by considering more flexible probability densities or mixtures of distributions. Approaches based on nonparametric smoothing techniques are more flexible since the shape of a nonparametric SPD estimate is not fixed in advance and the method controls only the smoothness of the estimate. For example,

* Corresponding author. Tel.: +420 221 913 284; fax: +420 283 073 341.

E-mail addresses: haerdle@wiwi.hu-berlin.de (W. Härdle), hlavka@karlin.mff.cuni.cz (Z. Hlávka).

Table 1
Summary of properties of parametric and nonparametric estimators.

	Methods			
	Parametric	Standard smoothing method	Nonparametric under constraints	This paper
Shape	Fixed	Flexible	Flexible	Flexible
Control	Choice of family	Smoothness	Smoothness	None
SPD support	Infinite	Restricted	Restricted	Restricted
Constraints	By design	Local	Yes	Yes

the smoothness of a kernel regression estimator depends mostly on the choice of the bandwidth parameter, the smoothness of the PCA estimator (Bondarenko, 2003) depends on the choice of the kernel, and the smoothness of the NNLS estimator (Yatchew and Härdle, 2006) is controlled by constraining the Sobolev norm of the SPD; using these nonparametric estimators, systematic bias may typically occur in the case of oversmoothing. Constraints on estimators are more easily implemented for globally valid parametric models than for local (nonparametric) models. The use of a standard smoothing technique which does not account for the constraints is not advisable. The value of the nonparametric estimate cannot be calculated in regions without any data and, therefore, the support of nonparametrically estimated SPDs is limited by the range of the observed strike prices even for nonparametric-under-constraints techniques.

Most of the commonly used estimation techniques do not specify explicitly the source of random error in the observed option prices; see Renault (1997) for an extensive review of this subject. A common approach in SPD estimation is to use either the closing option prices or to correct the intraday option prices by the current value of the underlying asset. Both approaches lack interpretation if the shape of the SPD changes rapidly. This can be made clear by a gedankenexperiment: if the shape of the SPD changes dramatically during the day, correcting the observed option prices by the value of the underlying asset and then estimating the SPD would lead to an estimate of some (nonexisting) daily average of the true SPDs. We try to circumvent this problem by introducing a simple model for the intraday covariance structure of option prices which allows us to estimate the value of the true SPD at an arbitrarily chosen fixed time; see also Hlávka and Svojík (2008). Most often, we are interested in the estimation of the current SPD.

We develop a simple estimation technique in order to construct constrained SPD estimates from the observed intraday option prices which are treated as repeated observations collected during a certain time period. The proposed technique involves constrained LS-estimation, it enables us to construct confidence intervals for the current value of the SPD and prediction intervals for its future development, and it does not depend on any tuning (smoothness) parameter. The construction of a simple approximation of the covariance structure of the observed option prices follows naturally from the derivation of our nonparametric constrained estimator. This covariance structure is interesting in itself; it separates two sources of random errors, and it is applicable to other SPD estimators.

We study the development of the estimated SPDs in Germany over 8 years. A no-arbitrage argument is imposed at each time point, leading (mathematically) to the above-mentioned no-arbitrage constraints. This, of course, is a vital feature for trading purposes where the derived (implied) volatility surfaces for different strikes and maturities are needed for proper judgment of risk and return.

The resulting SPDs and implied volatility surfaces are not smooth per se. In most applications, this is not a disadvantage though, since, first, we may smooth the resulting SPD estimates (Hlávka and Svojík, 2008) and, second, we are mostly interested in functionals of the estimated SPD like, for example, the expected payoff or the forward price. Another important feature that can be

easily estimated from the nonsmooth SPDs are the quantiles; see Section 6.2 for an application.

In Section 2, we introduce the notation, discuss constraints that are necessary for estimating SPDs, and we construct a very simple unconstrained SPD estimator using simple linear regression. In Section 3, this estimator is modified so that it satisfies the shape constraints given in Section 2.1. We demonstrate that the covariance structure of the option prices exhibits correlations depending both on the strike price and time of the trade in Section 4. In Section 5, we apply our estimation technique on option prices observed in the year 1995, and we show that the proposed approximation of the covariance structure removes the dependency and heteroscedasticity of the residuals. The dynamics of the estimated SPDs in years 1995–2003 is studied in Section 6.

2. Construction of the estimate

The fair price of a European call option with payoff $(S_T - K)_+ = \max(S_T - K, 0)$, with S_T denoting the price of the stock at time T , t the current time, K the strike price, and r the risk-free interest rate, can be written as

$$C_t(K, T) = \exp\{-r(T - t)\} \int_0^\infty (S_T - K)_+ f(S_T) dS_T, \tag{2}$$

i.e., as the discounted expected value of the payoff with respect to the SPD $f(\cdot)$. For the sake of simplicity of the following presentation, we assume in the rest of the paper that the discount factor $\exp\{-r(T - t)\} = 1$. In applications, this is achieved by correcting the observed option prices by the known risk-free interest rate r and the time to maturity $(T - t)$ in (2). At the time of the trade, the current index price and volatility are common to all options and, hence, do not appear explicitly in Eq. (2).

Let us denote the i -th observation of the strike price by K_i and the corresponding option price, divided by the discount factor $\exp\{-r(T - t)\}$ from (2), by $C_i = C_{t,i}(K_i, T)$. In practice, on any given day t , one observes option prices repeatedly for a small number of distinct strike prices. Therefore, it is useful to adopt the following notation. Let $\mathcal{C} = (C_1, \dots, C_n)^\top$ be the vector of the observed option prices on day t sorted by strike price. Then, the vector of strike prices has the following structure:

$$\mathcal{K} = \begin{pmatrix} K_1 \\ K_2 \\ \vdots \\ K_n \end{pmatrix} = \begin{pmatrix} k_1 \mathbf{1}_{n_1} \\ k_2 \mathbf{1}_{n_2} \\ \vdots \\ k_p \mathbf{1}_{n_p} \end{pmatrix},$$

where $k_1 < k_2 < \dots < k_p$, $n_j = \sum_{i=1}^n \mathbf{I}(K_i = k_j)$, with $\mathbf{I}(\cdot)$ denoting the indicator function and $\mathbf{1}_n$ a vector of ones of length n .

2.1. Assumptions and constraints

Let us now concentrate on options corresponding to a single maturity T observed at fixed time t . Let us assume that the i -th observed option price (corresponding to strike price K_i) follows the model

$$C_{t,i}(K_i, T) = \mu(K_i) + \varepsilon_i, \tag{3}$$

where ε_i are iid random variables with zero mean and variance σ^2 . In practice, one might expect that the errors exhibit correlations depending on the strike price and time. Heteroscedasticity can

be incorporated in model (3) if we assume that the random errors ε_i have variance $\text{Var } \varepsilon_i = \sigma_{k_i}^2$, leading to weighted least squares. The assumptions on the distribution of random errors will be investigated in more detail in Section 5.3. Following Renault (1997), we interpret the observed option price as the price given by a pricing formula plus an error term, and in Section 4 we suggest a covariance structure for the observed option prices taking into account the dependencies across strike prices and times of trade.

Harrison and Pliska (1981) characterized the absence of arbitrage by the existence of a unique risk neutral SPD $f(\cdot)$. From formula (2) and the properties of a probability density it follows that, in a continuous setting, the function $\mu(\cdot)$, defined on \mathbb{R}^+ , has to satisfy the following no-arbitrage constraints:

- 1': it is positive,
- 2': it is decreasing in K ,
- 3': it is convex,
- 4': its second derivative exists and it is a density (i.e., nonnegative and it integrates to one).

Let us now have a look at functions satisfying Constraints 1'–4'.

Lemma 1. Suppose that $\mu : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ satisfies Constraints 1'–4'. Then the first derivative, $\mu^{(1)}(\cdot)$, is nondecreasing and such that $\lim_{x \rightarrow 0} \mu^{(1)}(x) = -1$ and $\lim_{x \rightarrow +\infty} \mu^{(1)}(x) = 0$.

Proof. Constraint 4' implies that the first derivative, $\mu^{(1)}$, exists and that it is differentiable. $\lim_{x \rightarrow +\infty} \mu^{(1)}(x)$ exists since the function $\mu^{(1)}$ is nondecreasing (Constraint 3') and bounded (Constraint 2'). Next, $\lim_{x \rightarrow \infty} \mu^{(1)}(x) = 0$ since a negative limit would violate Constraint 1' for large x ($\mu^{(1)}(x)$ cannot be positive since $\mu(x)$ is decreasing). Finally, Constraint 4', $1 = \int_0^\infty \mu^{(2)}(x) dx = \lim_{x \rightarrow +\infty} \mu^{(1)}(x) - \lim_{x \rightarrow 0} \mu^{(1)}(x)$, implies that $\lim_{x \rightarrow 0} \mu^{(1)}(x) = -1$. \square

Remark 1. Lemma 1 allows us to restate Constraints 3' and 4' in terms of $\mu^{(1)}(\cdot)$ by assuming that $\mu^{(1)}(\cdot)$ is differentiable, nondecreasing, and such that $\lim_{x \rightarrow 0} \mu^{(1)}(x) = -1$ and $\lim_{x \rightarrow +\infty} \mu^{(1)}(x) = 0$.

In this section, we stated only constraints guaranteeing that the SPD estimate will be a probability density. Constraints for the expected value of the SPD estimate are discussed in Section 3.6.

2.2. Existence and uniqueness

In this subsection we address the issue of existence and uniqueness of a regression function, $\hat{C}(\cdot)$, satisfying the required assumptions and constraints. In practice, we do not deal with a continuous function. Hence, we restate Constraints 1'–4' for discrete functions, defined only on a finite set of distinct points, say $k_1 < \dots < k_p$, in terms of their function values, $C(k_i)$, and their scaled first differences, $C_{k_i, k_j}^{(1)} = \{C(k_i) - C(k_j)\} / \{k_i - k_j\}$.

- 1: $C(k_i) \geq 0, i = 1, \dots, p$,
- 2: $k_i < k_j$ implies that $C(k_i) \geq C(k_j)$,
- 3: $k_i < k_j < k_l$ implies that $-1 \leq C_{k_i, k_j}^{(1)} \leq C_{k_j, k_l}^{(1)} \leq 0$.

It is easy to see that Constraints 1–2 are discrete versions of Constraints 1' and 2'. Constraint 3 is a discrete version of Constraints 3' and 4'; see Remark 1.

From now on, similarly as in Robertson et al. (1988), we think of the collection, \mathcal{C} , of functions satisfying Constraints 1–3 as a subset of a p -dimensional Euclidean space, where p is the number of distinct k_i 's. The constrained regression, \hat{C} , is in this setting the closest point of \mathcal{C} to the vector C of the observed option prices with distances measured by the usual Euclidean distance

$$d(f, C) = (f - C)^\top (f - C) = \sum_{i=1}^n \{f(K_i) - C(K_i)\}^2. \quad (4)$$

From this point of view, the regression function, \hat{C} , consists only of the values of the function in the points k_1, \dots, k_p . The first and second differences are used to approximate the first and the second derivatives, respectively.

We claim that the set, \mathcal{C} , of functions satisfying Constraints 1–3 is closed in the topology induced by the metric given by Euclidean distance and it is convex, i.e., if $f, g \in \mathcal{C}$ and $0 \leq a \leq 1$, then $af + (1 - a)g \in \mathcal{C}$.

Lemma 2. If $\hat{C} \in \mathcal{C}$ is the regression of $C(K_i), i = 1, \dots, n$, on $k_1 < \dots < k_p$ under Constraints 1–3 and if a and b are constants such that $a \leq C(K_i) \leq b, \forall i$, then $a \leq \hat{C}(k_i) \leq b + (k_p - k_i)$.

Proof. It is not possible that $\hat{C}(k_i)$ lies above b for all k_i 's (otherwise we would get a better fit only by shifting $\hat{C}(k_i)$). The upper bound now follows from Constraint 3.

The validity of the lower bound may be demonstrated similarly. Clearly, it is not possible that $\hat{C}(k_i)$ lie below a for all k_i 's. Moreover, it is not possible that $\hat{C}(k_1) \geq \dots \geq \hat{C}(k_i) \geq a > \hat{C}(k_{i+1}) \geq \dots \geq \hat{C}(k_p)$ for any i , since in such a situation the fit could be trivially improved by increasing $\hat{C}(k_{i+1}), \dots, \hat{C}(k_p)$ by some small amount, for example, by $a - \hat{C}(k_{i+1})$, without violating any of the Constraints 1–3. \square

Theorem 1. A regression, $\hat{C} = \arg \min_{f \in \mathcal{C}} d(f, C)$, satisfying Constraints 1–3, exists and it is unique.

Proof. Lemma 2 implies that \hat{C} belongs to a subset, \mathcal{B} , of \mathcal{C} bounded below by a and above by $b + (k_p - k_1)$. Thinking of the functions as points in Euclidean space, it is clear that the continuous function $d(f, C)$ attains its minimum on the closed and bounded set \mathcal{B} . The uniqueness of \hat{C} follows from the convexity of \mathcal{B} using, for example, Robertson et al. (1988, Theorem 1.3.1). \square

2.3. Linear model

With the given option data, Constraints 1–3 of Section 2.2 can be reformulated using linear regression models with constraints.

In the following, we fix the time t and the expiry date T and we omit these symbols from the notation. In Section 2.2 we have noted that the option prices are repeatedly observed for a small number p of distinct strike prices. Defining the expected values of the option prices for a given strike price, $\mu_j = \mu(k_j) = E\{C(k_j)\}$, we can write

$$\begin{aligned} \mu_p &= \beta_0, \\ \mu_{p-1} &= \beta_0 + \beta_1, \\ \mu_{p-2} &= \beta_0 + 2\beta_1 + \beta_2, \\ \mu_{p-3} &= \beta_0 + 3\beta_1 + 2\beta_2 + \beta_3, \\ &\vdots \\ \mu_1 &= \beta_0 + (p-1)\beta_1 + (p-2)\beta_2 + \dots + \beta_{p-1}. \end{aligned}$$

Thus, we fit our data using coefficients $\beta_j, j = 1, \dots, p$. The conditional means $\mu_i, i = 1, \dots, p$ are replaced by the same number of parameters $\beta_j, j = 0, \dots, p-1$, which allow us to impose the shape constraints in a more natural way.

The interpretation of the coefficients β_j can be seen in Fig. 1, which shows a simple situation with only four distinct strike prices ($p = 4$). β_0 is the mean option price at point 4. Constraint 1', Section 2.1, implies that it has to be positive. β_1 is the difference between the mean option prices at point 4 and point 3; Constraint 2' implies that it has to be positive. The next coefficient, β_2 , approximates the change in first derivative in point 3 and it can be interpreted as an approximation of the second derivative in point 3. Constraint 3' implies that β_2 has to be positive. Similarly, β_3 is an estimate of the (positive) second derivative in point 2. Constraint 4' can be rewritten as $\beta_2 + \beta_3 \leq 1$.




Fig. 1. Illustration of the dummy variables for call options.

In practice, we start with the construction of a design matrix which allows us to write the above model in the following linear form. For simplicity of presentation, we again set $p = 4$:

$$\begin{pmatrix} \mu_1 \\ \mu_2 \\ \mu_3 \\ \mu_4 \end{pmatrix} = \begin{pmatrix} 1 & 3 & 2 & 1 \\ 1 & 2 & 1 & 0 \\ 1 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \end{pmatrix} \begin{pmatrix} \beta_0 \\ \beta_1 \\ \beta_2 \\ \beta_3 \end{pmatrix}. \tag{5}$$

Ignoring the constraints on the coefficients would lead to a simple linear regression problem. Unfortunately, this approach does not have to lead, and usually does not, to interpretable and stable results.

Model (5) in the above form can be reasonably interpreted only if the observed strike prices are equidistant and if the distances between the neighboring observed strike prices are equal to one. If we want to keep the interpretation of the parameters β_j as the derivatives of the estimated function, we should use the design matrix

$$\Delta = \begin{pmatrix} 1 & \Delta_p^1 & \Delta_{p-1}^1 & \Delta_{p-2}^1 & \cdots & \Delta_3^1 & \Delta_2^1 \\ 1 & \Delta_p^2 & \Delta_{p-1}^2 & \Delta_{p-2}^2 & \cdots & \Delta_3^2 & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ 1 & \Delta_p^{p-2} & \Delta_{p-1}^{p-2} & 0 & \cdots & 0 & 0 \\ 1 & \Delta_p^{p-1} & 0 & 0 & \cdots & 0 & 0 \\ 1 & 0 & 0 & 0 & \cdots & 0 & 0 \end{pmatrix}, \tag{6}$$

where $\Delta_j^i = \max(k_j - k_i, 0)$ denotes the positive part of the distance between k_i and k_j , the i -th and the j -th ($1 \leq i \leq j \leq p$) sorted distinct observed values of the strike price.

The vector of conditional means μ can be written in terms of the parameters β as follows:

$$\begin{pmatrix} \mu_1 \\ \mu_2 \\ \vdots \\ \mu_p \end{pmatrix} = \mu = \Delta\beta = \Delta \begin{pmatrix} \beta_0 \\ \beta_1 \\ \vdots \\ \beta_{p-1} \end{pmatrix}. \tag{7}$$

The constraints on the conditional means μ_j can now be expressed as conditions on the parameters of the model (7). Namely, it suffices to request that $\beta_i > 0$, $i = 0, \dots, p - 1$ and that $\sum_{j=2}^{p-1} \beta_j \leq 1$.

The model for the option prices can now be written as

$$C(\mathcal{K}) = \mathcal{X}_\Delta \beta + \varepsilon, \tag{8}$$

where \mathcal{X}_Δ is the design matrix obtained by repeating each row of matrix Δ n_i times, $i = 1, \dots, p$.

3. Implementing the constraints

In order to impose Constraints 1–3 on parameters β_i , $i = 0, \dots, p - 1$, we propose the following reparameterization of the

model in terms of parameters $\theta = (\theta_0, \dots, \theta_{p-1})^\top$:

$$\beta_0(\theta) = \exp(\theta_0),$$

$$\beta_1(\theta) = \exp(\theta_1),$$

\vdots

$$\beta_{p-1}(\theta) = \exp(\theta_{p-1}),$$

under the constraint that $\sum_{j=2}^{p-1} \exp(\theta_j) < 1$. Clearly, the parameters $\beta_i(\theta)$ satisfy the constraints

$$\beta_i(\theta) > 0, \quad i = 0, \dots, p - 1,$$

$$\sum_{j=2}^{p-1} \beta_j(\theta) < 1.$$

This means that the parameters $\beta_2(\theta), \dots, \beta_{p-1}(\theta)$ can be considered as point estimates of the state price density (the estimates have to be positive and integrate to less than one). Furthermore, in view of Lemma 1, it is worthwhile to note that the parameters also satisfy

$$-\sum_{j=1}^k \beta_j \in (-1, 0), \quad \text{for } k = 1, \dots, p - 1.$$

The model (8) rewritten in terms of parameters θ_i , $i = 0, \dots, p$, is a nonlinear regression model which can be estimated using standard nonlinear least squares or maximum likelihood methods (Seber and Wild, 2003). The main advantage of these methods is that the asymptotic distribution is well known and that the asymptotic variance of the estimator can be approximated using numerical methods implemented in many statistical packages.

3.1. Reparameterization

The following reparameterization of the model in terms of parameters $\xi = (\xi_0, \dots, \xi_p)^\top$ simplifies the calculation of the estimates because it guarantees that all constraints are automatically satisfied:

$$\beta_0(\xi) = \exp(\xi_0),$$

$$\beta_1(\xi) = \frac{\exp(\xi_1)}{\sum_{j=1}^p \exp(\xi_j)},$$

\vdots

$$\beta_{p-1}(\xi) = \frac{\exp(\xi_{p-1})}{\sum_{j=1}^p \exp(\xi_j)}.$$

This property simplifies the numerical minimization algorithm needed for the calculation of the estimates.

The equality

$$\frac{1}{\sum_{j=1}^{p-1} \beta_j(\xi)} = 1 + \frac{\exp(\xi_p)}{\sum_{j=1}^{p-1} \exp(\xi_j)}$$

shows the meaning of the additional parameter ξ_p . Setting this parameter to $-\infty$ would be the same as requiring that $\sum_{j=1}^{p-1} \beta_j(\xi) = 1$. Large values of the parameter ξ_p indicate that the estimated coefficients sum to less than one or, in other words, the observed strike prices do not cover the support of the estimated SPD. Notice that, by setting $\xi_p = -\infty$, we could easily modify our procedure and impose the equality constraint $\sum_{j=1}^{p-1} \beta_j(\xi) = 1$.

3.2. Inverse transformation of model parameters

For the numerical algorithm, it is useful to know how to calculate ξ 's from given β 's. This is needed, for example, to obtain reasonable starting points for the iterative procedure maximizing the likelihood.

Lemma 3. Given $\beta = (\beta_1, \dots, \beta_p)^\top$, where $\beta_p = 1 - \sum_{i=1}^{p-1} \beta_i$, the parameters $\xi = (\xi_1, \dots, \xi_p)^\top$ satisfy the system of equations

$$(\beta \mathbf{1}_p^\top - \mathbf{I}_p) \exp \xi^\top = \mathcal{A} \exp \xi^\top = 0, \tag{9}$$

where \mathbf{I}_p is the $(p \times p)$ identity matrix. Furthermore,

$$\text{rank} \mathcal{A} = p - 1. \tag{10}$$

The system of Eq. (9) has infinitely many solutions, which can be expressed as

$$\exp(\xi) = (\mathcal{A}^- \mathcal{A} - \mathbf{I}_p) z, \tag{11}$$

where \mathcal{A}^- denotes a generalized inverse of \mathcal{A} and where z is an arbitrary vector in \mathbb{R}^p such that the right-hand side of (11) is positive.

Proof. Parts (9) and (10) follow from the definition of $\beta(\xi)$ and from simple algebra (notice that the sum of rows of \mathcal{A} is equal to zero). Part (11) follows, for example, from Anděl (1985, Theorem IV.18). \square

It remains to choose the vector z in (11) so that the solution of the system of Eq. (9) is positive.

Proposition 1. The rank of the matrix $\mathcal{A}^- \mathcal{A} - \mathbf{I}_p$ is 1. Hence, any solution of the system of Eq. (9) is a multiple of the first column of the matrix $\mathcal{A}^- \mathcal{A} - \mathbf{I}_p$. The vector z in (11) can be chosen, for example, as $z = \pm \mathbf{1}_p$, where the sign is chosen so that the resulting solution is positive.

Proof. The definition of a generalized inverse is

$$\mathcal{A} \mathcal{A}^- \mathcal{A} - \mathcal{A} = \mathcal{A} (\mathcal{A}^- \mathcal{A} - \mathbf{I}_p) = 0. \tag{12}$$

Lemma 3 says that $\text{rank} \mathcal{A} = p - 1$ and, hence, Eq. (12) implies that $\text{rank}(\mathcal{A}^- \mathcal{A} - \mathbf{I}_p) \leq 1$. Noticing that $\mathcal{A}^- \mathcal{A} \neq \mathbf{I}_p$ means that $\text{rank}(\mathcal{A}^- \mathcal{A} - \mathbf{I}_p) > 0$, and concludes the proof. \square

3.3. The algorithm

The proposed algorithm consists of the following steps:

- 1: obtain a reasonable initial estimate $\hat{\beta}$, for example, by running the Pool-Adjacent-Violators algorithm (Robertson et al., 1988, Chapter 1) on the unconstrained least squares estimates of the first derivative of the curve,
- 2: transform the initial estimate $\hat{\beta}$ into the estimate $\hat{\xi}$ using the method described in Section 3.2,
- 3: estimate the parameters of the model (8) by minimizing the sum of squares $\{C(\mathcal{K}) - \mathcal{X}_\Delta \beta(\xi)\}^\top \{C(\mathcal{K}) - \mathcal{X}_\Delta \beta(\xi)\}$ in terms of ξ (see Section 3.1) using numerical methods.

An application of this simple algorithm on real data is given in Section 5.1.

3.4. Asymptotic confidence intervals

We construct confidence intervals based on the parameterization $\beta(\theta)$ introduced at the beginning of this section. The confidence limits for parameters θ_i are exponentiated in order to obtain valid pointwise confidence bounds for the true SPD. The main advantage of this approach is that such confidence bounds are always positive.

An alternative approach would be to construct confidence intervals based on the parameterizations in terms of β_i (Section 2.3) or ξ_i (Section 3.1). However, the limits of confidence intervals for β_i may be negative and confidence intervals for the SPD based on parameters ξ_i would have very complicated shapes in high-dimensional space and could not be easily calculated and interpreted.

Another approach to the construction of the asymptotic confidence intervals can be based on the maximum likelihood theory. Assuming normality, the log-likelihood for the model (8) can be written as

$$l(C, \mathcal{X}_\Delta, \theta, \sigma) = -n \log \sigma - \frac{1}{2\sigma^2} \{C - \mathcal{X}_\Delta \beta(\theta)\}^\top \times \{C - \mathcal{X}_\Delta \beta(\theta)\}, \tag{13}$$

where \mathcal{X}_Δ is the design matrix given in (8). This normality assumption is justified later by a residual analysis. The maximum likelihood estimator is defined as

$$\hat{\theta} = \arg \max_{\theta} l(C, \mathcal{X}_\Delta, \theta, \sigma), \tag{14}$$

and it has asymptotically a p -dimensional normal distribution with mean θ and the variance given by the inverse of the Fisher information matrix:

$$\mathcal{F}_n^{-1} = \left\{ -E \left(\frac{\partial^2}{\partial \theta \partial \theta^\top} l(C, \mathcal{X}_\Delta, \theta, \sigma) \right) \right\}^{-1}. \tag{15}$$

More precisely, $n^{1/2}(\hat{\theta} - \theta) \xrightarrow{\mathcal{L}} N_p(0, \mathcal{F}_n^{-1})$. In this framework, the Fisher information matrix can be estimated by using the numerically differentiated Hessian matrix of the log-likelihood. For details we refer, for example, to Serfling (1980, Chapter 4). The confidence intervals calculated for parameters θ may be transformed (exponentiated) to a confidence intervals for the SPD (β). We have not pursued the maximum likelihood approach since it was numerically less stable in this situation.

Note that, under the assumptions of normality, the maximum likelihood estimate is equal to the nonlinear least squares estimate (Seber and Wild, 2003, Section 2.2), and the asymptotic variance of $\hat{\theta} = \exp(\beta)$ may be approximated by $\text{Var} \hat{\theta} = \{\text{diag}(\exp \hat{\theta}) \mathcal{X}_\Delta^\top \mathcal{X}_\Delta \text{diag}(\exp \hat{\theta})\}^{-1} \hat{\sigma}^2$. Hence, asymptotic confidence intervals for θ_i may be calculated as $(\hat{\theta}_i \pm u_{1-\alpha/2} \hat{s}_{ii})$, where $u_{1-\alpha/2}$ is the $1 - \alpha/2$ quantile of the standard Normal distribution and \hat{s}_{ii} denotes the i -th diagonal element of $\text{Var} \hat{\theta}$. By exponentiating both limits of this confidence interval, we immediately obtain the $1 - \alpha$ confidence interval for $\beta_i = \exp \theta_i$.

The construction of the estimator guarantees that the matrix \mathcal{X}_Δ has full rank—this implies that $\mathcal{X}_\Delta^\top \mathcal{X}_\Delta$ is invertible and the asymptotic variance matrix $\text{Var} \hat{\theta}$ always exists. If the number of observations is equal to the number of distinct strike prices (if there is only one option price for each strike price), it may happen that $\hat{\sigma}^2 = 0$ and the confidence intervals degenerate to a single point.

3.5. Put–Call parity

The prices of put options can be easily included in our estimation technique by applying the Put–Call parity of the option prices. Assuming that there are no dividends or costs connected with the ownership of the stock, each put option with price $P_t(K, T)$ corresponds to a call option with price

$$C_t(K, T) = P_t(K, T) + S_t - Ke^{-r(T-t)}.$$

In this way, the prices of the put options can be converted into the prices of call options and used in our model (Stoll, 1969). Statistically speaking, these additional observations will increase the precision of the SPD and will lead to more stable results.

In Germany, the Put–Call parity might be biased by an effect of the DAX index calculation which is based on the assumption that




Fig. 2. Illustration of the dummy variables for both call (β) and put (α) options.

the dividends are reinvested after deduction of corporate income tax. As the income tax of some investors might be different, the value of the DAX has to be corrected before using Put–Call parity in subsequent analysis. For the exact description of this correction we refer to Hafner and Wallmeier (2000) who were analyzing the same data set.

The construction of our estimates allows us to include the put option prices in a more direct way by fitting the two curves separately using two sets of parameters. The situation is displayed in Fig. 2. Our assumption that the same SPD drives both the put and call option prices is naturally translated in terms of the coefficients α_i and β_i :

$$\alpha_i = \beta_{p-i+1}, \quad \text{for } i = 2, \dots, p - 1$$

$$\alpha_1 = 1 - \sum_{i=1}^{p-1} \beta_i.$$

The problem of estimating regression functions under such linear equality constraints is solved, for example, in Rao (1973). In Section 4.3, we will also investigate the covariance of the observed call and put option prices, and the suggested model will be presented in detail.

3.6. Expected value constraints

In Section 2.3, we have explained that the parameters $\beta_2, \dots, \beta_{p-1}$ can be interpreted as estimates of the state price density in points k_2, \dots, k_{p-1} . From the construction of the estimator, see also Fig. 1, it follows that parameter β_1 can be interpreted as the mass of the SPD lying to the right of k_{p-1} . Assuming that the observed strike prices entirely cover the support of the SPD, the mass β_1 could be attributed to the point k_p . Notice that the reparameterization introduced in Section 3 guarantees that $\sum_{i=1}^{p-1} \beta_i(\xi) < 1$, and it immediately follows that interpreting β_1 as the estimate of the SPD in point k_p does not violate any constraints described in Section 2.2.

Referring to Section 3.5, it is clear that the parameter $\beta_p \equiv \alpha_1 = 1 - \sum_{i=1}^{p-1} \beta_i$ can be interpreted as the estimator of the SPD in k_1 . The parameterization of the problem now guarantees that $\sum_{i=1}^p \beta_i = 1$.

The expected value of the underlying stock under the risk-neutral measure can now be estimated as $E^{\text{SPD}} = \sum_{i=1}^p k_i \beta_{p-i+1}$. From economic theory it follows that E^{SPD} has to be equal to the forward price of the stock. This constraint can be easily implemented by using the fact that β_1 and β_p estimate the mass of the SPD respectively to the right of k_{p-1} and to the left of k_2 .

If E^{SPD} is smaller than the forward price $\exp\{r(T-t)\}S_t$ of the stock, it suffices to move the mass β_1 further to the right. If E^{SPD} is too large, we move the mass β_p to the left. More precisely, setting $\tilde{k}_1 = k_1 - \mathbf{I}(E^{\text{SPD}} > \exp\{r(T-t)\}S_t)(E^{\text{SPD}} - \exp\{r(T-t)\}S_t)/\beta_p$, $\tilde{k}_p = k_p + \mathbf{I}(E^{\text{SPD}} < \exp\{r(T-t)\}S_t)(\exp\{r(T-t)\}S_t - E^{\text{SPD}})/\beta_1$,

we get

$$\exp\{r(T-t)\}S_t = \tilde{k}_1 \beta_p + \sum_{i=2}^{p-1} k_i \beta_{p-i+1} + \tilde{k}_p \beta_1.$$

This choice of \tilde{k}_1 and \tilde{k}_p guarantees that the expected value corresponding to the estimator β_1, \dots, β_p is equal to the forward price S_t of the stock; see the beginning of Section 6 for an application of this technique.

In Sections 4 and 5, we will concentrate on the properties of $\beta_2, \dots, \beta_{p-1}$ and further improvements in the estimation procedure.

4. Covariance structure

In this section, we use a model for the SPD development throughout the day to derive the covariance structure of the observed option prices depending on the strike prices and time of the trade. Considering the covariance structure in the estimation procedure solves the problems with heteroscedasticity and correlation of residuals that will be demonstrated in Section 5.3.

In this model, most recent option prices have the smallest variance and thus the largest weight in the estimation procedure. Similarly, the covariance of two option prices with the same strike price at approximately the same time is larger than the covariances of prices of some more dissimilar options.

We start by rewriting the model with iid error terms so that it can be more easily generalized. In Section 4.1, we present a model that accounts for heteroscedasticity and which is further developed in Sections 4.2 and 4.3, where an approximation of the covariance is calculated for any two options prices using only their strike prices and time of the trade. In Section 4.4, we suggest decomposing the error term into two parts, and we show how to estimate these additional parameters by the maximum likelihood method. The analysis of the resulting standardized residuals in Section 5.4 suggests that this covariance structure is applicable to our dataset.

Until now, we have assumed that the i -th option price (on a fixed day t) satisfies

$$C_i(k_j) = \Delta_j \tilde{\beta} + \varepsilon_i \tag{16}$$

or

$$C_i(k_j) = \Delta_j \tilde{\beta}_i + \varepsilon_i, \tag{17}$$

$$\tilde{\beta}_i = \tilde{\beta}_{i-1},$$

where ε_i are iid random errors with zero mean and constant variance σ^2 , $\tilde{\beta} = \tilde{\beta}_1 = \dots = \tilde{\beta}_i$ denotes the column vector of the unknown parameters, and Δ_j denotes the j -th row of the matrix Δ defined in (6), i.e.,

$$\Delta_j = (1, \Delta_p^j, \Delta_{p-1}^j, \dots, \Delta_{j+1}^j, \underbrace{0, \dots, 0}_{(j-1)}).$$

The residual analysis in Section 5.3 clearly demonstrates that the random errors ε_i are not independent and homoscedastic, and we have to consider some generalizations that lead to a better fit of the data set.

4.1. Heteroscedasticity

Assume that the i -th observation, corresponding to the j -th smallest exercise price k_j , can be written as

$$C_i(k_j) = \Delta_j \tilde{\beta}_i, \tag{18}$$

$$\tilde{\beta}_i = \tilde{\beta} + \varepsilon_i, \tag{19}$$

i.e., there are iid random vectors ε_i having iid components with zero mean and variances σ^2 in the state price density $\tilde{\beta}_i$. Clearly, the variance matrix of the vector of the observed option prices C is then

$$\text{Var } C = \sigma^2 \text{diag}(\mathcal{X}_\Delta \mathcal{X}_\Delta^T), \tag{20}$$

where \mathcal{X}_Δ is the design matrix in which each row of the matrix Δ is repeated n_j times, $j = 1, \dots, p$.

Remark 2. Assuming that the observed option prices have the covariance structure (20), the least squares estimates do not change, and

$$\text{Var } \hat{\beta} = \sigma^2 \{ \mathcal{X}_\Delta^\top \text{diag}(\mathcal{X}_\Delta \mathcal{X}_\Delta^\top)^{-1} \mathcal{X}_\Delta \}.$$

Another possible model for the heteroscedasticity would assume that the changes are multiplicative rather than additive.

$$C_i(k_j) = \Delta_j \tilde{\beta}_i$$

$$\log \tilde{\beta}_i = \log \tilde{\beta} + \varepsilon_i.$$

This model leads to a variance of $C_i(k_j)$ that depends on the value of the SPD:

$$\text{Var } C_i(k_j) = \sigma^2 \{ \beta_0^2 + (\Delta_p^j)^2 \beta_1^2 + (\Delta_{p-1}^j)^2 \beta_2^2 + (\Delta_{p-2}^j)^2 \beta_3^2 + \dots + (\Delta_{j+1}^j)^2 \beta_j^2 \}.$$

It is straightforward that Remark 2 also applies in this situation.

4.2. Covariance

Let us now assume that there are random changes in the state price density coefficients β_i over time so that we have

$$C_i(k_j) = \Delta_j \tilde{\beta}_i,$$

$$\tilde{\beta}_i = \tilde{\beta}_{i-1} + \varepsilon_i, \tag{21}$$

where, for fixed i , $\tilde{\beta}_i$ is the parameter vector and ε_k , $k = i, i-1, \dots$, are iid random vectors having iid components with zero mean and variances σ^2 . For nonequidistant time points, let δ_i denote the time between the i -th and $(i-1)$ -th observation. The model is

$$C_i(k_j) = \Delta_j \tilde{\beta}_i,$$

$$\tilde{\beta}_i = \tilde{\beta}_{i-1} + \delta_i^{1/2} \varepsilon_i, \tag{22}$$

and it leads to the covariance matrix with elements

$$\text{Cov}\{C_{i-u}(k_j), C_{i-v}(k_i)\} = \text{Cov}(\Delta_j \tilde{\beta}_{i-u}, \Delta_i \tilde{\beta}_{i-v})$$

$$= \sigma^2 \Delta_j \Delta_i^\top \sum_{l=1}^{\min(u,v)} \delta_{i+1-l}. \tag{23}$$

When we observe the i -th observation, we are usually interested in the estimation of the current value of the vector of parameters $\tilde{\beta}_i$.

4.3. Including put options

Similarly, we obtain the covariance for the price of the put options, $P_i(k_j)$. Using the relations between the α and β parameters, $\alpha_k = \beta_{p-k+1}$, for $k = 2, \dots, p-1$, and after some simplifications, we can write the model for the price of the put options, $P_i(k_j)$, as

$$P_i(k_j) = \Delta_j \tilde{\alpha}_i,$$

$$\tilde{\alpha}_i = \tilde{\alpha}_{i-1} + \delta_i^{1/2} \varepsilon_i, \tag{24}$$

where $\tilde{\alpha} = (\alpha_0, \alpha_1, \beta_{p-1}, \beta_{p-2}, \dots, \beta_2)^\top$ and Δ_j^p denotes the corresponding row of the design matrix, i.e.,

$$\Delta_j^p = (1, \Delta_j^1, \Delta_j^2, \dots, \Delta_j^{j-1}, \underbrace{0, \dots, 0}_{(p-j)}).$$

In this way, we obtain a joint estimation strategy for both the call and put option prices:

$$C_i(k_j) = \Delta_j \tilde{\beta}_i,$$

$$P_i(k_j) = \Delta_j^p \tilde{\alpha}_i,$$

$$\begin{pmatrix} \tilde{\beta}_i \\ \tilde{\alpha}_i \end{pmatrix} = \begin{pmatrix} \tilde{\beta}_{i-1} \\ \tilde{\alpha}_{i-1} \end{pmatrix} + \delta_i^{1/2} \varepsilon_i, \tag{25}$$

which directly leads to covariances

$$\text{Cov}\{P_{i-u}(k_j), P_{i-v}(k_i)\} = \text{Cov}(\Delta_j^p \tilde{\alpha}_{i-u}, \Delta_i^p \tilde{\alpha}_{i-v})$$

$$= \sigma^2 \Delta_j^p (\Delta_i^p)^\top \sum_{l=1}^{\min(u,v)} \delta_{i+1-l} \tag{26}$$

and

$$\text{Cov}\{C_{i-u}(k_j), P_{i-v}(k_i)\} = \text{Cov}(\Delta_j \tilde{\beta}_{i-u}, \Delta_i^p \tilde{\alpha}_{i-v})$$

$$= \sigma^2 \sum_{l=1}^{\min(u,v)} \delta_{i+1-l} \sum_{k=2}^{p-1} \Delta_{p+1-k}^j \Delta_i^{p+1-k}. \tag{27}$$

Together with (23), Eq. (26) and (27) allow us to calculate the covariance matrix of all observed option prices using only their strike prices and the times between the transactions.

4.4. Error term for option prices

Using the model (25) would mean that all changes observed in the option prices are due only to changes in the SPD. It seems natural to add another error term, η_i , as a description of the error in the option price:

$$C_i(k_j) = \Delta_j \tilde{\beta}_i + \eta_i,$$

$$P_i(k_j) = \Delta_j^p \tilde{\alpha}_i + \eta_i,$$

$$\begin{pmatrix} \tilde{\beta}_i \\ \tilde{\alpha}_i \end{pmatrix} = \begin{pmatrix} \tilde{\beta}_{i-1} \\ \tilde{\alpha}_{i-1} \end{pmatrix} + \delta_i^{1/2} \varepsilon_i, \tag{28}$$

where $\eta_i \sim N(0, v^2)$ are iid random variables independent of the random vectors ε_i . Here, normality assumptions are added both for η_i and ε_i so that the variance components parameters v^2 and σ^2 may be estimated by the maximum likelihood method.

Next, in order to simplify the notation, let us fix the index i , and let Y denote the vector of observed call and put option prices, \mathcal{X}_Δ the corresponding design matrix consisting of the corresponding rows Δ_j and Δ_j^p , and $\tilde{\gamma}$ the combined vector of unknown parameters. Denoting by Σ_i the matrix containing the covariances defined in (23), (26) and (27), we can rewrite model (25) as

$$Y = \mathcal{X}_\Delta \tilde{\gamma} + \xi, \tag{29}$$

where $\text{Var } \xi = \text{Var } Y = \sigma^2 \Sigma_i + v^2 \mathbf{I}_n = \sigma^2 (\Sigma_i + \psi^2 \mathbf{I}_n) = \sigma^2 V$, where $\psi^2 = v^2 / \sigma^2$. Differentiating the log-likelihood

$$l(\beta, \sigma^2, \psi^2) = -\frac{n}{2} \log(2\pi) - \frac{1}{2} \log |\sigma^2 V|$$

$$- \frac{1}{2\sigma^2} (Y - \mathcal{X}_\Delta \tilde{\gamma})^\top V^{-1} (Y - \mathcal{X}_\Delta \tilde{\gamma}),$$

we obtain

$$\frac{\partial l(\beta, \sigma^2, \psi^2)}{\partial \psi^2}$$

$$= -\frac{1}{2} \text{tr}(V^{-1}) + \frac{1}{2\sigma^2} (Y - \mathcal{X}_\Delta \tilde{\gamma})^\top V^{-2} (Y - \mathcal{X}_\Delta \tilde{\gamma}). \tag{30}$$

For any fixed value of the parameter ψ^2 , it is straightforward to calculate the optimal σ^2 and $\tilde{\gamma}$. Hence, the numerical maximization of the log-likelihood can be based on a search for a root (zero) of the one-dimensional function (30).

Moreover, the variance components parameters σ^2 and $v^2 = \psi^2 \sigma^2$ have a very natural econometric interpretation: σ^2 describes the speed of change of the SPD and v^2 the error in observed option prices.

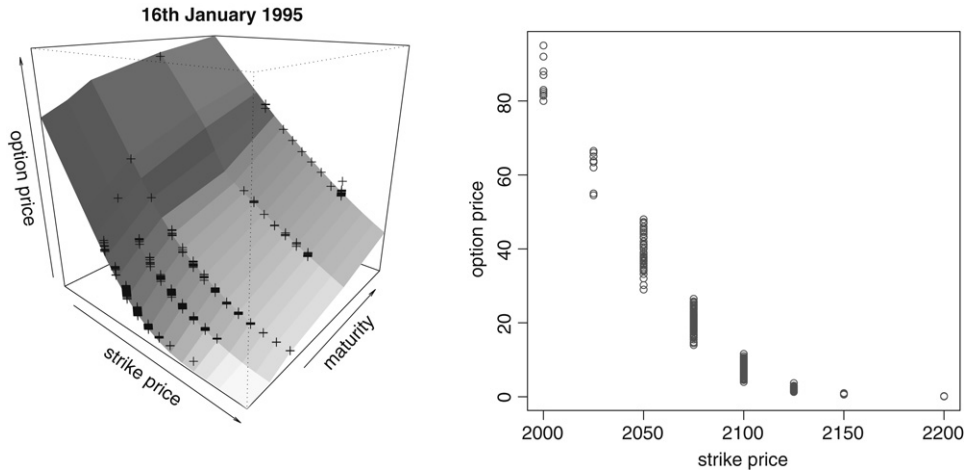


Fig. 3. Option prices plotted against strike price and time to maturity with a two-dimensional kernel regression surface (left) in January 1995 and the ensemble of the call option prices with shortest time to expiry against strike price (right) on 16 January 1995. SFB and CASE data base: sfb649.wiwi.hu-berlin.de.




Fig. 4. On 16 January 1995, the unconstrained estimate satisfies the constraints. Hence, it is equal to the constrained estimate. The top panel shows the original data with the fitted call pricing functions. The second and the third panels show the estimates of the first and second derivatives, respectively.

5. Application to DAX data

We analyze a data set containing observed option prices for various strike prices and maturities. Other variables are the interest rate, date, and time. In 1995, one observed every day about 500 trades; in today's more liquid option markets this number has increased approximately 10 times. In our empirical study we will consider the time period from 1995 to 2003, thus also covering more recent liquid option market.

Fig. 3 displays the observed prices of European call options written on the DAX for the 16 January 1995. The left panel shows the ensemble of call option prices for different strikes and maturities as a free structure together with a smooth surface. The typical shape of dependency of the option price on the strike price can be observed in the right panel, containing the option prices only for the shortest time to expiry, $\tau = T - t = 4$ days.

In order to illustrate the method, we apply it to DAX option prices on two consecutive days. These days (16 and 17 January

1995) were selected since they provide a nice insight into the behavior of the presented methods.

5.1. Estimator with iid random errors

We start by a comparison of the unconstrained and constrained estimator described respectively in Sections 2.3 and 3.1.

For the European call option prices displayed in the right-hand plot in Fig. 3, we obtain the estimates plotted in Fig. 4. The top plot displays the original data, the second plot shows the estimate of the first derivative, and the third plot shows the estimate of the second derivative, i.e., the state price density. Actually, all plots contain two curves, both obtained using model (8). The thick line is calculated using the parameters β_i without constraints, whereas the thin line uses the reparameterization $\beta_i(\xi)$ given in Section 3.1. In Fig. 4, these two estimates coincide since the model maximizing the likelihood without constraints, by chance, fulfills the constraints ($\exists \xi : \beta_i = \beta_i(\xi), i = 0, \dots, p - 1$), and hence it is clear that the same parameters also maximize the constrained likelihood.

The situation, in which the call pricing functions fitted with and without constraints differ, is displayed in Fig. 5. Notice that the difference between the two regression curves is small, whereas the difference between the estimates of the state price density (i.e., the second derivative of the curve) is surprisingly large. The unconstrained estimate shows very unstable behavior on the left-hand side of the plot. The constrained version behaves more reasonably. Very small differences between the fitted call pricing functions in the top plot in Fig. 5 lead to huge differences in the estimates of the second derivative.

We therefore conclude that a small error in the estimate of the call pricing function may lead to large scale error in the estimates of the first and second derivatives. The scale of this type of error seems to be limited by imposing the shape constraints given in Section 2.2.

5.2. Confidence intervals

In Figs. 6 and 7, we plot both estimates together with the 95% confidence intervals. Notice that, in the unconstrained model, the estimates of the values of the SPD are just the parameters of the linear regression model. Hence, the confidence intervals for the parameters are, at the same time, also confidence intervals for the SPD. These confidence intervals for 16 and 17 January 1995 are displayed in the upper plots in Figs. 6 and 7. The drawbacks of this method are clearly visible. In Fig. 6, the lower bounds of the confidence intervals only asymptotically satisfy the condition of




Fig. 5. On 17 January 1995, the unconstrained estimate, displayed using the thin line, does not satisfy the constraints. The top panel shows the original data with the two fitted call pricing functions. The estimates of the first derivative in the second panel look rather different. The constrained estimate of the second derivative in the bottom panel is clearly much more stable than the unconstrained estimate.




Fig. 6. The unconstrained and constrained confidence intervals for the SPD on 16 January 1995. The description on the x-axis shows the number of observations in each point.

positivity. In Fig. 7, we observe large variability on the left-hand side of the plot (the region with low number of observations). Again, some of the lower bounds are not positive. Clearly, the confidence intervals based on the unconstrained model make sense only if the constraints are, by chance, satisfied. Even if this is the case, there is no guarantee that the lower bounds will be positive. The lower panels in Figs. 6 and 7 display the nonnegative asymptotic confidence intervals calculated according to Section 3.4.

In Fig. 6, both types of confidence interval provide very similar results. The only difference is at the minimum and maximum value




Fig. 7. Confidence intervals for SPD on 17 January 1995. The description on the x-axis shows the number of observations in each point.




Fig. 8. The time dependency and the heteroscedasticity of the residuals during one day. The circle, square, and star denote the trades carried out in the morning, midday, and afternoon, respectively. The size of the symbols denotes the number of residuals.

of the independent variable (strike price), where the unconstrained method provides negative lower bounds and the conditional method leads to very large upper bounds of the confidence intervals.

In Fig. 7, we plot the confidence intervals for 17 January 1995. In the central region of the graphics, both types of confidence interval are quite similar. On the left-hand and right-hand sides, both methods tend to provide confidence intervals that seem to be overly wide. For the constrained method, we observe that the length of the confidence intervals explodes when the estimated value of the SPD is very close to zero and, at the same time, the number of observation in that region (see the description of the horizontal axis) is small.

5.3. Residual analysis

The residuals on 17 January 1995 are plotted in Fig. 8. The time of trade (in hours) is denoted by the plotting symbol. The circle, square, and star denote the trades carried out in the morning, midday, and afternoon, respectively. The size of the symbols





Fig. 9. Estimate using the covariance structure (28) on 17 January 1995. The upper plot shows the observed option prices and the constrained estimate. The size of the plotting symbols corresponds to the weight of the observations. The lower plot shows the estimated SPD with confidence intervals.

Fig. 10. The development of the standardized residuals resulting from the model with the covariance structure (28) on 17 January 1995 during the day, where circles, squares, and stars denote the residuals from morning, midday, and afternoon, and a histogram of the standardized residuals.

corresponds to the number of residuals lying in the respective areas.

The majority of the residuals correspond to the strike prices of 2075DEM and 2100DEM. The variance of the residuals is very low on the right-hand side of the plot and it rapidly increases when moving towards smaller strike prices. On the left-hand side of the plot, for strike prices smaller than 2000, we have only very few observations, and cannot judge the residual variability reliably.

Apart from the obvious heteroscedasticity we also observe a very strong systematic movement in the SPD throughout the day: the circles, corresponding to the first third of the day, are positive, and all stars, denoting the afternoon residuals, are negative. Similar patterns can be observed every day—residuals corresponding to the same time have the same sign.

We conclude that the assumption of iid random errors is obviously not fulfilled as the option prices tend to follow the changes of the market during the day.

5.4. Application of the covariance structure

In Fig. 9, we present the estimator combining both put and call option prices and using the covariance structure proposed in Section 4.4. In comparison with the results plotted in Fig. 7, we observe shorter length of the confidence intervals.

The estimates of the variance components parameters are $\hat{\psi}^2 = 17.77$, $\hat{\sigma}^2 = 0.0041$, and $\hat{\nu}^2 = 0.0722$. For interpretation, it is more natural to consider $\hat{\nu} = 0.2687$, suggesting that 95% of the option prices were on 17 January 1995 not further than 0.5DEM from the correct option price implied by the current (unobserved) SPD.




Fig. 11. SPD estimate on 17 January 1995 with prediction intervals for the next 5 h calculated for every 30 min.

The standardized residuals in the top panel of Fig. 10 were plotted using the same technique as the residuals in Fig. 8. Whereas the residuals for the iid model showed strong correlations and heteroscedasticity, the structure of the standardized residuals looks much better. It is natural that the residuals are larger in the central part since more than 90% of observations have strike price




Fig. 12. Daily development of the expected value of the uncorrected SPD from January to March 1995. The circles denote the corresponding closing value of the DAX.




Fig. 13. Daily development of the SPD variance from January to March 1995.

between 2050 and 2100. The largest residuals were omitted in the residual plot so that the structure in the central part is more visible, but the lower panel of Fig. 10 displays the histogram of all residuals. The distribution of the residuals seems to be symmetric, and its shape is not too far from Normal distribution. However, the kurtosis of this distribution is too large, and formal tests reject normality.

In Fig. 11, we plot prediction intervals for the SPD obtained only by recalculating the covariance structure (28) with respect to some future time. More precisely, the prediction intervals are obtained from option prices observed until \hat{i} . Then, using the notation of Section 4.4, we have, for the future $\tilde{\beta}_{i+1}$ and $\tilde{\alpha}_{i+1}$,

$$\begin{aligned} C_i(k_j) &= \Delta_j \tilde{\beta}_i + \eta_i, \\ P_i(k_j) &= \Delta_j^p \tilde{\alpha}_i + \eta_i, \\ \begin{pmatrix} \tilde{\beta}_{i+1} \\ \tilde{\alpha}_{i+1} \end{pmatrix} &= \begin{pmatrix} \tilde{\beta}_i \\ \tilde{\alpha}_i \end{pmatrix} + \delta_{i+1}^{1/2} \varepsilon_{i+1}. \end{aligned} \tag{31}$$

It is now easy to see that the only modification that has to be done for estimating $\tilde{\beta}_{i+1}$ is to add the length of the forecasting horizon δ_{i+1} to the sum in (23), (26) and (27), and to recalculate the confidence regions using this variance matrix with the same estimates of the variance parameters σ^2 and ν^2 . In Fig. 11, the 95% confidence intervals for the true SPD are denoted by the black

dashed line. The grey dashed lines denote the prediction intervals calculated for each 30 min for the next 5 h. In this way, we can obtain a simple approximation for future short-term fluctuations of the SPD. In the long run, the prediction intervals become too wide to be informative.

6. Dynamics of the SPD

In order to study the dynamics of SPDs, we calculated the basic moment characteristics of the estimated SPDs. Note that the estimator does not allow one to estimate the SPD in the tails of the distribution. We can only estimate the probability mass lying to the left ($1 - \sum_{i=1}^{p-1} \beta_i$) and to the right (β_1) of the available strike price range. Hence, the moments calculated in this section are only approximations which cannot be calculated more precisely without additional assumptions, for example, on the tail behavior or parametric shape of the SPD.

The estimated mean and variance in the first quarter of 1995 are plotted as lines in Figs. 12–13. Note that the SPDs in this period were always estimated using the options with shortest time to maturity. This means that the time to maturity is decreasing linearly in both plots, but it jumps up whenever the option with the shortest time to maturity expires. These jumps occurred at days 16, 36, and 56.

From no-arbitrage considerations, it follows that the mean of the SPD should correspond to the value of the DAX,

$$\widehat{E}^{SPD} = \int S_T f(S_T) dS_T = \exp\{r(T - t)\} S_t.$$

See also the discussion in Section 3.6. In Fig. 12, the observed values of the DAX multiplied by the factor $\exp\{r(T - t)\}$ are plotted as circles for the first 65 trading days in 1995, and we observe that the estimated means of the SPD estimates, displayed as the line, follow the theoretical value very closely. A small difference is mainly due to the fact that, in 1995, the observed strike prices do not entirely cover the support of the SPD. For example, on day 16, the difference between the SPD mean (2018.7) and the DAX multiplied by the discount factor (2012.1) is equal to 6.6. The fact that there are not any trades for strike prices smaller than 1925 means that we only know that the probability mass lying to the left from 1950 is equal to 0.25. In the calculation of the estimate of the SPD mean plotted in Fig. 12, this probability mass is assigned to the value 1925, as this is the leftmost observed strike price. Obviously, assigning this probability mass rather to the value $1925 - (6.6/0.25) = 1898.6$ leads a more realistic estimate of the SPD and to the equality of the SPD mean and the discounted DAX.

In Fig. 13, we see that the variance of the SPD decreases linearly as the option moves closer to its maturity. This observation suggests that SPD estimates calculated for neighboring maturities can be linearly interpolated in order to obtain an SPD estimate with arbitrary time to maturity. Such an estimate is important for making the SPD estimates comparable and for studying the development of the market expectations.

6.1. Estimate with the fixed time to expiry

The variances displayed in Fig. 13 suggest that the variance of the SPD estimates changes approximately linearly in time when moving closer to the date of expiry.

Hence, from the estimates $f_{\tau_1}(\cdot)$ and $f_{\tau_2}(\cdot)$ of centered SPDs corresponding to the times of expiry $\tau_1 < \tau_2$, we construct an estimate $f_\tau(\cdot)$ for any time of expiry $\tau \in (\tau_1, \tau_2)$ as

$$f_\tau(\cdot) = \frac{(\tau_2 - \tau)f_{\tau_1}(\cdot) + (\tau - \tau_1)f_{\tau_2}(\cdot)}{\tau_2 - \tau_1}. \tag{32}$$




Fig. 14. Prediction intervals for the DAX based on SPDs and historical simulation from January 1995 to March 2003.

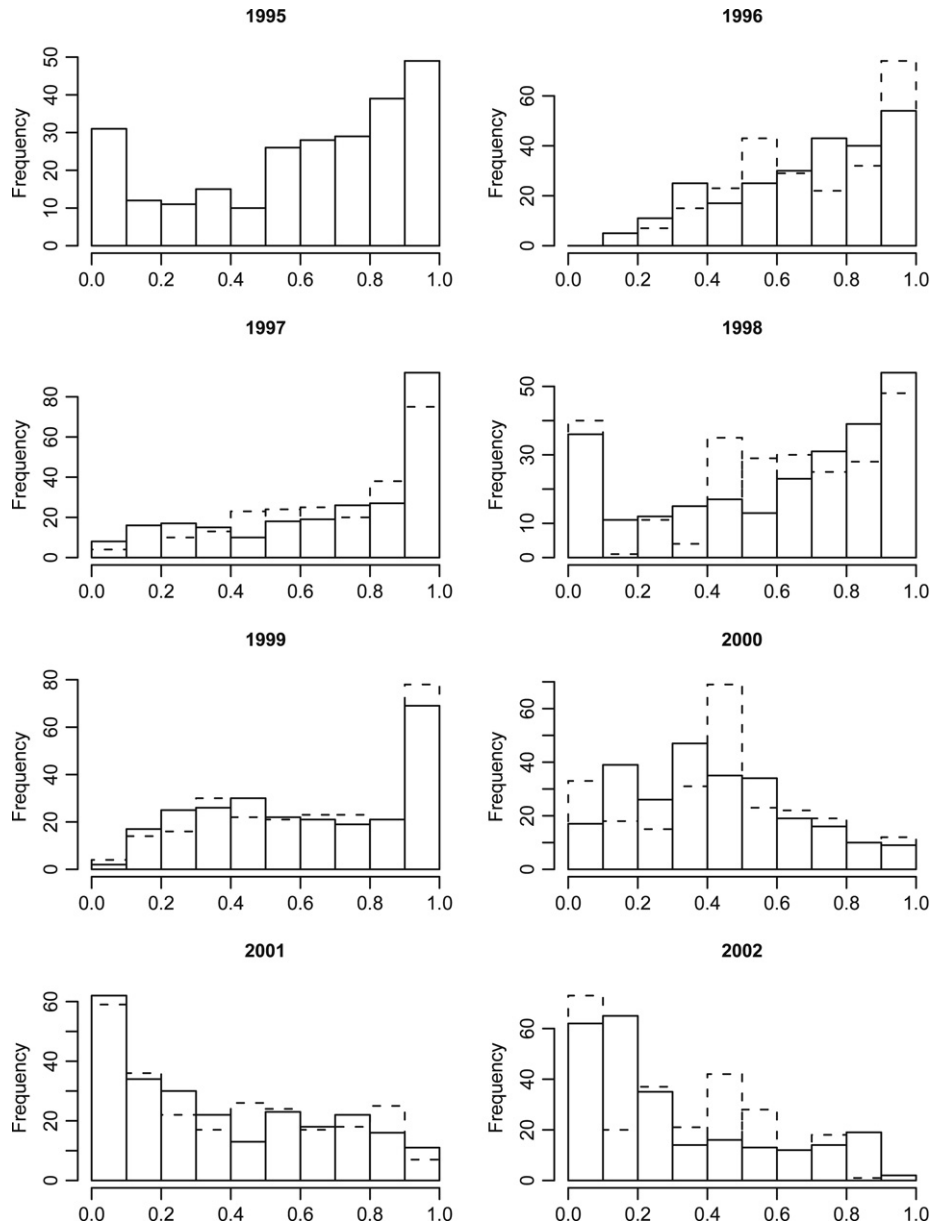


Fig. 15. Histograms for the SPDs (full line) and historical simulation (dashed line).




Fig. 16. Integral transformation for estimated SPDs.

In this way, the variance, V_τ , of the centered SPD with time to expiry equal to τ can be expressed as

$$\begin{aligned} V_\tau &= \int x^2 f_\tau(x) dx \\ &= \int x^2 \frac{(\tau_2 - \tau) f_{\tau_1}(x) + (\tau - \tau_1) f_{\tau_2}(x)}{\tau_2 - \tau_1} dx \\ &= \frac{(\tau_2 - \tau) V_{\tau_1} + (\tau - \tau_1) V_{\tau_2}}{\tau_2 - \tau_1}. \end{aligned}$$

We argue that such an estimate is reasonable since we observed in Fig. 13 that the SPD variances change linearly in time.

6.2. Verification of the market's expectations

Under the risk neutral (equivalent martingale) measure, the SPD reflects the market's expectation of the behavior of the value of the DAX in 45 days. Hence, it is interesting to use our data set to verify how these expectations compare with reality. In the left plot in Fig. 14, we plot intervals based on the SPD together with the true future value of the DAX: the black lines display the 2.5% and

97.5% quantiles of the estimated SPD; the future value of the DAX is displayed as a grey line. In the right plot, we show in the same way the 45-day ahead predictions based on the historical distribution of the 45-day absolute returns in the last 100 trading days; the 2.5% and 97.5% quantiles of this distribution are plotted as black lines.

Fig. 14 suggests that the method works well and that the DAX mostly stays well within the quantiles calculated from the estimated SPDs. The DAX was sometimes rising faster than the market expected from 1995 to mid-1998. After a fast decrease in the second half of 1998, the market increased again till the beginning of year 2000. Since then, the market has decreased. However, the changes stay mostly within or very close to the bounds predicted by our SPD estimates. The only exception is the large shock observed in September 2001, caused by the terrorist attack on the World Trade Center.

The upper quantiles, 97.5%, of the historical distribution of the 45-day absolute returns mostly agree with the upper quantiles of the SPD. The lower quantiles, 2.5%, of the SPDs seem to be much more variable than the same quantiles of the historical distribution. Both the lower and the upper quantiles of the historical distribution lie mostly above the corresponding quantiles of the estimated SPD, respectively in 69.44% and 81.75%.

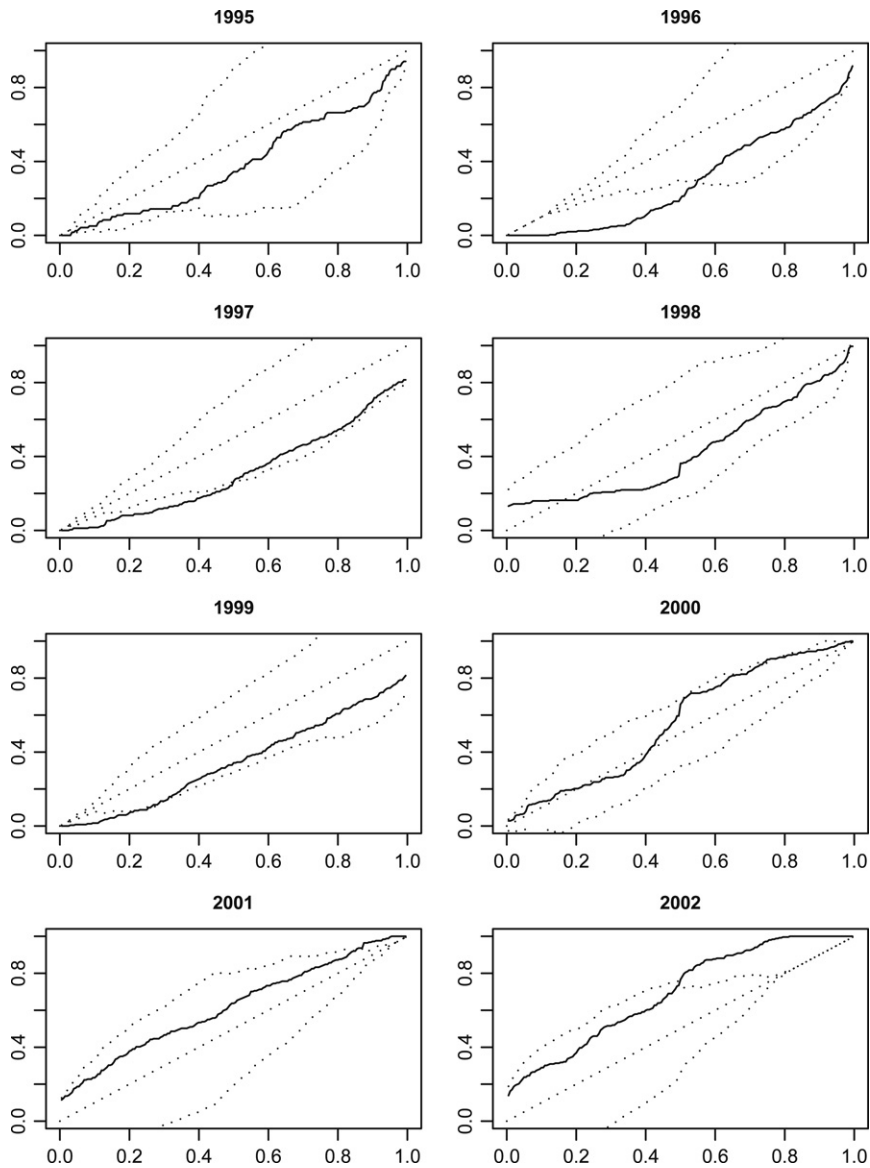


Fig. 17. Integral transformation for historical simulation.

Table 2
 Fraction of the year that the DAX stays in the prediction corridor.

Year	1996	1997	1998	1999	2000	2001	2002
SPD (%)	84.40	66.13	75.30	74.60	97.22	85.66	94.84
Historical (%)	82.00	79.44	76.89	77.38	93.25	86.06	80.56

This observation just confirms the fact that the observed SPD includes effects of risk aversion.

In Table 2, we show the fraction of the year that the DAX stays in the prediction corridor. This suggests that the coverage is slightly better for the historical simulation if the DAX is increasing and better for the SPD based prediction if the DAX is decreasing (years 2000 and 2002).

6.3. Evaluation of the quality of the forecasts

The quality of the forecasts can be evaluated by comparing the true future observation with its predicted distribution (the SPD). Diebold et al. (1998) propose to evaluate density forecasts using the probability integral transformed observations $z_{h,t}$, where t denotes the time and h the forecasting horizon. More precisely,

we define

$$z_{h,t} = \int_{-\infty}^{X_{t+h}} \hat{f}_{h,t}(u) du,$$

where $\hat{f}_{h,t}(\cdot)$ denotes our estimate of the SPD h days ahead at time t and X_{t+h} is the future observation. In other words, $z_{h,t}$ is the probability value of X_{t+h} with respect to $\hat{f}_{h,t}(\cdot)$. Clearly, the $z_{h,t}$ should be uniformly $U(0, 1)$ distributed if the estimated SPD $\hat{f}_{h,t}(\cdot)$ is equal to the true density of X_{t+h} . In Fig. 15, we display the histograms of $z_{h,t}$'s for each year for the estimated SPDs and historical simulation using full and dashed histograms, respectively. Clearly, in the ideal case, the histograms should not be too far from a Uniform $U(0, 1)$ distribution. In our data, for the prediction horizon $h = 45$ days, we observe that the histograms look quite different from what we would expect. Especially in

years 1995–1999, the DAX was moving mainly in the upper quantiles of the predicted SPD. The forecasts based on the historical distribution of the 45-day returns behave similarly.

In order to account for the overlapping forecasting periods, we calculate the confidence limits for the empirical distribution function

$$\widehat{F}(u) = \frac{1}{T} \sum_{t=1}^T \mathbf{I}(z_{h,t} \leq u)$$

of $z_{h,t}$'s that take into account the autocorrelation structure.

$$\widehat{\text{Var}}\{\widehat{F}(u)\} = \frac{1}{T} \left\{ \widehat{\gamma}_u(0) + 2 \sum_{j=1}^h \left(1 - \frac{j}{T}\right) \widehat{\gamma}_u(j) \right\}, \quad (33)$$

where $\gamma_u(j)$ is the sample autocovariance of order j :

$$\gamma_u(j) = \frac{1}{T} \sum_{t=j+1}^T \{ \mathbf{I}(z_{h,t} \leq u) - \widehat{F}(u) \} \{ \mathbf{I}(z_{h,t-j} \leq u) - \widehat{F}(u) \}.$$

The empirical distribution functions $\widehat{F}(\cdot)$ are plotted separately for years 1995–2002 in Fig. 16. The distribution function of $U(0, 1)$ and the limits following from (33) are displayed as dotted lines. The year 2003 was not included since our dataset contains only two months of the year 2003, which did not leave enough observations to confirm the forecasts.

In 1996 and 1997, the market was growing much faster than the SPDs were indicating. In 1996, it never happened that the DAX fell below the 10% quantile of the SPD, and there were only a few days when this value was below 20%. The situation in 1998 and 1999 was less extreme even though the fast growth of the DAX continued. The distribution given by the SPD estimate $\widehat{f}_{t,h}(\cdot)$ for the horizon $h = 45$ days does not differ significantly from the true distribution of X_{t+h} in 2000–2001, but in 2002 we again observe significant differences. Thus, the DAX was growing faster than the option market expected in 1996, 1997, and 1999 and it was falling faster in 2002.

Fig. 17 shows the same graphics for the forecast based on the historical distribution of the returns. The deviations are more clearly visible but the overall picture is very similar; the only difference arises in 2001 when the predictions did not stay between the limits.

7. Conclusion

We have proposed a simple nonparametric model for arbitrage-free estimation of the SPD. Our procedure takes care of the daily changing covariance structure and involves both types of European option. Moreover, the covariance structure allows us to calculate prediction intervals capturing future behavior of the SPD. We analyze the moment dynamics of the SPD from 1995–2003. An application to DAX EUREX data for the years 1995–2003 produces a corridor that is compared to the future DAX index value. The proposed technique enables us not only to price exotic options but also to measure the risk and volatility ahead of us.

Acknowledgments

We thank Volker Krätschmer for useful comments concerning the existence and uniqueness of the constrained regression function and the anonymous referee for many insightful comments leading to substantial improvements in both the presentation and the content of the paper. The research was supported by Deutsche Forschungsgemeinschaft, SFB 649 “Ökonomisches Risiko”, by MSM0021620839, GAČR GA201/08/0486, and by MŠMT 1K04018.

References

- Ait-Sahalia, Y., Duarte, J., 2003. Nonparametric option pricing under shape restrictions. *Journal of Econometrics* 116, 9–47.
- Ait-Sahalia, Y., Lo, A.W., 1998. Nonparametric estimation of state-price densities implicit in financial asset prices. *Journal of Finance* 53, 499–547.
- Ait-Sahalia, Y., Lo, A.W., 2000. Nonparametric risk management and implied risk aversion. *Journal of Econometrics* 94, 9–51.
- Ait-Sahalia, Y., Wang, Y., Yared, F., 2000. Do option markets correctly price the probabilities of movement of the underlying asset? *Journal of Econometrics* 102, 67–110.
- Anděl, J., 1985. *Mathematical Statistics*. SNTL/Alfa, Prague (in Czech).
- Andersen, L.B.G., Brotherton-Ratcliffe, R., 1997. The equity option volatility smile: An implicit finite-difference approach. *Journal of Computational Finance* 1 (2), 5–37.
- Bondarenko, O., 2003. Estimation of risk-neutral densities using positive convolution approximation. *Journal of Econometrics* 116, 85–112.
- Breeden, D., Litzenberger, R., 1978. Prices of state-contingent claims implicit in option prices. *Journal of Business* 51, 621–651.
- Buehler, H., 2006. Expensive martingales. *Quantitative Finance* 6 (3), 207–218.
- Diebold, F.X., Gunther, T., Tay, A., 1998. Evaluating density forecasts, with applications to financial risk management. *International Economic Review* 39, 863–883.
- Dupire, B., 1994. Pricing with a smile. *RISK* 7 (1), 18–20.
- Fengler, M.R., 2005. *Semiparametric Modeling of Implied Volatility*. Springer, Heidelberg.
- Fengler, M.R., Härdle, W., Mammen, E., 2007. A dynamic semiparametric factor model for implied volatility string dynamics. *Journal of Financial Econometrics* 5 (2), 189–218.
- Hafner, R., Wallmeier, M., 2000. *The Dynamics of DAX Implied Volatilities*. University of Augsburg Working Paper. Available at SSRN: <http://ssrn.com/abstract=234829> or doi: 10.2139/ssrn.234829.
- Harrison, J., Pliska, S., 1981. Martingale and stochastic integral in the theory of continuous trading. *Stochastic Processes and their Applications* 11, 215–260.
- Hlávka, Z., Svojk, M., 2008. Application of extended Kalman filter to SPD estimation. In: Härdle, W., Hautsch, N., Overbeck, L. (Eds.), *Applied Quantitative Finance*. Springer, Berlin, pp. 233–247.
- Huynh, K., Kervella, P., Zheng, J., 2002. Estimating state-price densities with nonparametric regression. In: Härdle, W., Kleinow, T., Stahl, G. (Eds.), *Applied Quantitative Finance*. Springer, Heidelberg, pp. 171–196.
- Jackwerth, J.C., 1999. Option-implied risk-neutral distributions and implied binomial trees: A literature review. *Journal of Derivatives* 7, 66–82.
- Kahalé, N., 2004. An arbitrage-free interpolation of volatilities. *RISK* 17 (5), 102–106.
- Rao, C.R., 1973. *Linear Statistical Inference and Its Applications*. Wiley, New York.
- Renault, E., 1997. Econometric models of option pricing errors. In: Kreps, D.M., Wallis, K.F. (Eds.), *Advances in Economics and Econometrics: Theory and Applications*, Seventh World Congress, vol. III. Cambridge University Press, Cambridge, pp. 223–278.
- Robertson, T., Wright, F.T., Dykstra, R.L., 1988. *Order Restricted Statistical Inference*. Wiley, Chichester.
- Seber, G.A.F., Wild, C.J., 2003. *Nonlinear Regression*. Wiley, Hoboken, New Jersey.
- Serfling, R., 1980. *Approximation Theorems of Mathematical Statistics*. Wiley, New York.
- Stoll, H.R., 1969. The relationship between put and call option prices. *Journal of Finance* 24, 801–824.
- Yatchew, A., Härdle, W., 2006. Nonparametric state price density estimation using constrained least squares and the bootstrap. *Journal of Econometrics* 133 (2), 579–599.

Variable Selection and Oversampling in the Use of Smooth Support Vector Machines for Predicting the Default Risk of Companies

WOLFGANG HÄRDLE,¹ YUH-JYE LEE,²
DOROTHEA SCHÄFER^{3*} AND YI-REN YEH²

¹ CASE, Humboldt University, Berlin, Germany

² Department of Computer Science Information Engineering,
National Taiwan University of Science and Technology, Taipei,
Taiwan

³ German Institute of Economic Research, Berlin, Germany

ABSTRACT

In the era of Basel II a powerful tool for bankruptcy prognosis is vital for banks. The tool must be precise but also easily adaptable to the bank's objectives regarding the relation of false acceptances (Type I error) and false rejections (Type II error). We explore the suitability of smooth support vector machines (SSVM), and investigate how important factors such as the selection of appropriate accounting ratios (predictors), length of training period and structure of the training sample influence the precision of prediction. Moreover, we show that oversampling can be employed to control the trade-off between error types, and we compare SSVM with both logistic and discriminant analysis. Finally, we illustrate graphically how different models can be used jointly to support the decision-making process of loan officers. Copyright © 2008 John Wiley & Sons, Ltd.

KEY WORDS insolvency prognosis; support vector machines; statistical learning theory; non-parametric classification

INTRODUCTION

Default prediction is at the core of credit risk management and has therefore always attracted special attention. It has become even more important since the Basel Committee on Banking Supervision (Basel II) established borrowers' rating as the crucial criterion for minimum capital requirements of banks. The methods for generating rating figures have developed significantly over the last 10 years (Krahn and Weber, 2001). The rationale behind the increased sophistication in predicting borrowers' default risk is the aim of banks to minimize their cost of capital and to mitigate their own bankruptcy risks.

*Correspondence to: Dorothea Schäfer, German Institute for Economic Research (DIW) Berlin, Mohrenstrasse 58, 10117 Berlin, Germany. E-mail: dschaefer@diw.de

In this paper we intend to contribute to the increasing sophistication by exploring the predicting power of smooth support vector machines (SSVM). SSVM are a variant of the conventional support vector machines (SVM). The working principle of SVM in general can be described very easily. Imagine a group of observations in distinct classes such as balance sheet data from solvent and insolvent companies. Assume that the observations are such that they cannot be separated by a linear function. Rather than fitting nonlinear curves to the data, SVM handle this problem by using a specific transformation function—the kernel function—that maps the data from the original space into a higher-dimensional space where a hyperplane can do the separation linearly. The constrained optimization calculus of SVM gives a unique optimal separating hyperplane and adjusts it in such a way that the elements of distinct classes possess the largest distance to the hyperplane. By re-transforming the separating hyperplane into the original space of variables, the typical nonlinear separating function emerges (Vapnik, 1995). The main difference between SSVM and SVM is the following: the SSVM technique formulates the problem as an unconstrained minimization problem. This formulation has mathematical properties such as strong convexity and desirable infinite differentiability.

Our aim is threefold when using SSVM. Firstly, we examine the power of the SSVM in predicting company defaults; secondly, we investigate how important factors that are exogenous to the model, such as selecting the appropriate set of accounting ratios, length of training period and structure of the training sample, influence the precision; and thirdly, we explore how oversampling and downsampling affect the trade-off between Type I and Type II errors. In addition, we illustrate graphically how loan officers can benefit from jointly considering the prediction results of different SSVM variants and different models.

There are basically three distinct approaches in predicting the risk of default: option theory-based approaches, parametric models and non-parametric methods. While the first class relies on the rule of no arbitrage, the latter two are based purely on statistic principles. The popular (Merton, 1974) model treats the company's equity as the underlying asset of a call option held by shareholders. In case of insolvency shareholders deny exercising. The probability of default is derived from an adapted Black–Scholes formula. Later, several authors (e.g., Longstaff and Schwartz, 1995; Mella-Barral and Perraudin, 1997; Leland and Toft, 1996; Zhou, 2001; to name only a few) proposed variations to ease the strict assumptions on the structure of the data imposed by the Merton model. These approaches are frequently denoted as structural models. However, the most challenging requirement is the knowledge of market values of debt and equity. This precondition is a severe obstacle to using the Merton model adequately as it is only satisfied in a minority of cases.

Parametric statistical models can be applied to any type of data, whether they are market based or book based. The first model introduced was discriminant analysis (DA) for univariate (Beaver, 1966) and multivariate models (Altman, 1968). After DA usage of the logit and probit approach for predicting default was proposed in Martin (1977) and Ohlson (1980). These approaches rely on the a priori assumed functional dependence between risk of default and predictor. DA requires a linear functional dependence, or a pre-shaped polynomial functional dependence in advanced versions. Logit and probit tools work with monotonic relationships between default event and predictors such as accounting ratios. However, such restrictions often fail to meet the reality of observed data. This fact makes it clear that there is a need for an approach that, in contrast to conventional methods, relaxes the requirements on data and/or lowers the dependence on heuristics. Semi-parametric models as in Hwang *et al.* (2007) are between conventional linear models and non-parametric approaches. Nonlinear classification methods such as support vector machines (SVM) or neural networks are even stronger candidates to meet these demands as they go beyond conventional

discrimination methods. Tam and Kiang (1992) and Altman *et al.* (1994) focus on neural networks. In contrast, we concentrate on SVM exclusively.

The SVM method is a relatively new technique and builds on the principles of statistical learning theory. It is easier to handle compared to neural networks. Furthermore, SVM have a wider scope of application as the class of SVM models includes neural networks (Schölkopf and Smola, 2002). The power of SVM technology becomes evident in a situation as depicted in Figure 1 where operating profit margin and equity ratio are used as explanatory variables. A separating function similar to a parabola (in black) appears in the two-dimensional space. The accompanying light-grey lines represent the margin boundaries whose shape and location determine the distance of elements from the separating function. In contrast, the logit approach and discriminant DA yield the (white) linear separating function (Härdle *et al.*, 2007a).

Selecting the best accounting ratios for executing the task of predicting is an important issue in practice but has not received appropriate attention in research. We address this issue of how important the chosen set of predictors is for the outcome. For this purpose we explore the prediction potential of SSVM within a two-step approach. First, we derive alternative sets of accounting ratios that are used as predictors. The benchmark set comes from Chen *et al.* (2006). A second set is defined by a 1-norm SVM, and the third set is based on the principle of adding only those variables that contain the most contrary information with respect to an initial set that is a priori chosen. We call the latter procedure the incremental forward selection of variables. As a result we are working with three variants of SSVM. In the second step, these variants are compared with respect to their prediction power. We also compare SSVM with two traditional methods: the logit model and linear discriminant analysis.

The analysis is built on 28 accounting ratios of 20,000 solvent and 1000 insolvent German companies. Our findings show that the different SSVM types have an overall good performance with the means of correct predictions ranging from 70% to 78%. The SSVM on the basis of incremental




Figure 1. SVM-separating function (black) with margin in a two-dimensional space

forward selection clearly outperform the SSVM based on predictors selected by the 1-norm SVM. It is also found that oversampling influences the trade-off between Type I and Type II errors. Thus, oversampling can be used to make the relation of the two error types an issue of bank policy.

The rest of the paper is organized as follows. The following two sections describe the data, performance measures and SVM methodology. In the fourth section the variable selection technique and outcome are explained. The fifth section presents the experimental settings, estimation procedure and findings, and illustrates selected results. The sixth section concludes.

DATA AND MEASURES OF ACCURACY

In this study of the potential virtues of SVM in insolvency prognosis the CreditReform database is employed. The database consists of 20,000 financially and economically solvent and 1000 insolvent German companies observed once in the period from 1997 to 2002. Although the companies were randomly selected, accounting information dates most frequently in 2001 and 2002. Approximately 50% of the observations come from this period. The industry distribution of the insolvent companies is as follows: manufacturing 25.7%, wholesale and retail trade 20.1%, real estate 9.4%, construction 39.7% and others 5.1%. The latter includes businesses in agriculture, mining, electricity, gas and water supply, transport and communication, financial intermediation social service activities and hotels and restaurants. The 20,000 solvent companies belong to manufacturing (27.4%), wholesale and retail trade (24.8%), real estate (16.9%), construction (13.9%) and others (17.1%). There is only low coincidence between the industries represented in the insolvent and the solvent group of 'others'. The latter comprises many companies in industries such as publication administration and defense, education and health. Figure 2 shows the distribution of solvent and insolvent companies across industries. A set of balance sheet and income statement items describes each company. The ones we use for further analysis are described below:

- AD (amortization and depreciation)
- AP (accounts payable)
- AR (account receivable)




Figure 2. The distribution of solvent and insolvent companies across industries

- CA (current assets)
- CASH (cash and cash equivalents)
- CL (current liabilities)
- DEBT (debt)
- EBIT (earnings before interest and tax)
- EQUITY (equity)
- IDINV (growth of inventories)
- IDL (growth of liabilities)
- INTE (interest expense)
- INV (inventories)
- ITGA (intangible assets)
- LB (lands and buildings)
- NI (net income)
- OI (operating income)
- QA (quick assets)
- SALE (sales)
- TA (total assets)
- TL (total liabilities)
- WC (working capital (= CA – CL))

The companies appear in the database several times in different years; however, each year of balance sheet information is treated as a single observation. The data of the insolvent companies were collected 2 years prior to insolvency. The company sizes are measured by total assets. We construct 28 ratios to condense the balance sheet information (see Table I). However, before dealing with the CreditReform dataset, some companies whose behavior is very different from other ones are filtered out in order to make the dataset more compact. The data pre-processing procedure is described as follows:

1. We excluded companies whose total assets were not in the range of 10^5 – 10^7 EUR (remaining insolvent: 967; solvent: 15,834).
2. In order to compute the accounting ratios AP/SALE, OI/TA, TL/TA, CASH/TA, IDINV/INV, INV/SALE, EBIT/TA and NI/SALE, we have removed companies with zero denominators (remaining insolvent: 816; solvent 11,005).
3. We dropped outliers, that is, in the insolvent class companies with extreme values of financial indices have been removed (remaining insolvent: 811; solvent: 10,468).

After pre-processing, the dataset consists of 11,279 companies (811 insolvent and 10,468 solvent). In the following analysis, we focus on the revised dataset.

The performance of the SSVM is evaluated on the basis of three measures of accuracy: Type I error rate (%), Type II error rate (%) and total error rate (%). The Type I error is the ratio of the number of insolvent companies predicted as solvent ones to the number of insolvent companies. The Type II error is the ratio of the number of solvent companies predicted as insolvent ones to the number of solvent companies. Accordingly, the error-type rates (in percentage) are defined as follows

- Type I error rate = $FN/(FN + TP) \times 100$ (%);
- Type II error rate = $FP/(FP + TN) \times 100$ (%);
- Total error rate = $(FN + FP)/(TP + TN + FP + FN) \times 100$ (%);

Table I. Definitions of accounting ratios used in the analysis

Variable	Ratio	Indicator for
X1	NI/TA	Profitability
X2	NI/SALE	Profitability
X3	OI/TA	Profitability
X4	OI/SALE	Profitability
X5	EBIT/TA	Profitability
X6	(EBIT + AD)/TA	Profitability
X7	EBIT/SALE	Profitability
X8	EQUITY/TA	Leverage
X9	(EQUITY-ITGA)/ (TA-ITGA-CASH-LB)	Leverage
X10	CL/TA	Leverage
X11	(CL-CASH)/TA	Leverage
X12	TL/TA	Leverage
X13	DEBT/TA	Leverage
X14	EBIT/INTE	Leverage
X15	CASH/TA	Liquidity
X16	CASH/CL	Liquidity
X17	QA/CL	Liquidity
X18	CA/CL	Liquidity
X19	WC/TA	Liquidity
X20	CL/TL	Liquidity
X21	TA/SALE	Activity
X22	INV/SALE	Activity
X23	AR/SALE	Activity
X24	AP/SALE	Activity
X25	Log(TA)	Size
X26	IDINV/INV	Growth
X27	IDL/TL	Growth
X28	IDCASH/CASH	Growth

where

- True positive (TP)*: Predict insolvent companies as insolvent ones
- False positive (FP)*: Predict solvent companies as insolvent ones
- True negative (TN)*: Predict solvent companies as solvent ones
- False negative (FN)*: Predict insolvent companies as solvent ones

The following matrix explains the terms used in the definition of error rates:

		Predicted class	
		Positive	Negative
Actual Class	Positive	<i>True positive (TP)</i>	<i>False negative (FN)</i>
	Negative	<i>False positive (FP)</i>	<i>True negative (TN)</i>

SVM METHODOLOGY

In recent years, the so-called support vector machines (SVM), which have their roots in the theory of statistical learning (Burgess, 1998; Christianini and Shawe-Taylor, 2000; Vapnik, 1995) have

become one of the most successful learning algorithms for classification as well as for regression (Drucker *et al.*, 1997; Mangasarian and Musicant, 2000; Smola and Schölkopf, 2004). Some features of SVM make them particularly attractive for predicting the default risk of companies. SVM are a non-parametric technique that learn the separating function from the data; they are based on a sound theoretical concept, do not require a particular distribution of the data, and deliver an optimal solution for the expected loss from misclassification. SVM estimate the separating hyperplane between defaulting and non-defaulting companies under the constraint of a maximal margin between the two classes (Vapnik, 1995; Schölkopf and Smola, 2002).

SVM can be formulated differently. However, in all variants either a constrained minimization problem or an unconstrained minimization problem is solved. The objective function in these optimization problems basically consists of two parts: a misclassification penalty part which stands for *model bias* and a regularization part which controls the *model variance*. We briefly introduce three different models: the smooth support vector machines (SSVM) (Lee and Mangasarian, 2001), the smooth support vector machines with reduced kernel technique (RSVM) and the 1-norm SVM. The SSVM will be used for classification and the 1-norm SVM will be employed for variable selection. The RSVM are applied for oversampling in order to mitigate the computational burden due to increasing the number of instances in the training sample.

Smooth support vector machines

The aim of the SVM technique is to find the separating hyperplane with the largest margin from the training data. This hyperplane is ‘optimal’ in the sense of statistical learning: it strikes a balance between overfitting and underfitting. Overfitting means that the classification boundary is too curved and therefore has less ability to classify unseen data correctly. Underfitting, on the other hand, gives a too simple classification boundary and leaves too many misclassified observations (Vapnik, 1995). We begin with linear support vector machines. Given a training dataset $S = \{(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_n, y_n)\} \subseteq \mathbb{R}^d \times \mathbb{R}$, where $\mathbf{x}_i \in \mathbb{R}^d$ is the input data and $y_i \in \{-1, 1\}$ is the corresponding class label, a conventional SVM separating hyperplane is generated by solving a convex optimization problem given as follows:

$$\begin{aligned} \min_{(w, b, \xi) \in \mathbb{R}^{d+1+n}} \quad & C \sum_{i=1}^n \xi_i + \frac{1}{2} \|w\|_2^2 \\ \text{s.t.} \quad & y_i(w^\top \mathbf{x}_i + b) + \xi_i \geq 1 \\ & \xi_i \geq 0, \quad \text{for } i = 1, 2, \dots, n \end{aligned} \quad (1)$$

where C is a positive parameter controlling the trade-off between the training error (model bias) and the part of maximizing the margin (model variance) that is achieved by minimizing $\|w\|_2^2$. In contrast to the conventional SVM of (1), smooth support vector machines minimize the square of the slack vector ξ with weight $\frac{C}{2}$. In addition, the SSVM methodology appends $\frac{b^2}{2}$ to the term that is to be minimized. This expansion results in the following minimization problem:

$$\begin{aligned} \min_{(w, b, \xi) \in \mathbb{R}^{d+1+n}} \quad & \frac{C}{2} \sum_{i=1}^n \xi_i^2 + \frac{1}{2} (\|w\|_2^2 + b^2) \\ \text{s.t.} \quad & y_i(w^\top \mathbf{x}_i + b) + \xi_i \geq 1 \\ & \xi_i \geq 0, \quad \text{for } i = 1, 2, \dots, n \end{aligned} \quad (2)$$

In a solution of (2), ξ is given by $\xi_i = \{1 - y_i(w^\top \mathbf{x}_i + b)\}_+$ for all i where the *plus* function x_+ is defined as $x_+ = \max\{0, x\}$. Thus, we can replace ξ_i in (2) by $\{1 - y_i(w^\top \mathbf{x}_i + b)\}_+$. This will convert the problem (2) into an unconstrained minimization problem as follows:

$$\min_{(w, b) \in \mathbb{R}^{d+1}} \frac{C}{2} \sum_{i=1}^n \{1 - y_i(w^\top \mathbf{x}_i + b)\}_+^2 + \frac{1}{2} (\|w\|_2^2 + b^2) \tag{3}$$

This formulation reduces the number of variables from $d + 1 + n$ to $d + 1$. However, the objective function to be minimized is not twice differentiable, which precludes the use of a fast Newton method. In the SSVM, the plus function x_+ is approximated by a smooth *p-function*, $p(x, \alpha) = x + \frac{1}{\alpha} \log(1 + e^{-\alpha x})$, $\alpha > 0$. Replacing the plus function with a very accurate smooth approximation *p-function* gives the smooth support vector machine formulation:

$$\min_{(w, b) \in \mathbb{R}^{d+1}} \frac{C}{2} \sum_{i=1}^n p(\{1 - y_i(w^\top \mathbf{x}_i + b)\}, \alpha)^2 + \frac{1}{2} (\|w\|_2^2 + b^2) \tag{4}$$

where $\alpha > 0$ is the smooth parameter. The objective function in problem (4) is strongly convex and infinitely differentiable. Hence, it has a unique solution and can be solved by using a fast Newton–Armijo algorithm. For the nonlinear case, this formulation can be extended to the nonlinear SVM by using the kernel trick as follows:

$$\min_{(u, b) \in \mathbb{R}^{n+1}} \frac{C}{2} \sum_{i=1}^n p\left(\left[1 - y_i \left\{ \sum_{j=1}^n u_j K(\mathbf{x}_i, \mathbf{x}_j) + b \right\}\right], \alpha\right)^2 + \frac{1}{2} (\|u\|_2^2 + b^2) \tag{5}$$

where $K(\mathbf{x}_i, \mathbf{x}_j)$ is a kernel function. This kernel function represents the inner product of $\phi(\mathbf{x}_i)$ and $\phi(\mathbf{x}_j)$, where ϕ is a certain mapping from input space \mathbb{R}^d to a feature space \mathcal{F} . We do not need to know the mapping of ϕ explicitly. This is the so-called kernel trick. The nonlinear SSVM classifier can be expressed in matrix form as follows:

$$\sum_{u_j \neq 0} u_j K(A_j^\top, \mathbf{x}) + b = K(\mathbf{x}, A^\top) u + b \tag{6}$$

where $A = [\mathbf{x}_1^\top; \dots; \mathbf{x}_n^\top]$ and $A_j = \mathbf{x}_j^\top$.

Reduced support vector machine

In large-scale problems, the full kernel matrix will be very large so it may not be appropriate to use the full kernel matrix when dealing with (5). In order to avoid facing such a big full kernel matrix, we brought in the reduced kernel technique (Lee and Huang, 2007). The key idea of the reduced kernel technique is to randomly select a portion of data and to generate a thin rectangular kernel matrix, then to use this much smaller rectangular kernel matrix to replace the full kernel matrix. In the process of replacing the full kernel matrix by a reduced kernel, we use the Nyström approximation (Smola and Schölkopf, 2000) for the full kernel matrix:

$$K(A, A^\top) \approx K(A, \tilde{A}^\top) K(\tilde{A}, \tilde{A}^\top)^{-1} K(\tilde{A}, A^\top) \tag{7}$$

where $K(A, A^\top) = K_{n \times n}$, $\tilde{A}_{\tilde{n} \times d}$ is a subset of A and $K(A, \tilde{A}) = \tilde{K}_{n \times \tilde{n}}$ is a reduced kernel. Thus, we have

$$K(A, A^\top)u \approx K(A, \tilde{A}^\top)K(\tilde{A}, \tilde{A}^\top)^{-1}K(\tilde{A}^\top, A)u = K(A, \tilde{A}^\top)\tilde{u} \quad (8)$$

where $\tilde{u} \in \mathbb{R}^{\tilde{n}}$ is an approximated solution of u via the reduced kernel technique. The reduced kernel method constructs a compressed model and cuts down the computational cost from $\mathcal{O}(n^3)$ to $\mathcal{O}(\tilde{n}^3)$. It has been shown that the solution of reduced kernel matrix approximates the solution of full kernel matrix well. The SSVM with the reduced kernel are called RSVM.

1-Norm support vector machine

The 1-norm support vector machine replaces the regularization term $\|w\|_2^2$ in (1) with the ℓ_1 -norm of w . The ℓ_1 -norm regularization term is also called the LASSO penalty (Tibshirani, 1996). It tends to shrink the coefficients w 's towards zeros in particular for those coefficients corresponding to redundant noise features (Zhu *et al.*, 2003; Williams and Seeger, 2001). This nice feature will lead to a way of selecting the important ratios in our prediction model. The formulation of 1-norm SVM is described as follows:

$$\begin{aligned} \min_{(w, b, \xi) \in \mathbb{R}^{d+1+n}} \quad & C \sum_{i=1}^n \xi_i + \|w\|_1 \\ \text{s.t.} \quad & y_i(w^\top \mathbf{x}_i + b) + \xi_i \geq 1 \\ & \xi_i \geq 0, \quad \text{for } i = 1, 2, \dots, n. \end{aligned} \quad (9)$$

The objective function of (9) is a piecewise linear convex function. We can reformulate it as the following linear programming problem:

$$\begin{aligned} \min_{(w, s, b, \xi) \in \mathbb{R}^{d+d+1+n}} \quad & C \sum_{i=1}^n \xi_i + \sum_{j=1}^d s_j \\ \text{s.t.} \quad & y_i(w^\top \mathbf{x}_i + b) + \xi_i \geq 1 \\ & -s_j \leq w_j \leq s_j, \quad \text{for } j = 1, 2, \dots, d, \\ & \xi_i \geq 0, \quad \text{for } i = 1, 2, \dots, n \end{aligned} \quad (10)$$

where s_j is the upper bound of the absolute value of w_j . In the optimal solution of (10) the sum of s_j is equal to $\|w\|_1$.

The 1-norm SVM can generate a very sparse solution w and lead to a parsimonious model. In a linear SVM classifier, solution sparsity means that the separating function $f(\mathbf{x}) = w^\top \mathbf{x} + b$ depends on very few input attributes. This characteristic can significantly suppress the number of nonzero coefficient w 's, especially when there are many redundant noise features (Fung and Mangasarian, 2004; Zhu *et al.*, 2003). Therefore the 1-norm SVM can be a very promising tool for the variable selection tasks. We will use it to choose the important financial indices for our bankruptcy prognosis model.

SELECTION OF ACCOUNTING RATIOS

In principle any possible combination of accounting ratios could be used as explanatory variables in a bankruptcy prognosis model. Therefore, appropriate performance measures are needed to gear the process of variable selection towards picking the ratios with the highest separating power. In

Chen *et al.* (2006) accuracy ratio (AR) and conditional information entropy ratio (CIER) determine the selection procedure's outcome. It turned out that the ratio 'accounts payable divided by sales', X24 (AP/SALE), has the best performance values for a univariate SVM model. The second selected variable was the one combined with X24 that had the best performance in a bivariate SVM model. This is the analogue of forward selection in linear regression modeling. Typically, improvement declines if new variables are added consecutively. In Chen *et al.* (2006) the performance indicators started to decrease after the model included eight variables. The described selection procedure is quite lengthy, since there are at least 216 accounting ratio combinations to be considered. We will not employ the procedure here but use the chosen set of eight variables as the benchmark set V1. Table II presents V1 in the first column.

We propose two different approaches for variable selection that will simplify the selection procedure. The first one is based on 1-norm SVM introduced above. The SVM were applied to the period from 1997 to 1999. We selected the variables according to the size of the absolute values of the coefficients w from the solution of the 1-norm SVM. Table II displays the eight selected variables as V2. We obtain eight variables out of 28. Note that five variables, X2, X3, X5, X15 and X24, are also in the benchmark set V1.

The second variable selection scheme is incremental forward variable selection. The intuition behind this scheme is that a new variable will be added into the already selected set, if it brings in the most extra information. We measure the extra information for an accounting ratio using the distance between this new ratio vector and the space spanned by the current selected ratio subset. This distance can be computed by solving a least-squares problem (Lee *et al.*, 2008). The ratio with the farthest distance will be added into the selected accounting ratio set. We repeat this procedure until a certain stopping criterion is satisfied. The accounting ratio X24 (AP/SALE) is used as the initial selected accounting ratio. Then we follow the procedure seven times to select seven more extra accounting ratios. The variable set generated is called V3. We will use these three variable sets, V1, V2 and V3, for further data analysis in the next section. The symbol $+$ denotes the variables that are common to all sets: X2, X3, X5 and X24.

Table II. Selected variables

Variable	Definition	V1	V2	V3
X2 ⁺	NI/SALE	x	x	x
X3 ⁺	OI/TA	x	x	x
X4	OI/SALE			x
X5 ⁺	EBIT/TA	x	x	x
X6	(EBIT + AD)/TA		x	
X7	EBIT/SALE			x
X8	EQUITY/TA		x	
X12	TL/TA	x		
X13	DEBT/TA			x
X15	CASH/TA	x	x	
X21	TA/SALE			x
X22	INV/SALE	x		
X23	AR/SALE		x	
X24 ⁺	AP/SALE	x	x	x
X26	IDINV/INV	x		

EXPERIMENTAL SETTING AND RESULTS

In this section we present our experimental setting and results. We compare the performance of three sets of accounting ratios, V1, V2 and V3, in our SSVM-based insolvency prognosis model. The performance is measured by Type I error rate, Type II error rate and total error rate. Fortunately, in reality, there is only a small number of insolvent companies compared to the number of solvent companies. Due to the small share in a sample that reflects reality, a simple classification such as naive Bayesian or a decision tree tends to classify every company as solvent. Such a classification would imply accepting all companies' loan applications and would thus lead to a very high Type I error rate while the total error rate and the Type II error rate are very small. Such models are useless in practice.

Our cleaned dataset consists of around 10% of insolvent companies. Thus, the sample is fairly unbalanced although the share of insolvent companies is higher than in reality. In order to deal with this problem, insolvency prognosis models usually start off with more balanced training and testing samples than reality can provide. For example, Härdle *et al.* (2007b) employ a downsampling strategy and work with balanced (50%/50%) samples. The chosen bootstrap procedure repeatedly randomly selects a fixed number of insolvent companies from the training set and adds the same number of randomly selected solvent companies. However, in this paper we adopt an oversampling strategy, to balance the size between the solvent and the insolvent companies, and refer to the downsampling procedure primarily for reasons of reference.

Oversampling duplicates the number of insolvent companies a certain number of times. In this experiment, we duplicate in each scenario the number of insolvent companies as many times as necessary to reach a balanced sample. Note that in our oversampling scheme every solvent and insolvent company's information is utilized. This increases the computational burden due to increasing the number of training instances. We employ the reduced kernel technique introduced above to mediate this problem.

All classifiers we need in these experiments are reduced SSVM with the Gaussian kernel, which is defined as

$$K(\mathbf{x}, \mathbf{z}) = e^{-\gamma \|\mathbf{x} - \mathbf{z}\|_2^2}$$

where γ is the width parameter. In nonlinear SSVM, we need to determine two parameters: the penalty term C and γ . The 2D grid search will consume a lot of time. In order to cut down the search time, we adopt the uniform design model selection method (Huang *et al.*, 2007) to search an appropriate pair of parameters.

Performance of SSVM

We conduct the experiments in a scenario in which we always train the SSVM bankruptcy prognosis model from the data at hand and then use the trained SSVM to predict the following year's cases. This strategy simulates the real task of prediction which binds the analyst to use past data for forecasting future outcomes. The experimental setting is described in Table III. The number of periods which enter the training set changes from 1 year (S1) to 5 years (S5).

In Tables IV and V we report the results for the oversampling and downsampling strategy respectively. Mean and standard deviation of Type I, Type II and total error rates (misclassification rates) are shown. We perform these experiments for the three variable sets, V1 to V3, and compare the oversampling and downsampling scheme in each experiment. All experiments are repeated 30 times

Table III. The scenario of our experiments

Scenario	Observation period of training set	Observation period of testing set
S1	1997	1998
S2	1997–1998	1999
S3	1997–1999	2000
S4	1997–2000	2001
S5	1997–2001	2002

Table IV. Results of oversampling for three variable sets (RSVM)

Set of accounting ratios	Scenario	Type I error rate		Type II error rate		Total error rate	
		Mean	SD	Mean	SD	Mean	SD
V1	S1	33.16	0.55	26.15	0.13	26.75	0.12
	S2	31.58	0.01	29.10	0.07	29.35	0.07
	S3	28.11	0.73	26.73	0.16	26.83	0.16
	S4	30.14	0.62	25.66	0.17	25.93	0.15
	S5	24.24	0.56	23.44	0.13	23.48	0.13
V2	S1	29.28	0.92	27.20	0.24	27.38	0.23
	S2	28.20	0.29	30.18	0.18	29.98	0.16
	S3	27.41	0.61	29.67	0.19	29.50	0.17
	S4	28.12	0.74	28.32	0.19	28.31	0.15
	S5	23.91	0.62	24.99	0.10	24.94	0.10
V3	S1	29.28	0.83	25.11	0.25	25.46	0.21
	S2	31.27	0.62	29.79	0.34	29.94	0.35
	S3	30.91	0.13	27.21	0.19	27.48	0.18
	S4	32.00	0.54	25.19	0.17	25.61	0.14
	S5	26.98	0.42	22.90	0.11	23.08	0.11

Table V. Results of downsampling for three variable sets (SSVM with Gaussian kernel)

Set of accounting ratios	Scenario	Type I error rate		Type II error rate		Total error rate	
		Mean	SD	Mean	SD	Mean	SD
V1	S1	32.20	3.12	28.98	1.70	29.26	1.46
	S2	29.74	2.29	28.77	1.97	28.87	1.57
	S3	30.46	1.88	26.23	1.33	26.54	1.17
	S4	31.55	1.52	23.89	0.97	24.37	0.87
	S5	28.81	1.53	23.09	0.73	23.34	0.69
V2	S1	29.94	2.91	28.07	2.15	28.23	1.79
	S2	28.77	2.58	29.80	1.89	29.70	1.52
	S3	29.88	1.88	27.19	1.32	27.39	1.19
	S4	29.06	1.68	26.26	1.00	26.43	0.86
	S5	26.92	1.94	25.30	1.17	25.37	1.06
V3	S1	30.87	3.25	26.61	2.45	26.98	2.11
	S2	33.31	2.16	28.60	2.01	29.08	1.65
	S3	31.82	1.52	26.41	1.45	26.80	1.31
	S4	35.0	2.13	24.29	0.77	24.96	0.68
	S5	30.66	1.60	21.92	0.96	22.30	0.92

because of the randomness in the experiments. The randomness is very obvious in the downsampling scheme (see Table V). Each time we only choose negative instances with the same size of the whole positive instances. The observed randomness in our oversampling scheme (Table IV) is due to applying the reduced kernel technique to solving the problem. We use the training set in the downsampling scheme as the reduced set. That is, we use all the insolvent instances and the equal number of solvent instances as our reduced set in generating the reduced kernel. Then we duplicate the insolvent part of the kernel matrix to balance the size of insolvent and solvent companies.

Both tables reveal that different variable selection schemes produce dissimilar results with respect to both precision and deviation of predicting. The oversampling scheme shows better results in the Type I error rate but has slightly bigger total error rates. It is also obvious that in almost all models a longer training period works in favor of accuracy of prediction. Clearly, the oversampling schemes have much smaller standard deviations in the Type I error rate, Type II error rate, and total error rate than the downsampling one. According to this observation, we conclude that the oversampling scheme will generate a more robust model than the downsampling scheme.

Figure 3 illustrates the development (learning curve) of the Type I error rate and total error rate with regard to variable set V3 for both oversampling and downsampling. The bullets on the lines




Figure 3. Learning curve for variables set V3

mark the different training scenarios. For example, the first bullets from the left represent S1 (training set from 1997, testing set from 1998), the second bullets illustrate S2 (training set from 1997 to 1998, testing set from 1999) etc. For the purpose of better visibility, the Type I error rate is only indirectly displayed as $100 - \text{Type I error rate}$. The upper solid line in gray represents the oversampling scheme and the black solid line the downsampling one. Note that the performance in terms of the Type I error rate is worse the higher the distance between the upper end of the diagram and the solid lines. The learning curve over the time frame the training sample covers shows an upward tendency between S1 and S5 for the number $100 - \text{Type I error rate}$. However, the curves are non-monotonic. There is a disturbance for the forecast of year 1999 that is based on training samples that cover 1997 to 1998, and also one for the forecast of year 2001 based on training samples covering 1997 to 2000. Both disturbances may have been caused by the reform of the German insolvency code that came into force in 1999. The most important objective of the reform was to allow for more company restructuring and less liquidation than before. This reform considerably changed the behavior of German companies towards declaring insolvency, and thus most likely the nature of balance sheets that are associated with insolvent companies.

The disturbances are less visible with respect to the overall performance. The dashed lines near the lower edge of the diagram box show total error rates, gray for the oversampling and black for the downsampling scheme. There is a clear tendency towards a lower total error rate from S2 to S5 for both schemes. The downsampling line is slightly below the oversampling one, representing a slightly better performance in terms of the mean of the total error rate. However, this result has to be seen in the light of the trade-off between magnitude and stability of results, as oversampling yields much more stable results. The standard deviations for V3 are only a small portion of the numbers generated by the downsampling procedure across all training scenarios (Tables IV and V).

Table VI presents the comparison between the sets by focusing on the total error rate. It indicates by an asterisk whether the differences in means are significant at the 10% level via t -test and, in addition, gives the set which is superior in the dual comparison. Variable set V2 is nearly absent in Table VI. Thus V2 is clearly outperformed by both sets V1 and V3. There is no clear distinction between V1 and V3 except for Scenario S5. Given the long training period V3 is superior in both the downsampling and oversampling scenarios and generates the lowest total error rate in absolute terms.

In order to investigate the effect of the oversampling versus the downsampling scheme we follow the setting as above, but we use the V3 variable set. For each training–test pair, we carry out oversampling for positive instances from 6 to 15 times. We show the trend and effect in Figure 4. It is

Table VI. Statistical significance in differences in means (10% level) between the three variable sets: total error

Sets	S1	S2	S3	S4	S5
<i>Oversampling</i>					
V1 vs. V2	V1*	V1*	V1*	V1*	V1*
V1 vs. V3	V3*	V1*	V1*	V3*	V3*
V2 vs. V3	V3*		V3*	V3*	V3*
<i>Downsampling</i>					
V1 vs. V2	V2*	V1*	V1*	V1*	V1*
V1 vs. V3	V3*			V1*	V3*
V2 vs. V3	V3*		V3*	V3*	V3*




Figure 4. The effect of oversampling on Type I and Type II error rates for scenario S5 and variables set V3

easy to see that the Type I (II) error rate decreases (increases) as the oversampling times increase. This feature implies that the machine would have a tendency of classifying all companies as solvent if the training sample had realistic shares of insolvent and solvent companies. Such behavior would produce a Type I error rate of 100%. The more balanced the sample is, the higher the penalty for classifying insolvent companies as solvent. This fact is illustrated in Figure 4 by the decreasing curve with respect to the number of duplications of insolvent companies.

Often banks favor a strategy that allows them to minimize the Type II errors for a given number of Type I errors. The impact of oversampling on the trade-off between the two types of errors—shown in Figure 4—implies that the number of oversampling times is a strategic variable in training the machine. This number can be determined by the bank's aim regarding the relation of Type I and Type II errors.

Comparison with logit and linear discriminant analysis

The examination of SSVM is incomplete without comparing it to highly used traditional methods such as the logistic model (LM) and linear discriminant analysis (DA). Therefore, we replicate the research design of the previous section with both traditional models. In addition, we test whether the difference in means in the total error rate is statistically significant. The comparison of means with regard to the total error rate is presented in Tables VII and VIII for the oversampling and downsampling strategy respectively. Table IX summarizes the comparison of the approaches and displays the statistical significance of their mean differences. Asterisks indicate the out-performance

Table VII. Comparison of the total error rate (%) as generated by SSVM with LM and DA: oversampling for three variable sets

Set of accounting ratios	Scenario	SSVM	LM	DA
		Mean	Mean	Mean
V1	S1	26.75	26.50	25.60
	S2	29.35	28.96	27.22
	S3	26.83	28.94	27.42
	S4	25.93	26.20	25.55
	S5	23.48	26.95	28.23
V2	S1	27.38	26.80	26.20
	S2	29.98	28.63	28.70
	S3	29.50	29.52	29.46
	S4	28.31	28.43	28.08
	S5	24.94	29.22	31.42
V3	S1	25.46	25.07	23.65
	S2	29.94	28.29	27.02
	S3	27.48	27.89	25.84
	S4	25.61	26.60	24.85
	S5	23.08	25.32	26.15

Table VIII. Comparison of the total error rate (%) as generated by SSVM with LM and DA: downsampling for three variable sets

Set of accounting ratios	Scenario	SSVM	LM	DA
		Mean	Mean	Mean
V1	S1	29.26	26.86	27.34
	S2	28.87	28.62	28.26
	S3	26.54	27.54	28.22
	S4	24.37	24.80	25.47
	S5	23.34	24.81	25.86
V2	S1	28.23	27.28	28.62
	S2	29.70	29.29	29.65
	S3	27.39	28.56	29.58
	S4	26.43	26.41	27.96
	S5	25.37	26.52	29.69
V3	S1	26.98	26.03	25.47
	S2	29.08	28.04	27.22
	S3	26.80	26.60	26.51
	S4	24.96	25.25	25.44
	S5	22.30	23.96	24.31

Table IX. Statistical significance in differences of means (10% level) between SSVM and LM and SSVM and DA, respectively, for the sets V1 to V3: total error rate

V1	S1	S2	S3	S4	S5
<i>Oversampling</i>					
SSVM vs. LM			*	*	*
SSVM vs. DA			*		*
<i>Downsampling</i>					
SSVM vs. LM			*	*	*
SSVM vs. DA			*	*	*
V2	S1	S2	S3	S4	S5
<i>Oversampling</i>					
SSVM vs. LM				*	*
SSVM vs. DA					*
<i>Downsampling</i>					
SSVM vs. LM			*		*
SSVM vs. DA			*	*	*
V3	S1	S2	S3	S4	S5
<i>Oversampling</i>					
SSVM vs. LM			*	*	*
SSVM vs. DA					*
<i>Downsampling</i>					
SSVM vs. LM					*
SSVM vs. DA				*	*

of the logistic model or discriminant analysis by SSVMs at the 10% level via *t*-test. It is obvious that the SSVM technique yields the better results, the longer the period is from which the training observations are taken. In fact, the results show that the SSVM works significantly better than LM and DA in most cases in S3 to S5, with the clearest advantage for testing sets S4 and S5, where the accounting information of the predicted companies dates most frequently in 2001 and 2002.

We also investigate the effect of oversampling on LM and DA. We follow the same setting in the previous section, doing oversampling for positive instances from 6 to 15 times. Unlike the SSVM-based insolvency prognosis model, the DA approach is insensitive in both Type I and Type II error rates to the replication of positive instances. The result for DA is illustrated in Figure 5. The LM approach has very similar results to the SSVM model. We will not show the result here.

More data visualization

Each SSVM model has its own output value. We use this output to construct 2D coordinate systems. Figure 6 shows an example for scenario S5 where the scores of the SSVM_{V3} model (SSVM_{V1} model) are represented by the horizontal (vertical) line. A positive (negative) value indicates predicted insolvency (solvency). We then map all insolvent companies in the testing set onto the coordinate systems. There are 132 insolvent companies and 2866 solvent companies in this testing set. We also randomly choose the same amount of solvent companies from the testing set.

The plus points in the lower left quadrant and the circle points in the upper right quadrant show the number of Type I errors and Type II errors, respectively, in both models. Plus points in the upper right quadrant and circle points in the lower left quadrant reflect those companies that are predicted




Figure 5. The effect of oversampling on Type I and Type II error rates for scenario S5 and variables set V3 in DA

correctly by both models. Circles and plus points in the lower right quadrant (upper left quadrant) represent conflicting prognoses. We also report the number of insolvent companies and the number of solvent companies in each quadrant of Figure 6. The two different insolvency prognosis models based on V1 and V3, respectively, can be considered as alternative experts. The two forecasts for each instance in the testing set is plotted in the diagram. The proposed visualization scheme could be used to support loan officers in their final decision on accepting or rejecting a client's application. Furthermore, this data visualization scheme can also be applied to two different learning algorithms, such as $SSVM_{V3}$ vs. LM_{V3} and $SSVM_{V3}$ vs. DA_{V3} . We show these data visualization plots in Figures 7 and 8. If the loan application has been classified as solvent or insolvent by alternative machines, it is most likely that the prognosis meets reality (the plus points in the upper right quadrant and the circle points in the lower left quadrant). Opposing forecasts, however, should be taken as a hint to evaluate the particular company more thoroughly, for example by employing an expert team, or even by using a third model.




Figure 6. Data visualization via model one (generated by SSVM with V3) and model two (generated by SSVM with V1) in scenario S5

CONCLUSION

In this paper we apply different variants of support vector machines to a unique dataset of German solvent and insolvent companies. We use a priori a given set of predictors as a benchmark, and suggest two further variable selection procedures; the first procedure uses the 1-norm SVM and the second, incremental way consecutively selects the variable that is the farthest one from the column space of the current variable set. Given the three SSVM based on distinct variable sets, the relative performance of the types of smooth support vector machines is tested. The performance is measured by error rates. The two sets of variables newly created here lead to a dissimilar performance of SSVM. The selection of variables by the 1-norm SVM clearly underperforms compared to the incremental selection scheme. This difference in accuracy hints at the need for further research with respect to the variable selection. The training period makes a clear difference, though. Results improve considerably if more years of observation are used in training the machine. The SSVM




Figure 7. Data visualization via model one (generated by SSVM with V3) and model two (generated by LM with V3) in scenario S5

model benefits more from longer training periods than traditional methods do. As a consequence the logit model and discriminant analysis are both outperformed by the SSVM in long-term training scenarios. Moreover, the oversampling scheme works very well in dealing with unbalanced datasets. It provides flexibility to control the trade-off between Type I and Type II errors, and is therefore a strategic instrument in a bank's hand. The results generated are very stable in terms of small deviations of Type I, Type II and total error rates.

Finally, we want to stress that SSVM should be considered not as a substitute for traditional methods but rather as a complement which, when employed side by side with either the logit model or discriminant analysis, can generate new information that helps practitioners select those companies that are difficult to predict and, therefore, need more attention and further treatment.




Figure 8, Data visualization via model one (generated by SSVM with V3) and model two (generated by DA with V3) in scenario S5

ACKNOWLEDGEMENTS

This research was supported by the ‘Stiftung Geld und Währung’ and by the Deutsche Forschungsgemeinschaft through the SFB 649 ‘Economic Risk’.

REFERENCES

Altman E. 1968. Financial ratios, discriminant analysis and the prediction of corporate bankruptcy. *Journal of Finance* **23**(4): 589–609.
 Altman E, Marco G, Varetto F. 1994. Corporate distress diagnosis: comparisons using linear discriminant analysis and neural networks (the italian experience). *Journal of Banking and Finance* **18**: 505–529.

- Beaver W. 1966. Financial ratios as predictors of failures: empirical research in accounting: selected studies. *Journal of Accounting Research* **4**: 71–111.
- Burges CJC. 1998. A tutorial on support vector machines for pattern recognition. *Data Mining and Knowledge Discovery* **2**(2): 121–167.
- Chen S, Härdle W, Moro RA. 2006. Estimation of default probabilities with support vector machines. *SFB 649 Discussion Paper 2006-077*.
- Cristianini N, Shawe-Taylor J. 2000. *An Introduction to Support Vector Machines*. Cambridge University Press: Cambridge, UK.
- Drucker H, Burges CJC, Kaufman L, Smola A, Vapnik V. 1997. Support vector regression machines. In *Advances in Neural Information Processing Systems 9*, Mozer MC, Jordan MI, Petsche T (eds). MIT Press: Cambridge, MA; 155–161.
- Fung G, Mangasarian OL. 2004. A feature selection Newton method for support vector machine classification. *Computational Optimization and Applications* **28**(2): 185–202.
- Härdle W, Moro R, Schäfer D. 2007a. Graphical data representation in bankruptcy analysis based on support vector machines. In *Handbook of Data Visualization*, Chen C, Härdle W, Unwin A (eds). Springer: Heidelberg; 853–872.
- Härdle W, Moro RA, Schäfer D. 2007b. Estimating probabilities of default with support vector machines. *SFB 649 Discussion Paper 2007-035*.
- Huang CM, Lee YJ, Lin DKJ, Huang SY. 2007. Model selection for support vector machines via uniform design. *Computational Statistics and Data Analysis* **52**: 335–346. Special Issue on Machine Learning and Robust Data Mining (to appear).
- Hwang RC, Cheng KF, Lee JC. 2007. A semiparametric method for predicting bankruptcy. *Journal of Forecasting* **26**(5): 317–342.
- Krahen JP, Weber M. 2001. Generally accepted rating principles: a primer. *Journal of Banking and Finance* **25**(1): 3–23.
- Lee YJ, Huang SY. 2007. Reduced support vector machines: a statistical theory. *IEEE Transactions on Neural Networks* **18**: 1–13.
- Lee YJ, Mangasarian OL. 2001. SSVM: a smooth support vector machine. *Computational Optimization and Applications* **20**: 5–22.
- Lee YJ, Chang CC, Chao CH. 2008. Incremental forward feature selection with application to microarray gene expression. *Journal of Biopharmaceutical Statistics* **18**(5): 824–840.
- Leland H, Toft K. 1996. Optimal capital structure, endogenous bankruptcy, and the term structure of credit spreads. *Journal of Finance* **51**: 987–1019.
- Longstaff FA, Schwartz ES. 1995. A simple approach to valuating risky fixed and floating rate debt. *Journal of Finance* **50**: 789–819.
- Mangasarian OL, Musicant DR. 2000. Robust linear and support vector regression. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **22**(9): 950–955.
- Martin D. 1977. Early warning of bank failure: a logit regression approach. *Journal of Banking and Finance* **1**: 249–276.
- Mella-Barral P, Perraudin W. 1997. Strategic debt service. *Journal of Finance* **52**: 531–556.
- Merton R. 1974. On the pricing of corporate debt: the risk structure of interest rates. *Journal of Finance* **29**(2): 449–470.
- Ohlson J. 1980. Financial ratios and the probabilistic prediction of bankruptcy. *Journal of Accounting Research* **18**(1): 109–131.
- Schölkopf B, Smola AJ. 2002. *Learning with Kernels*. MIT Press: Cambridge, MA.
- Smola A, Schölkopf B. 2000. Sparse greedy matrix approximation for machine learning. In *Proceedings of the 17th International Conference on Machine Learning*, San Francisco, CA.
- Smola A, Schölkopf B. 2004. A tutorial on support vector regression. *Statistics and Computing* **14**: 199–222.
- Tam K, Kiang M. 1992. Managerial application of neural networks: the case of bank failure prediction. *Management Science* **38**(7): 926–947.
- Tibshirani R. 1996. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society* **58**(1): 267–288.
- Vapnik VN. 1995. *The Nature of Statistical Learning Theory*. Springer: New York.
- Williams CKI, Seeger M. 2001. Using the Nyström method to speed up kernel machines. *Advances in Neural Information Processing Systems* **13**: 682–688.

- Zhou C. 2001. The term structure of credit spreads with jump risk. *Journal of Banking and Finance* **25**: 2015–2040.
- Zhu J, Rosset S, Hastie T, Tibshirani R. 2003. 1-Norm support vector machines. In *Advances in Neural Information Processing Systems* **16**: 49–56.

Authors' biographies:

Wolfgang Härdle did in 1982 his Dr. rer. nat. in Mathematics at Universität Heidelberg and in 1988 his Habilitation at Universität Bonn. He is currently chair professor of statistics at the Dept. of Economics and Business Administration, Humboldt-Universität zu Berlin. He is also director of CASE—Center for Applied Statistics & Economics and of the Collaborative Research Center 'Economic Risk'. His research focuses on dimension reduction techniques, computational statistics and quantitative finance. He has published 34 books and more than 200 papers in top statistical, econometrics and finance journals. He is one of the 'Highly cited Scientist' according to the Institute of Scientific Information.

Yuh-Jye Lee received his Master degree in Applied Mathematics from the National Tsing Hua University, Taiwan in 1992 and PhD degree in computer sciences from the University of Wisconsin-Madison in 2001. In 2002, Dr. Lee joined the Computer Science and Information Engineering Department, National Taiwan University of Science and Technology. He is an associate professor now. His research interests are in machine learning, data mining, optimization, information security and operations research. He developed new algorithms for large data mining problems such as classification problem, clustering, feature selection and dimension reduction. These algorithms have been used in intrusion detection systems (IDS), face detection, micro array gene expression analysis and breast cancer diagnosis and prognosis.

Dorothea Schäfer did in 1992 her Dr. rer. pol. in Economics and in the year 2000 her Habilitation at Freie Universität Berlin. She is currently coordinator of the research group Financial Markets and Financial Institutions and senior researcher at the German Institute for Economic Research (DIW) Berlin which she joined in 2002. She is managing editor of the Quarterly Journal of Economic Research (Vierteljahreshefte zur Wirtschaftsforschung) and adjunct lecturer at Freie Universität Berlin. Her research focuses on insolvency risk, financial management of firms and banks, and on behavioural finance.

Yi-Ren Yeh received the M.S. degree from the Department of Computer Science and Information Engineering, National Taiwan University of Science and Technology, Taipei, Taiwan, R.O.C., in 2006. He is currently working toward the PhD degree in the Department of Computer Science and Information Engineering, National Taiwan University of Science and Technology. His research interests include machine learning, data mining, and information security.

Authors' addresses:

Wolfgang Härdle, Center for Applied Statistics and Economics, Humboldt-Universität zu Berlin, Spandauer Straße 1, 10178 Berlin, Germany.

Yuh-Jye Lee and **Yi-Ren Yeh**, Department of Computer Science Information Engineering, National Taiwan University of Science and Technology, Taipei 106, Taiwan.

Dorothea Schäfer, German Institute for Economic Research (DIW) Berlin, Mohrenstrasse 58, 10117 Berlin, Germany.

Time Series Modelling With Semiparametric Factor Dynamics

Byeong U. PARK, Enno MAMMEN, Wolfgang HÄRDLE, and Szymon BORAK

High-dimensional regression problems, which reveal dynamic behavior, are typically analyzed by time propagation of a few number of factors. The inference on the whole system is then based on the low-dimensional time series analysis. Such high-dimensional problems occur frequently in many different fields of science. In this article we address the problem of inference when the factors and factor loadings are estimated by semiparametric methods. This more flexible modeling approach poses an important question: Is it justified, from an inferential point of view, to base statistical inference on the estimated time series factors? We show that the difference of the inference based on the estimated time series and “true” unobserved time series is asymptotically negligible. Our results justify fitting vector autoregressive processes to the estimated factors, which allows one to study the dynamics of the whole high-dimensional system with a low-dimensional representation. We illustrate the theory with a simulation study. Also, we apply the method to a study of the dynamic behavior of implied volatilities and to the analysis of functional magnetic resonance imaging (fMRI) data.

KEY WORDS: Asymptotic inference; Factor models; Implied volatility surface; Semiparametric models; Vector autoregressive process.

1 INTRODUCTION

Modeling for high-dimensional data is a challenging task in statistics especially when the data comes in a dynamic context and is observed at changing locations with different sample sizes. Such modeling challenges appear in many different fields. Examples are Stock and Watson (2005) in empirical macroeconomics, Lee and Carter (1992) in mortality analysis, Nelson and Siegel (1987) and Diebold and Li (2006) in bond portfolio risk management or derivative pricing, Martinussen and Scheike (2000) in biomedical research. Other examples include the studies of radiation treatment of prostate cancer by Kauermann (2000) and evoked potentials in Electroencephalogram (EEG) analysis by Gasser, Möcks, and Verleger (1983). In financial engineering, it is common to analyze the dynamics of implied volatility surface for risk management. For functional magnetic resonance imaging data (fMRI), one may be interested in analyzing the brain’s response over time as well as identifying its activation area, see Worsley et al. (2002).

A successful modeling approach utilizes factor type models, which allow low-dimensional representation of the data. In an orthogonal L -factor model an observable J -dimensional random vector $Y_t = (Y_{t,1}, \dots, Y_{t,J})^T$ can be represented as

$$Y_{t,j} = m_{0,j} + Z_{t,1}m_{1,j} + \dots + Z_{t,L}m_{L,j} + \varepsilon_{t,j}, \quad (1)$$

where $Z_{t,l}$ are common factors, $\varepsilon_{t,j}$ are errors or specific factors, and the coefficients $m_{l,j}$ are factor loadings. In most applications, the index $t = 1, \dots, T$ reflects the time evolution of the whole system, and Y_t can be considered as a multidimensional time series. For a method to identify common factors in this model we refer to Peña and Box (1987). The study of high-dimensional Y_t is then simplified to the modeling of $Z_t = (Z_{t,1},$

$\dots, Z_{t,L})^T$, which is a more feasible task when $L \ll J$. The model (1) reduces to a special case of the generalized dynamic factor model considered by Forni, Hallin, Lippi, and Reichlin (2000), Forni and Lippi (2001) and Hallin and Liska (2007), when $Z_{t,l} = a_{l,1}(B)U_{t,1} + \dots + a_{l,q}(B)U_{t,q}$ where the q -dimensional vector process $U_t = (U_{t,1}, \dots, U_{t,q})^T$ is an orthonormal white noise and B stands for the lag operator. In this case, the model (1) is expressed as $Y_{t,j} = m_{0,j} + \sum_{k=1}^q b_{k,j}(B)U_{t,k} + \varepsilon_{t,j}$, where $b_{k,j}(B) = \sum_{l=1}^L a_{l,k}(B)m_{l,j}$.

In a variety of applications, one has explanatory variables $X_{t,j} \in \mathbb{R}^d$ at hand that may influence the factor loadings m_l . An important refinement of the model (1) is to incorporate the existence of observable covariates $X_{t,j}$. The factor loadings are now generalized to functions of $X_{t,j}$, so that the model (1) is generalized to

$$Y_{t,j} = m_0(X_{t,j}) + \sum_{l=1}^L Z_{t,l}m_l(X_{t,j}) + \varepsilon_{t,j}, \quad 1 \leq j \leq J_t. \quad (2)$$

In this model, $Z_{t,l}$ for each $l: 1 \leq l \leq L$ enters into all $Y_{t,j}$ for j such that $m_l(X_{t,j}) \neq 0$. Note that the probability of the event that $m_l(X_{t,j}) = 0$ for some $1 \leq j \leq J$ equals zero if $m_l(x) = 0$ at countably many points of x and the density f_l of $X_{t,j}$ is supported on an interval with nonempty interior, as we assume at (A2) in Section 5.

The model (2) can be interpreted as a discrete version of the following functional extension of the model (1):

$$Y_t(x) = m_0(x) + \sum_{l=1}^L Z_{t,l}m_l(x) + \varepsilon_t(x), \quad (3)$$

where $\varepsilon_t(\cdot)$ is a mean zero stochastic process, and also regarded as a regression model with embedded time evolution. It is different from varying-coefficient models, such as in Fan, Yao, and Cai (2003) and Yang, Park, Xue, and Härdle (2006), because Z_t is unobservable. Our model also has some similarities to the one considered in Connor and Linton (2007) and Connor, Hagmann, and Linton (2007), which generalized the study of Fama and French (1992) on the common movements of stock price returns. There, the covariates, denoted by $X_{l,j}$, are

Byeong U. Park is Professor, Department of Statistics, Seoul National University Seoul 151-747, Korea (E-mail: bupark@stats.snu.ac.kr). Enno Mammen is Professor, Department of Economics, University of Mannheim, 68131 Mannheim, Germany (E-mail: emammen@rumms.uni-mannheim.de). Wolfgang Härdle is Professor, Institute for Statistics and Econometrics, Humboldt Universität zu Berlin, D-10178 Berlin, Germany (E-mail: haerdle@wiwi.hu-berlin.de). Szymon Borak is Ph.D. Student, Institute for Statistics and Econometrics, Humboldt Universität zu Berlin, D-10178 Berlin, Germany (E-mail: szymon.borak@gmail.de). The authors gratefully acknowledge financial support Deutsche Forschungsgemeinschaft and the Sonderforschungsbereich 649 “Ökonomisches Risiko.” Byeong U. Park’s research was supported by the Korea Research Foundation Grant funded by the Korean Government (MOEHRD) (KRF-2005-070-C00021). The authors thank the associate editor and referees for their helpful comments and suggestions.




Figure 1. The typical IV data design on two different days. In the maturity direction observations appear in the discrete points for each particular day. Bottom solid lines indicate the observed maturities. Left panel: observations on 2004.07.08, $J_t = 5,606$. Right panel: observations on 2004.08.19, $J_t = 8,152$.

time-invariant and are different for different m_t , which allows a direct application of backfitting procedures and makes the problem quite different from our setting. Some linear models, which allow time-varying coefficients, as considered in Hansen, Nielsen, and Nielsen (2004) and Bruback and Rice (1998), may be recognized as a special case of (2).

In this article we consider the model (2) with unknown nonparametric functions m_t . We call this model a dynamic semiparametric factor model (DSFM). The evolution of complex high-dimensional objects may be described by (2), so that their analysis can be reduced to the study of a low-dimensional vector of factors Z_t . In the present article, we consider an efficient nonparametric method of fitting the model. We provide relevant theory for the method as well as illustrate its empirical aspects through a simulation and a real data application. Fengler, Härdle, and Mammen (2007) used a kernel smoothing approach for the same model, but it was focused on a particular data application without offering any discussion of numerical issues, statistical theory, and simulation analysis.

One of the main motivations for the model (2) comes from a special structure of the implied volatility (IV) data, as is observed in Figure 1. The IV is a volatility parameter that matches the observed plain vanilla option prices with the theoretical ones given by the formula of Black and Scholes (1973). Figure 1 shows the special “string” structure of the IV data obtained from the European option prices on the German stock index DAX (ODAX) for two different days. The volatility strings shift toward expiry, which is indicated by the bottom line in the figure. Moreover the shape of the IV strings is subject to stochastic deformation. Fengler et al. (2007) proposed to use the model (2) to describe the dynamics of the IV data, where $Y_{t,j}$ are the values of IV or those of its transformation on the day t , and $X_{t,j}$ are the two-dimensional vectors of the moneyness and time-to-maturity. For more details on the data design and econometric motivation, we refer to Fengler et al. (2007).

One may find another application of the model (2) in the analysis of functional magnetic resonance imaging (fMRI) data. The fMRI is a noninvasive technique of recording brain’s signals on spatial area in every particular time period (usually 1–4 sec). One obtains a series of three-dimensional images of the blood-oxygen-level-dependent (BOLD) fMRI signals, whereas an exercised person is subject to certain stimuli. An example of the images in 15 different slices at one particular time point is presented in Figure 2. For the more detailed

description on the fMRI methodology we refer to Logothetis and Wandell (2004). The main aims of the statistical methods in this field are identification of the brain’s activation areas and analysis of its response over time. For this purpose the model (2) can be applied. DSFM may be applied to many other problems, such as modeling of yield curve evolution where the standard approach is to use the parametric factor model proposed by Nelson and Siegel (1987).

Our methods produce estimates of the true unobservable Z_t , say \hat{Z}_t , as well as estimates of the unknown functions m_t . In practice, one operates on these estimated values of Z_t for further statistical analysis of the data. In particular, for the IV application, one needs to fit an econometric model to the estimated factors \hat{Z}_t . For example, Hafner (2004) and Cont and da Fonseca (2002) fitted an AR(1) process to each factor, and Fengler et al. (2007) considered a multivariate VAR(2) model. The main question that arises from these applications is whether the inference based on \hat{Z}_t is equivalent to the one based on Z_t . Attempting to give an answer to this question forms the core of this article.

It is worthwhile to note here that Z_t is not identifiable in the model (2). There are many versions of (Z_t, m) , where $m = (m_0, \dots, m_L)^T$, that give the same distribution of Y_t . This means that estimates of Z_t and m_t are not uniquely defined. We show that for any version of $\{Z_t\}$ there exists a version of $\{\hat{Z}_t\}$ whose lagged covariances are asymptotically the same as those of $\{Z_t\}$. This justifies the inference based on $\{\hat{Z}_t\}$ when $\{Z_t\}$ is a VAR process, in particular. We confirm this theoretical result by a Monte Carlo simulation study. We also discuss fitting the model to the real ODAX IV and fMRI data.

The article is organized as follows. In the next section we propose a new method of fitting DSFM and an iterative algorithm that converges at a geometric rate. In Section 3 we present the results of a simulation study that illustrate the theoretical findings given in Section 5. In Section 4 we apply the model to the ODAX IV and fMRI data. Section 5 is devoted to the asymptotic analysis of the method. Technical details are provided in the Appendix.

2. METHODOLOGY

We observe $(X_{t,j}, Y_{t,j})$ for $j = 1, \dots, J_t$ and $t = 1, \dots, T$ such that

$$Y_{t,j} = Z_t^T m(X_{t,j}) + \varepsilon_{t,j}. \tag{4}$$




Figure 2. Typical fMRI data in one particular time point. The figure presents 15 parallel horizontal images. The brightness corresponds to the strength of the observed signals.

Here $Z_t^\top = (1, Z_t^\top)$ and $Z_t = (Z_{t,1}, \dots, Z_{t,L})^\top$ is an unobservable L -dimensional process. The function m is an $(L + 1)$ -tuple (m_0, \dots, m_L) of unknown real-valued functions m_l defined on a subset of \mathbb{R}^d . The variables $X_{1,1}, \dots, X_{T,J_T}, \varepsilon_{1,1}, \dots, \varepsilon_{T,J_T}$ are independent. The errors $\varepsilon_{t,j}$ have zero means and finite second moments. For simplicity of notation, we will assume that the covariates $X_{t,j}$ have support $[0, 1]^d$, and also that $J_t \equiv J$ do not depend on t .

For the estimation of m , we use a series estimator. For an integer $K \geq 1$, we choose functions $\psi_1, \dots, \psi_K: [0, 1]^d \rightarrow \mathbb{R}$, which are normalized so that $\int_{[0,1]^d} \psi_k^2(x) dx = 1$. For example, one may take $\{\psi_k: 1 \leq k \leq K\}$ to be a tensor B-spline basis (e.g., see de Boor 2001). Then, an $(L + 1)$ -tuple of functions $m = (m_0, \dots, m_L)^\top$ may be approximated by $\mathcal{A}\psi$, where $\mathcal{A} = (\alpha_{l,k})$ is an $(L + 1) \times K$ matrix and $\psi = (\psi_1, \dots, \psi_K)^\top$. We define the least squares estimators $\widehat{Z}_t = (\widehat{Z}_{t,1}, \dots, \widehat{Z}_{t,L})^\top$ and $\widehat{\mathcal{A}} = (\widehat{\alpha}_{l,k})$:

$$S(\mathcal{A}, z) \equiv \sum_{t=1}^T \sum_{j=1}^J \{Y_{t,j} - (1, z_t^\top) \mathcal{A} \psi(X_{t,j})\}^2 = \min_{\mathcal{A}, z}! \quad (5)$$

where $z = (z_1^\top, \dots, z_T^\top)^\top$ for L -dimensional vectors z_t . With $\widehat{\mathcal{A}}$ at hand, we estimate m by $\widehat{m} = \widehat{\mathcal{A}}\psi$.

We note that, given z or \mathcal{A} , the function S in (5) is quadratic with respect to the other variables, and thus has an explicit unique minimizer. However, minimization of S with respect to \mathcal{A} and z simultaneously is a fourth-order problem. The solution is neither unique nor explicit. It is unique only up to the values of $\widehat{z}_1^\top \widehat{\mathcal{A}}, \dots, \widehat{z}_T^\top \widehat{\mathcal{A}}$, where $\widehat{z}_t^\top = (1, \widehat{Z}_t^\top)$. We will come back to this identifiability issue later in this section.

To find a solution $(\widehat{\mathcal{A}}, \widehat{z})$ of the minimization problem (5), one might adopt the following iterative algorithm: (i) Given an initial choice $Z^{(0)}$, minimize $S(\mathcal{A}, Z^{(0)})$ with respect to \mathcal{A} , which is an ordinary least squares problem and thus has an explicit unique solution. Call it $\mathcal{A}^{(1)}$. (ii) Minimize $S(\mathcal{A}^{(1)}, z)$ with respect to z now, which is also an ordinary least squares

problem. (iii) Iterate (i) and (ii) until convergence. This is the approach taken by Fengler et al. (2007). However, the procedure is not guaranteed to converge to a solution of the original problem.

We propose to use a Newton-Raphson algorithm. Let $\alpha \equiv \alpha(\mathcal{A})$ denote the stack form of $\mathcal{A} = (\alpha_{l,k})$ [i.e., $\alpha = (\alpha_{0,1}, \dots, \alpha_{L,1}, \alpha_{0,2}, \dots, \alpha_{L,2}, \dots, \alpha_{0,K}, \dots, \alpha_{L,K})^\top$]. In a slight abuse of notation we write $S(\alpha, z)$ for $S(\mathcal{A}, z)$. Define

$$\begin{aligned} F_{10}(\alpha, z) &= \frac{\partial}{\partial \alpha} S(\alpha, z), & F_{01}(\alpha, z) &= \frac{\partial}{\partial z} S(\alpha, z), \\ F_{20}(\alpha, z) &= \frac{\partial^2}{\partial \alpha^2} S(\alpha, z), & F_{11}(\alpha, z) &= \frac{\partial^2}{\partial \alpha \partial z} S(\alpha, z), \\ F_{02}(\alpha, z) &= \frac{\partial^2}{\partial z^2} S(\alpha, z). \end{aligned}$$

Let $\Psi_t = [\psi(X_{t,1}), \dots, \psi(X_{t,J})]$ be a $K \times J$ matrix. Define A to be the $L \times K$ matrix obtained by deleting the first row of \mathcal{A} . Writing $\zeta_t^\top = (1, z_t^\top)$, it can be shown that

$$\begin{aligned} F_{10}(\alpha, z) &= 2 \sum_{t=1}^T [(\Psi_t \Psi_t^\top) \otimes (\zeta_t \zeta_t^\top)] \alpha - 2 \sum_{t=1}^T (\Psi_t Y_t) \otimes \zeta_t, \\ F_{20}(\alpha, z) &= 2 \sum_{t=1}^T [(\Psi_t \Psi_t^\top) \otimes (\zeta_t \zeta_t^\top)], \end{aligned}$$

$F_{01}(\alpha, z)^\top = 2(\zeta_1^\top \mathcal{A} \Psi_1 \Psi_1^\top A^\top - Y_1^\top \Psi_1^\top A^\top, \dots, \zeta_T^\top \mathcal{A} \Psi_T \Psi_T^\top A^\top - Y_T^\top \Psi_T^\top A^\top)$, and $F_{02}(\alpha, z)$ equals a $(TL) \times (TL)$ matrix that consists of T diagonal blocks $A \Psi_t \Psi_t^\top A^\top$ for $t = 1, \dots, T$. Here and later, \otimes denotes the Kronecker product operator. Also, by some algebraic manipulations it can be shown that

$$[(\Psi_t \Psi_t^\top) \otimes (\zeta_t \zeta_t^\top)] \alpha = (\Psi_t \Psi_t^\top A^\top \zeta_t) \otimes \zeta_t. \quad (6)$$

Let \mathcal{I} be an $(L + 1) \times L$ matrix such that $\mathcal{I}^\top = (0, I_L)$ and I_L denote the identity matrix of dimension L . Define $F_{11,t}(\alpha, z) = (\Psi_t \Psi_t^\top A^\top) \otimes \zeta_t + (\Psi_t \Psi_t^\top A^\top \zeta_t) \otimes \mathcal{I} - (\Psi_t Y_t) \otimes$

\mathcal{I} . Then, we get $F_{11}(\alpha, z) = 2(F_{11,1}(\alpha, z), F_{11,2}(\alpha, z), \dots, F_{11,T}(\alpha, z))$. Let

$$F(\alpha, z) = \begin{pmatrix} F_{10}(\alpha, z) \\ F_{01}(\alpha, z) \end{pmatrix}, F'(\alpha, z) = \begin{pmatrix} F_{20}(\alpha, z) & F_{11}(\alpha, z) \\ F_{11}(\alpha, z)^\top & F_{02}(\alpha, z) \end{pmatrix}.$$

We need to solve the equation $F(\alpha, z) = 0$ simultaneously for α and z . We note that the matrices $(\Psi_t \Psi_t^\top) \otimes (\zeta_t \zeta_t^\top) = (\Psi_t \otimes \zeta_t)(\Psi_t \otimes \zeta_t)^\top$ and $A \Psi_t \Psi_t^\top A^\top$ are nonnegative definite. Thus, by Miranda's existence theorem (for example, see Vrahatis 1989) the nonlinear system of equations $F(\alpha, z) = 0$ has a solution.

Given $(\alpha^{\text{OLD}}, Z^{\text{OLD}})$, the Newton-Raphson algorithm gives the updating equation for $(\alpha^{\text{NEW}}, Z^{\text{NEW}})$:

$$\begin{pmatrix} \alpha^{\text{NEW}} \\ Z^{\text{NEW}} \end{pmatrix} = \begin{pmatrix} \alpha^{\text{OLD}} \\ Z^{\text{OLD}} \end{pmatrix} - F'_*(\alpha^{\text{OLD}}, Z^{\text{OLD}})^{-1} F(\alpha^{\text{OLD}}, Z^{\text{OLD}}), \quad (7)$$

where $F'_*(\alpha, z)$ for each given (α, z) is the restriction to \mathcal{F}_* of the linear map defined by the matrix $F'(\alpha, z)$ and \mathcal{F}_* is the linear space of values of (α, z) with $\sum_{t=1}^T z_t = 0$ and $\sum_{t=1}^T Z_t^{(0)}(z_t - Z_t^{(0)})^\top = 0$. We denote the initial value of the algorithm by $(\alpha^{(0)}, Z^{(0)})$. We will argue later that under mild conditions, $(\hat{\alpha}, \hat{Z})$ can be chosen as an element of \mathcal{F}_* .

The algorithm (7) is shown to converge to a solution of (5) at a geometric rate under some weak conditions on the initial choice $(\alpha^{(0)}, Z^{(0)})$, as is demonstrated by Theorem 1 later. We collect the conditions for the theorem.

(C1) It holds that $\sum_{t=1}^T Z_t^{(0)} = 0$. The matrix $\sum_{t=1}^T Z_t^{(0)} Z_t^{(0)\top}$ and the map $F'_*(\alpha^{(0)}, Z^{(0)})$ are invertible.

(C2) There exists a version $(\hat{\alpha}, \hat{Z})$ with $\sum_{t=1}^T \hat{Z}_t = 0$ such that $\sum_{t=1}^T \hat{Z}_t Z_t^{(0)\top}$ is invertible. Also, $\hat{\alpha}_l = (\hat{\alpha}_{l1}, \dots, \hat{\alpha}_{lK})^\top$ for $l = 0, \dots, L$ are linearly independent.

Let $\alpha^{(k)}$ and $Z^{(k)}$ denote the k th updated vectors in the iteration with the algorithm (7). Also, we write $\mathcal{A}^{(k)}$ for the matrix that corresponds to $\alpha^{(k)}$, and $\mathcal{Z}_t^{(k)\top} = (1, Z_t^{(k)\top})$.

Theorem 1. Let T, J and K be held fixed. Suppose that the initial choice $(\alpha^{(0)}, Z^{(0)})$ satisfies (C1) and (C2). Then, for any constant $0 < \gamma < 1$ there exist $r > 0$ and $C > 0$, which are random variables depending on $\{(X_{t,j}, Y_{t,j})\}$, such that, if $\sum_{t=1}^T \|\mathcal{Z}_t^{(0)\top} \mathcal{A}^{(0)} - \hat{\mathcal{Z}}_t^\top \hat{\mathcal{A}}\|^2 \leq r$, then

$$\sum_{t=1}^T \|\mathcal{Z}_t^{(k)\top} \mathcal{A}^{(k)} - \hat{\mathcal{Z}}_t^\top \hat{\mathcal{A}}\|^2 \leq C 2^{-2(k-1)} \gamma^{2(2^k-1)}.$$

We now argue that under (C1) and (C2), $(\hat{\alpha}, \hat{Z})$ can be chosen as an element of \mathcal{F}_* . Note first that one can always take $Z_t^{(0)}$ and \hat{Z}_t so that $\sum_{t=1}^T Z_t^{(0)} = 0$ and $\sum_{t=1}^T \hat{Z}_t = 0$. This is because, for any version $(\hat{\alpha}, \hat{Z})$, one has

$$\begin{aligned} \hat{\mathcal{Z}}_t^\top \hat{\mathcal{A}} &= \hat{\alpha}_0^\top + \sum_{l=1}^L \hat{Z}_{t,l} \hat{\alpha}_l^\top = \left(\hat{\alpha}_0^\top + \sum_{l=1}^L \bar{\mathcal{Z}}_t \hat{\alpha}_l^\top \right) \\ &+ \sum_{l=1}^L (\hat{Z}_{t,l} - \bar{\mathcal{Z}}_t) \hat{\alpha}_l^\top \stackrel{\text{let}}{=} \hat{\alpha}_0^{*\top} + \sum_{l=1}^L \hat{Z}_{t,l}^* \hat{\alpha}_l^\top = \hat{\mathcal{Z}}_t^{*\top} \hat{\mathcal{A}}^*, \end{aligned}$$

where $\bar{\mathcal{Z}}_t = T^{-1} \sum_{t=1}^T \hat{\mathcal{Z}}_{t,l}$, $\hat{\mathcal{Z}}_t^{*\top} = (1, \hat{\mathcal{Z}}_t^{*\top})$ and $\hat{\mathcal{A}}^*$ is the matrix obtained from $\hat{\mathcal{A}}$ by replacing its first row by $\hat{\alpha}_0^{*\top}$. Furthermore, the minimization problem (5) has no unique solution. If $(\hat{Z}_t, \hat{\mathcal{A}})$ or $(\hat{Z}_t, \hat{m} = \hat{\mathcal{A}}\psi)$ is a minimizer, then also $(B^\top \hat{Z}_t, \tilde{B}^{-1} \hat{m})$ is a minimizer. Here

$$\tilde{B} = \begin{pmatrix} 1 & 0 \\ 0 & B \end{pmatrix} \quad (8)$$

and B is an arbitrary invertible matrix. The special structure of \tilde{B} assures that the first component of $\tilde{B}^\top \hat{Z}_t$ equals 1. In particular, with the choice $B = (\sum_{t=1}^T Z_t^{(0)} \hat{\mathcal{Z}}_t^\top)^{-1} \sum_{t=1}^T Z_t^{(0)} Z_t^{(0)\top}$ we get for $\hat{Z}_t^* = B^\top \hat{Z}_t$ that $\sum_{t=1}^T Z_t^{(0)} (\hat{Z}_t^* - Z_t^{(0)})^\top = 0$.

In Section 5, we will show that, for any solution \hat{Z}_t and for any version of true Z_t , there exists a random matrix B such that $\tilde{Z}_t = B^\top \hat{Z}_t$ has asymptotically the same covariance structure as Z_t . This means that the difference of the inferences based on \tilde{Z}_t and Z_t is asymptotically negligible.

We also note that one can always choose $\hat{m} = \hat{\mathcal{A}}\psi$ such that the components $\hat{m}_1, \dots, \hat{m}_L$ are orthonormal in $L_2([0, 1]^d)$ or in other L_2 [e.g., in $L_2(T^{-1} \sum_{t=1}^T \hat{f}_t)$ where \hat{f}_t is a kernel estimate of the density of $X_{t,j}$]. If one selects \hat{m} in this way, then the matrix B should be an orthogonal matrix and the underlying time series Z_t is estimated up to such transformations.

In practice one needs to choose an initial estimate $(\alpha^{(0)}, Z^{(0)})$ to run the algorithm. One may generate normal random variates for $Z_{t,l}^{(0)}$, and then find the initial $\alpha^{(0)}$ by solving the equation $F_{10}(\alpha, Z^{(0)})$. This initial choice was found to work well in our numerical study presented in Sections 3 and 4.

As an alternative way of fitting the model (2), one may extend the idea of the principal component method that is used to fit the orthogonal factor model (1). In this way, the data $\{Y_{t,j}; 1 \leq j \leq J\}$ are viewed as the values of a functional datum $Y_t(\cdot)$ observed at $x = X_{t,j}$, $1 \leq j \leq J$, and the functional factor model given at (3) may be fitted with smooth approximations of Y_t obtained from the original dataset. If one assumes $E Z_t = 0$, $\text{var}(Z_t) = I_L$, as is typically the case with the orthogonal factor model (1), then one can estimate m_t and Z_t by performing functional principal component analysis with the sample covariance function

$$\hat{K}(x, x') = T^{-1} \sum_{t=1}^T \{Y_t(x) - \bar{Y}(x)\} \{Y_t(x') - \bar{Y}(x')\},$$

where $\bar{Y}(x) = T^{-1} \sum_{t=1}^T Y_t(x)$. There are some limitations for this approach. First, it requires initial fits to get smooth approximations of $Y_t(\cdot)$, which may be difficult when the design points $X_{t,j}$ are sparse as is the case with the IV application. Our method avoids the preliminary estimation and shifts the discrete representation directly to the functions m_t . Second, for the method to work one needs at least stationarity of Z_t and ε_t , whereas our theory does not rely on these assumptions.

3. SIMULATION STUDY

In Theorem 3 we will argue that the inference based on the covariances of the unobserved factors Z_t is asymptotically equivalent to the one based on $B^\top \hat{Z}_t$ for some invertible B . In this section we illustrate the equivalence by a simulation study. We compare the covariances of Z_t and $\tilde{Z}_t \equiv B^\top \hat{Z}_t$, where




Figure 3. The boxplots based on 250 values of the entries of the scaled difference of the covariance matrices given at (10). The lengths of the series Z_t and \tilde{Z}_t were 500, 1,000, 2,000. The thick lines represent the upper and lower quartiles of (11).

$$B = \left(T^{-1} \sum_{t=1}^T Z_{c,t} \hat{Z}_{c,t}^\top \right)^{-1} T^{-1} \sum_{t=1}^T Z_{c,t} Z_{c,t}^\top, \quad (9)$$

$Z_{c,t} = Z_t - T^{-1} \sum_{s=1}^T Z_s$ and $\hat{Z}_{c,t} = \hat{Z}_t - T^{-1} \sum_{s=1}^T \hat{Z}_s$. Note that B at (9) minimizes $\sum_{t=1}^T \| \hat{Z}_{c,t} - (B^\top)^{-1} Z_{c,t} \|^2$. In the Appendix we will prove that Theorem 3 holds with the choice at (9).

We took $T = 500, 1,000, 2,000, J = 100, 250, 1,000$ and $K = 36, 49, 64$. We considered $d = 2, L = 3$ and the following tuple of 2-dimensional functions:

$$\begin{aligned} m_0(x_1, x_2) &= 1, & m_1(x_1, x_2) &= 3.46(x_1 - .5), \\ m_2(x_1, x_2) &= 9.45 \left\{ (x_1 - .5)^2 + (x_2 - .5)^2 \right\} - 1.6, \\ m_3(x_1, x_2) &= 1.41 \sin(2\pi x_2). \end{aligned}$$

The coefficients in these functions were chosen so that m_1, m_2, m_3 are close to orthogonal. We generated Z_t from a centered VAR(1) process $Z_t = \mathcal{R}Z_{t-1} + U_t$, where U_t is $N_3(0, \Sigma_U)$ random vector, the rows of \mathcal{R} from the top equal $(0.95, -0.2, 0), (0, 0.8, 0.1), (0.1, 0, 0.6)$, and $\Sigma_U = 10^{-4}I_3$. The design points $X_{t,j}$ were independently generated from a uniform distribution on the unit square, $\varepsilon_{t,j}$ were iid $N(0, \sigma^2)$ with $\sigma = 0.05$, and $Y_{t,j}$

were obtained according to the model (4). The simulation experiment was repeated 250 times for each combination of (T, J, K) . For the estimation we employed, for ψ_j , the tensor products of linear B-splines. The one-dimensional linear B-splines $\tilde{\psi}_k$ are defined on a consecutive equidistant knots x^k, x^{k+1}, x^{k+2} by $\tilde{\psi}_k(x) = (x - x^k)/(x^{k+1} - x^k)$ for $x \in (x^k, x^{k+1}]$, $\tilde{\psi}_k(x) = (x^{k+2} - x)/(x^{k+2} - x^{k+1})$ for $x \in (x^{k+1}, x^{k+2}]$, and $\tilde{\psi}_k(x) = 0$ otherwise. We chose $K = 8 \times 8 = 64$.

We plotted in Figure 3 the entries of the scaled difference of the covariance matrices

$$\tilde{D} = \frac{1}{\sqrt{T}} \left\{ \sum_{t=1}^T (\tilde{Z}_t - \bar{\tilde{Z}}) (\tilde{Z}_t - \bar{\tilde{Z}})^\top - \sum_{t=1}^T (Z_t - \bar{Z}) (Z_t - \bar{Z})^\top \right\}. \quad (10)$$

Each panel of Figure 3 corresponds to one entry of the matrix \tilde{D} , and the three boxplots in each panel represent the distributions of the 250 values of the corresponding entry for $T = 500, 1,000, 2,000$. In the figure we also depicted, by thick lines, the upper and lower quartiles of

$$D = \frac{1}{\sqrt{T}} \left\{ \sum_{t=1}^T (Z_t - \bar{Z}) (Z_t - \bar{Z})^\top - T I \right\}, \quad (11)$$

where Γ is the true covariance matrix of the simulated VAR process. We refer to Lütkepohl (1993) for a representation of Γ .

Our theory in Section 5 tells that the size of \tilde{D} is of smaller order than the normalized error D of the covariance estimator based on Z_t . It is known that the latter converges to a non-degenerate law as $T \rightarrow \infty$. This is well supported by the plots in Figure 3 showing that the distance between the two thick lines in each panel is almost invariant as T increases. The fact that the additional error incurred by using \tilde{Z}_t instead of Z_t is negligible for large T is also confirmed. In particular, the long stretches at tails of the distributions of \tilde{D} get shorter as T increases. Also, the upper and lower quartiles of each entry of \tilde{D} , represented by the boxes, lie within those of the corresponding entry of D , represented by the thick lines, when $T = 1,000$ and $2,000$.

4. APPLICATIONS

This section presents an application of DSFM. We fit the model to the intraday IV based on ODAX prices and to fMRI data.

For our analysis we chose the data observed from July 1, 2004 to June 29, 2005. The one year period corresponds to the financial regulatory requirements. The data were taken from Financial and Economic Data Center of Humboldt-Universität zu Berlin. The IV data were regressed on the two-dimensional space of future moneyness and time-to-maturity, denoted by $(\kappa_t, \tau_t)^\top$. The future moneyness κ_t is a monotone function of the strike price K : $\kappa_t = K/(S_t e^{-r_t \tau_t})$, where S_t is the spot price at time t and r_t is the interest rate. We chose r_t as a daily Euro Interbank Offered Rate (EURIBOR) taken from the Ecwin Reuters database. The time-to-maturity of the options were measured in years. We took all trades with $10/365 < \tau < 0.5$. We limit also the moneyness range to $\kappa \in [0.7, 1.2]$.

The structure of the IV data, described already in Section 1, requires a careful treatment. Apart from the dynamic degeneration, one may also observe nonuniform frequency of the trades with significantly greater market activities for the options closer to expiry or at-the-money. Here, ‘‘at-the-money’’ means a condition in which the strike price of an option equals the spot price of the underlying security (i.e., $K = S_t$). To avoid the computational problems with the highly skewed empirical distribution of $X_t = (\kappa_t, \tau_t)$, we transformed the initial space $[0.7, 1.2] \times [0.03, 0.5]$ to $[0, 1]^2$ by using the marginal empirical distribution functions. We applied the estimation algorithm to the transformed space, and then transformed back the results to the original space.

Because the model is not nested, the number of the dynamic functions needs to be determined in advance. For this, we used

$$RV(L) = \frac{\sum_t^T \sum_j^{J_t} \left\{ Y_{t,j} - \hat{m}_0(X_{t,j}) - \sum_{l=1}^L \hat{Z}_{t,l} \hat{m}_l(X_{t,j}) \right\}^2}{\sum_t^T \sum_j^{J_t} (Y_{t,j} - \bar{Y})^2}, \quad (12)$$

although one may construct an Akaike information (AIC) or Bayesian information (BIC) type of criterion, where one penalizes the number of the dynamic functions in the model, or performs some type of cross-validation. The quantity $1 - RV(L)$ can be interpreted as a proportion of the variation explained by

Table 1. Proportion of the explained variation by the models with $L = 1, \dots, 5$ dynamic factors

No. factors	$L = 1$	$L = 2$	$L = 3$	$L = 4$	$L = 5$
$1 - RV(L)$	0.848	0.969	0.976	0.978	0.980

the model among the total variation. The computed values of $RV(L)$ are given in Table 1 for various L . Because the third, fourth, and fifth factor made only a small improvement in the fit, we chose $L = 2$.

For the series estimators of \hat{m}_l we used tensor B-splines that are cubic in the moneyness and quadratic in the maturity direction. In the transformed space we placed 10×5 knots, 10 in the moneyness and 5 in the maturity direction. We found that the results were not sensitive to the choice of the number of knots and the orders of splines. For several choices of knots in the range $5 \times 5 - 15 \times 10$ and for the spline orders (2, 1), (2, 2), (3, 2), the values of $1 - RV(2)$ were between 0.949 and 0.974. Because the model is identifiable only up to the transformation (8), one has a freedom for the choice of factors. Here, we chose the approach taken by Fengler et al. (2007) with $L_2[0,1]^2$ norm. Specifically, we orthonormalized \hat{m}_l and transformed \hat{Z}_t according to their Equation (19) with $\Gamma = \int \hat{m}(x) \hat{m}(x)^\top dx$, where $\hat{m} = (\hat{m}_1, \dots, \hat{m}_L)^\top$. Call them \hat{m}_l^* and \hat{Z}_t^* , respectively. Then, we transformed them further by $\hat{m}_l^{**} = p_l^\top \hat{m}_l^*$ and $\hat{Z}_{t,l}^{**} = p_l^\top \hat{Z}_t^*$, where p_l were the orthonormal eigenvectors of the matrix $\sum_{t=1}^T \hat{Z}_t^* \hat{Z}_t^{*\top}$ that correspond to the eigenvalues $\lambda_1 > \lambda_2$. Note that $\hat{Z}_t^{*\top} \hat{m}^* = \hat{Z}_t^{**\top} \hat{m}^{**}$. In this way, $\{\hat{Z}_{t,1}^{**} \hat{m}_1^{**}\}$ makes a larger contribution than $\{\hat{Z}_{t,2}^{**} \hat{m}_2^{**}\}$ to the total variation $\sum_{t=1}^T \int (\hat{Z}_t^{**\top} \hat{m}^{**})^2$ because $\sum_{t=1}^T \int (\hat{Z}_{t,1}^{**} \hat{m}_1^{**})^2 = \lambda_1$ and $\sum_{t=1}^T \int (\hat{Z}_{t,2}^{**} \hat{m}_2^{**})^2 = \lambda_2$. Later, we continue to write \hat{Z}_t and \hat{m} for such \hat{Z}_t^{**} and \hat{m}^{**} , respectively.

The estimated functions \hat{m}_1 and \hat{m}_2 are plotted in Figure 4 in the transformed estimation space. The intercept function \hat{m}_0 was almost flat around zero, thus is not given. By construction, $\hat{m}_0 + \hat{Z}_{t,1} \hat{m}_1$ explain the principal movements of the surface. It was observed by Cont and da Fonseca (2002) and Fengler et al. (2007) that most dominant innovations of the entire surface are parallel level shifts. Note that VDAX is an estimated at-the-money IV for an option with 45 days to maturity, and thus indicates up-and-down shifts. The left panel of Figure 5 shows the values of VDAX together with $\hat{m}_0(X_{t,0}) + \hat{Z}_{t,1} \hat{m}_1(X_{t,0})$, where $X_{t,0}$ is the moneyness and maturity corresponding to an option at-the-money with 45 days to maturity. The right panel of Figure 5 depicts the factor \hat{Z}_t , where one can find that \hat{Z}_t shows almost the same dynamic behavior as the index VDAX. This similarity supports that DSFM catches leading dynamic effects successfully. Obviously the model in its full setting explains other effects, such as skew or term structure changes, which are not explicitly stated here.

Statistical analysis on the evolution of a high-dimensional system ruling the option prices can be simplified to a low-dimensional analysis of the \hat{Z}_t . In particular, as our theory in Section 5 and the simulation results in Section 3 assert, the inference based on the \hat{Z}_t is well justified in the VAR context. To select a VAR model we computed the Schwarz (SC), the Hannan-Quinn (HQ), and the Akaike criterion, as given in




Figure 4. The estimated factor functions for the ODAX IV data in the period 20040701–20050629.

Table 2. One can find that SC and HQ suggest a VAR(1) process, whereas AIC selects VAR(2). The parameter estimates for each selected model are given in Table 3. The roots of the characteristic polynomial lie inside the unit circle, so the specified models satisfy the stationarity condition. For each of VAR(1) and VAR(2) models, we conducted a portmanteau test for the hypothesis that the autocorrelations of the error term at lags up to 12 are all zero, and also a series of LM tests, each of which tests whether the autocorrelation at a particular lag up to 5 equals zero. Some details on selection of lags for these tests can be found in Hosking (1980, 1981) and Brüggemann, Lütkepohl, and Saikkonen (2006). We found that in any test the null hypothesis was not rejected at 5% level. A closer inspection on the autocorrelations of the residuals, however, revealed that the autocorrelation of $\widehat{Z}_{t,2}$ residuals at lag one is slightly significant in the VAR(1) model, see Figure 6. But, this effect disappears in the VAR(2) case, see Figure 7. Similar analyses of characteristic polynomials, portmanteau and Lagrange multiplier (LM) tests supported VAR(2) as a successful model for \widehat{Z}_t .

As a second application of the model, we considered fitting an fMRI dataset. The data were obtained at Max-Planck Institut für Kognitions-und-Neurowissenschaften Leipzig by scanning a subject’s brain using a standard head coil. The scanning was done every two seconds on the resolution of $3 \times 3 \times 2 \text{ mm}^3$ with 1 mm gap between the slices. During the experiment, the

subject was exposed to three types of objects (bench, phone and motorbike) and rotated around randomly changing axes for four seconds, followed by relaxation phase of six to ten seconds. Each stimulus was shown 16 times in pseudo-randomized order. As a result, a series of 290 images with $64 \times 64 \times 30$ voxels was obtained.

To apply the model (2) to the fMRI data, we took the voxel’s index (i_1, i_2, i_3) as covariate $X_{t,j}$, and the BOLD signal as $Y_{t,j}$. For numerical tractability we reduced the original data to a series of $32 \times 32 \times 15$ voxels by taking every second slice in each direction. Thus, $J_t \equiv 32 \times 32 \times 15$ and $T = 290$. The voxels’ indices (i_1, i_2, i_3) for $1 \leq i_1, i_2 \leq 32 ; 1 \leq i_3 \leq 15$ are associated with $32 \times 32 \times 15$ equidistant points in \mathbb{R}^3 . The function m_0 represents the “average” signal as a function of the three-dimensional location, and m_l for each $l \geq 1$ determines the effect of the l th common factor $Z_{t,l}$ on the brain’s signal. In Figure 8, each estimated function \widehat{m}_l is represented by its sections on the 15 slices in the direction of i_3 [i.e., by those $\widehat{m}_l(\cdot, \cdot, x_3)$ for which x_3 are fixed at the equidistant points corresponding to $i_3 = 1, \dots, 15$]. We used quadratic tensor B-splines on equidistant knots. The number of knots in each direction was 8, 8, 4, respectively, so that $K = 9 \times 9 \times 5 = 405$. For the model identification we used the same method as in the IV application, but normalized \widehat{Z} to have mean zero.

In contrast to the IV application, there was no significant difference between the values of $1 - RV(L)$ for different $L \geq 1$.




Figure 5. Left panel: VDX in the period 20040701–20050629 (solid) and the dynamics of the corresponding IV given by the submodel $\widehat{m}_0 + \widehat{Z}_{t,1}\widehat{m}_1$ (dashed). Right panel: The obtained time series \widehat{Z}_t on the ODAX IV data in the period 20040701–20050629. The solid line represents $\widehat{Z}_{t,1}$, the dashed line $\widehat{Z}_{t,2}$.

Table 2. The VAR model selection criteria. The smallest value for each criterion is marked by an asterisk

Order	AIC	SC	HQ
1	-14.06	-13.98*	-14.03*
2	-14.07*	-13.93	-14.02
3	-14.06	-13.86	-13.98
4	-14.06	-13.81	-13.96
5	-14.07	-13.76	-13.95

All the values for $L \geq 1$ were around 0.871. The fMRI signals $Y_{t,j}$ were explained mostly by $\hat{m}_0(X_{t,j}) + Z_{t,1}\hat{m}_1(X_{t,j})$, and the effects of the common factors $Z_{t,l}$ for $l \geq 2$ were relatively small. The slow increase in the value of $1 - RV(L)$ as $L \geq 1$ grows in the fMRI application, contrary to the case of the IV application, can be explained partly by the high complexity of human brain. Because the values of $1 - RV(L)$ were similar for $L \geq 1$, one might choose $L = 1$. However, we chose $L = 4$, which we think still allows relatively low complexity, to demonstrate some further analysis that might be possible with similar datasets. The estimated functions \hat{m}_l for $0 \leq l \leq 4$ and the time series $\hat{Z}_{t,l}$ for $1 \leq l \leq 4$ are plotted in Figures 8 and 9, respectively. The function \hat{m}_0 can be recognized as a smoothed version of the original signal. By construction the first factor and loadings incorporate the largest variation. One may see the strong positive trend in $\hat{Z}_{t,1}$ and relatively flat patterns of $\hat{Z}_{t,2}, \hat{Z}_{t,3}, \hat{Z}_{t,4}$. These effects could be typically explained by the mixture of several components, such as physiological pulsation, subtle head movement, machine noise, and so on. For a description of different artifacts, which significantly influence the fMRI signals, we refer to Biswal, Yetkin, Haughton, and Hyde (1995). The function estimates \hat{m}_l for $1 \leq l \leq 4$ appear to have a clear peak, and $\hat{Z}_{t,l}$ for $2 \leq l \leq 4$ show rather mild mean reverting behavior.

To see how the recovered signals interact with the given stimuli, we plotted $\hat{Z}_{t+s,l} - \hat{Z}_{s,l}$ against t in Figure 10, where s is the time when a stimulus appears. The mean changes of $\hat{Z}_{t,1}$ and $\hat{Z}_{t,3}$ show mild similarity, up to sign change, to the hemodynamic response (see Worsley et al. 2002). The case of $\hat{Z}_{t,4}$ has a similar pattern as those of $\hat{Z}_{t,1}$ and $\hat{Z}_{t,3}$ but with larger amplitude, whereas the changes in $\hat{Z}_{t,2}$ seem to be independent of the stimuli. In fitting the fMRI data, we did not use any external information on the signal. From the biological perspective it could be hardly expected that a pure statistical procedure gives full insight into understanding of the complex dynamics of MR images. For the latter one needs to incorporate into the procedure the shape of hemodynamic response, for example, or consider physiologically motivated identification of the fac-

tors. It goes however beyond the scope of this illustrative example.

5. ASYMPTOTIC ANALYSIS

In the simulation study and the real data application in Sections 3 and 4, we considered the case where Z_t is a VAR-process. Here, we only make some weak assumptions on the average behavior of the process. In our first theorem we allow that it is a deterministic sequence. In our second result we assume that it is a mixing sequence. For the asymptotic analysis, we let $K, J, T \rightarrow \infty$. This is a very natural assumption often also made in cross-sectional or panel data analysis. It is appropriate for data with many observations per data point that are available for many dates. It allows us to study how J and T have to grow with respect to each other for a good performance of a procedure. The distance between m and its best approximation $\mathcal{A}\psi$ does not tend to zero unless $K \rightarrow \infty$, see Assumption (A5) later. One needs to let $J \rightarrow \infty$ to get consistency of $\hat{Z}_t^\top \hat{A}$ and $\hat{m} = \hat{A}\psi$ as estimates of $Z_t^\top \mathcal{A}^*$ and m , respectively, where \mathcal{A}^* is defined at (A5). One should let $T \rightarrow \infty$ to describe the asymptotic equivalence between the lagged covariances of Z_t and those of \hat{Z}_t , see Theorem 3 below. In our analysis the dimension L is fixed. Clearly, one could also study our model with L growing to infinity. We treat the case where X_{it} are random. However, a theory for deterministic designs can be developed along the lines of our theory.

Our first result relies on the following assumptions.

- (A1) The variables $X_{1,1}, \dots, X_{T,J}, \varepsilon_{1,1}, \dots, \varepsilon_{T,J}$, and Z_1, \dots, Z_T are independent. The process Z_t is allowed to be nonrandom.
- (A2) For $t = 1, \dots, T$ the variables $X_{t,1}, \dots, X_{t,J}$ are identically distributed, have support $[0, 1]^d$ and a density f_t that is bounded from below and above on $[0, 1]^d$, uniformly over $t = 1, \dots, T$.
- (A3) We assume that $E\varepsilon_{t,j} = 0$ for $1 \leq t \leq T, 1 \leq j \leq J$, and for $c > 0$ small enough $\sup_{1 \leq t \leq T, 1 \leq j \leq J} E \exp(c\varepsilon_{t,j}^2) < \infty$.
- (A4) The functions ψ_k may depend on the increasing indices T and J , but are normed so that $\int_{[0,1]^d} \psi_k^2(x) dx = 1$ for $k = 1, \dots, K$. Furthermore, it holds that $\sup_{x \in [0,1]^d} \|\psi(x)\| = \mathcal{O}(K^{1/2})$.
- (A5) The vector of functions $m = (m_0, \dots, m_L)^\top$ can be approximated by ψ_k , i.e.,

$$\delta_K \equiv \sup_{x \in [0,1]^d} \inf_{\mathcal{A} \in \mathbb{R}^{(L+1) \times K}} \|m(x) - \mathcal{A}\psi(x)\| \rightarrow 0$$

as $K \rightarrow \infty$. We denote \mathcal{A} that fulfills $\sup_{x \in [0,1]^d} \|m(x) - \mathcal{A}\psi(x)\| \leq 2\delta_K$ by \mathcal{A}^* .

- (A6) There exist constants $0 < C_L < C_U < \infty$ such that all eigenvalues of the matrix $T^{-1} \sum_{t=1}^T Z_t Z_t^\top$ lie in the interval $[C_L, C_U]$ with probability tending to one.

Table 3. The estimated parameters for VAR(1) and VAR(2) models. Those that are not significant at 5% level are marked by asterisk

	VAR(1)			VAR(2)				
	$\hat{Z}_{t-1,1}$	$\hat{Z}_{t-1,2}$	Const.	$\hat{Z}_{t-1,1}$	$\hat{Z}_{t-1,2}$	$\hat{Z}_{t-2,1}$	$\hat{Z}_{t-2,2}$	Const.
$\hat{Z}_{t,1}$	0.984	-0.029*	-0.001	0.913	-0.025	0.071	-0.004	-0.001
$\hat{Z}_{t,2}$	0.055	0.739	0.005	0.124	0.880	-0.065	-0.187*	0.006

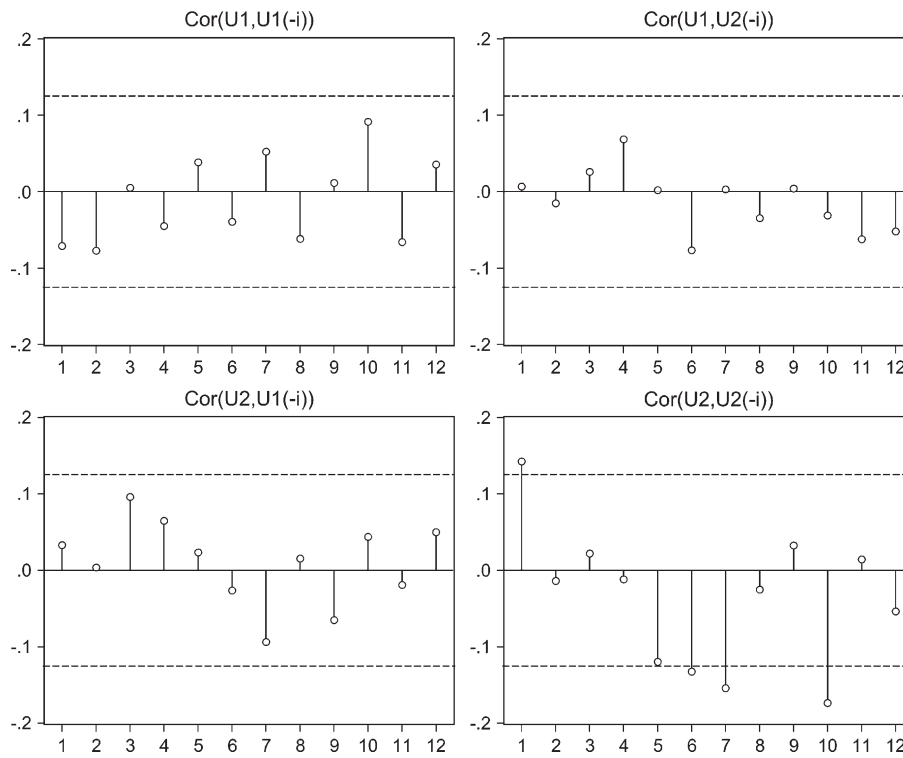


Figure 6. Cross-autocorrelogram for the VAR(1) residuals. The dashed line-bounds indicate $\pm 2 \times$ (standard deviations), which correspond to an approximate 95% confidence bound.

(A7) The minimization (5) runs over all values of (\mathcal{A}, z) with

$$\sup_{x \in [0,1]} \max_{1 \leq t \leq T} \| (1, z_t^T) \mathcal{A} \psi(x) \| \leq M_T,$$

where the constant M_T fulfils $\max_{1 \leq t \leq T} \|Z_t\| \leq M_T/C_m$ (with probability tending to one) for a constant C_m such that $\sup_{x \in [0, 1]} \|m(x)\| < C_m$.




Figure 7. Cross-autocorrelogram for the VAR(2) residuals. The dashed line-bounds indicate $\pm 2 \times$ (standard deviations), which correspond to an approximate 95% confidence bound.

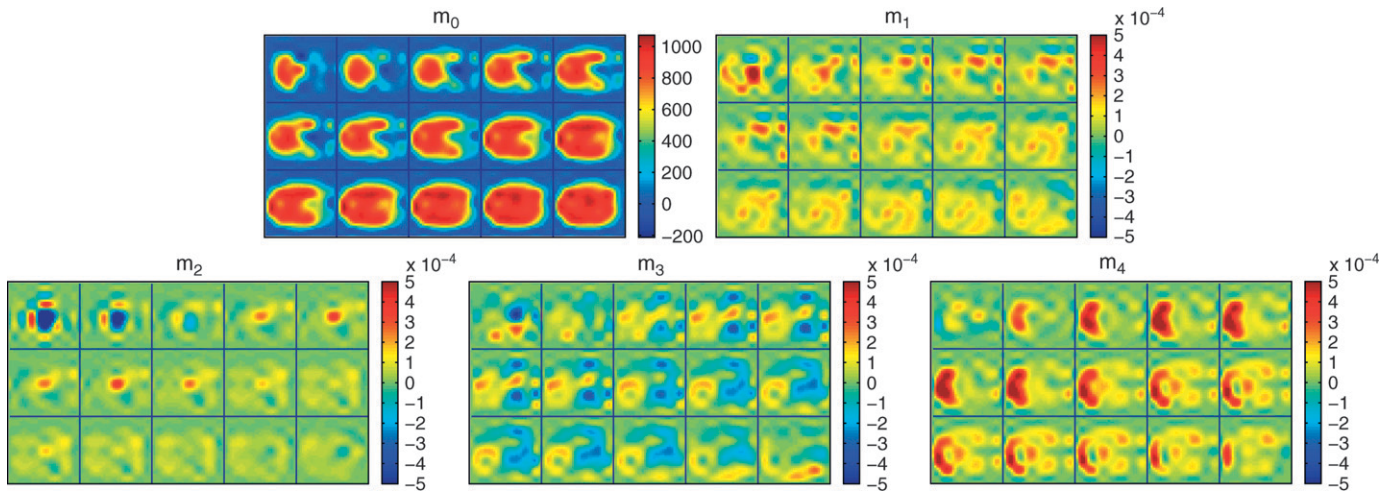


Figure 8. The estimated functions \hat{m}_i for the fMRI signals.

(A8) It holds that $\rho^2 = (K + T)M_T^2 \log(JTM_T)/(JT) \rightarrow 0$. The dimension L is fixed.

Assumption (A7) and the additional bound M_T in the minimization is introduced for purely technical reasons. We conjecture that to some extent the asymptotic theory of this article could be developed under weaker conditions. The independence assumptions in (A1) and Assumption (A3) could be relaxed to assuming that the errors $\epsilon_{t,j}$ have a conditional mean zero and have a conditional distribution with subgaussian tails, given the past values $X_{s,i}, Z_s$ ($1 \leq i \leq J, 1 \leq s \leq t$). Such a theory would

require an empirical process theory that is more explicitly designed for our model and it would also require a lot of more technical assumptions. We also expect that one could proceed with the assumption of subexponential instead of subgaussian tails, again at the cost of some additional conditions. Recall that the number of parameters to be estimated equals $TL + K(L + 1)$. Because L is fixed, Assumption (A8) requires basically that, neglecting the factor $M_T^2 \log(JTM_T)$, the number of parameters grows slower than the number of observations, JT .

Our first result gives rates of convergence for the least squares estimators \hat{Z}_t and \hat{A} .




Figure 9. The estimated time series $\hat{Z}_{t,l}$ for the fMRI signals.

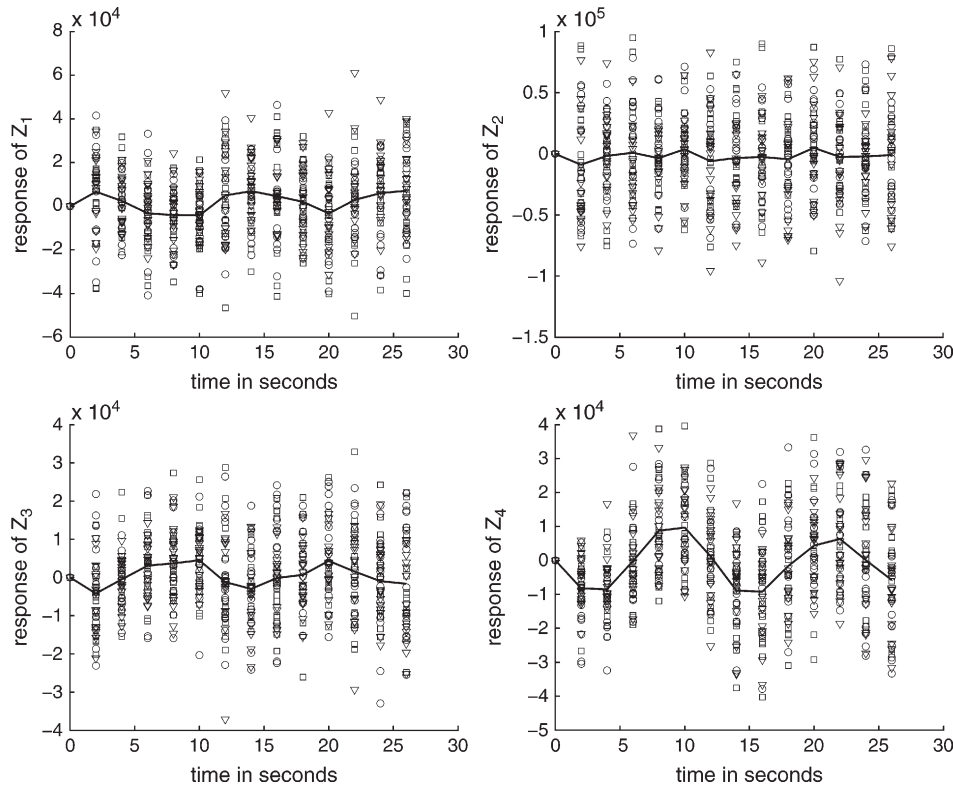


Figure 10. The responses of $\hat{Z}_{t,t}$ to the stimuli.

Theorem 2. Suppose that model (4) holds and that (\hat{Z}_t, \hat{A}) is defined by the minimization problem (5). Make the Assumptions (A1)–(A8). Then it holds that

$$\frac{1}{T} \sum_{1 \leq t \leq T} \left\| \hat{Z}_t^\top \hat{A} - Z_t^\top A^* \right\|^2 = \mathcal{O}_P(\rho^2 + \delta_K^2). \quad (13)$$

At this point we have made no assumptions on the sequence Z_t , $1 \leq t \leq T$, besides the bound in (A7). Up to now it is allowed to be a deterministic or a random sequence. We now assume that it is a random process. We discuss how a statistical analysis differs if inference on Z_t is based on \hat{Z}_t instead of using (the unobserved) process Z_t . We will show that the differences are asymptotically negligible (except an orthogonal transformation). This is the content of the following theorem, where we consider estimators of autocovariances and show that these estimators differ only by second order terms. This asymptotic equivalence carries over to classical estimation and testing procedures in the framework of fitting a vector autoregressive model. For the statement of the theorem we need the following assumptions:

- (A9) Z_t is a strictly stationary sequence with $E(Z_t) = 0$, $E(\|Z_t\|^\gamma) < \infty$ for some $\gamma > 2$. It is strongly mixing with $\sum_{i=1}^\infty \alpha(i)^{(\gamma-2)/\gamma} < \infty$. The matrix $E Z_t Z_t^\top$ has full rank. The process Z_t is independent of $X_{11}, \dots, X_{TJ}, \epsilon_{11}, \dots, \epsilon_{TJ}$.
- (A10) The functions m_0, \dots, m_L are linearly independent. In particular, no function is equal to 0.
- (A11) It holds that $[\log(KT)^2 \{ (KM_T/J)^{1/2} + T^{1/2} M_T^4 J^{-2} + K^{3/2} J^{-1} + K^{4/3} J^{-2/3} T^{-1/6} \} + 1] T^{1/2} (\rho^2 + \delta_K^2) = o(1)$.

Assumption (A11) poses very weak conditions on the growth of J , K , and T . Suppose, for example, that M_T is of logarithmic

order and that K is of order $(TJ)^{1/5}$ so that the variance and the bias are balanced for twice differentiable functions. In this setting, (A11) only requires that T/J^2 times a logarithmic factor converges to zero. Define $\tilde{Z}_t = B^\top \hat{Z}_t$,

$$\begin{aligned} \tilde{Z}_{c,t} &= \tilde{Z}_t - T^{-1} \sum_{s=1}^T \tilde{Z}_s, \\ Z_{c,t} &= Z_t - T^{-1} \sum_{s=1}^T Z_s, \\ \tilde{Z}_{n,t} &= (T^{-1} \sum_{s=1}^T \tilde{Z}_{c,s} \tilde{Z}_{c,s}^\top)^{-1/2} \tilde{Z}_{c,t}, \end{aligned}$$

$$\text{and } Z_{n,t} = (T^{-1} \sum_{s=1}^T Z_{c,s} Z_{c,s}^\top)^{-1/2} Z_{c,t}.$$

Theorem 3. Suppose that model (4) holds and that (\hat{Z}_t, \hat{A}) is defined by the minimization problem (5). Make the Assumptions (A1)–(A11). Then there exists a random matrix B such that for $h \neq 0$

$$\begin{aligned} \frac{1}{T} \sum_{t=\max[1,-h+1]}^{\min[T,T-h]} \tilde{Z}_{c,t} (\tilde{Z}_{c,t+h} - \tilde{Z}_{c,t})^\top - \frac{1}{T} \sum_{t=\max[1,-h+1]}^{\min[T,T-h]} Z_{c,t} \\ (Z_{c,t+h} - Z_{c,t})^\top &= o_P(T^{-1/2}), \\ \frac{1}{T} \sum_{t=\max[1,-h+1]}^{\min[T,T-h]} \tilde{Z}_{n,t} \tilde{Z}_{n,t+h}^\top - \frac{1}{T} \sum_{t=\max[1,-h+1]}^{\min[T,T-h]} Z_{n,t} Z_{n,t+h}^\top &= o_P(T^{-1/2}). \end{aligned}$$

To illustrate an implication of Theorem 3, suppose that the factor process Z_t in (4) is a stationary VAR(p) process in a mean adjusted form:

$$Z_t - \mu = \Theta_1(Z_{t-1} - \mu) + \dots + \Theta_p(Z_{t-p} - \mu) + U_t, \quad (14)$$

where $\mu = E(Z_t)$, Θ_j is a $L \times L$ matrix of coefficients and U_t is a white noise with a nonsingular covariance matrix. Let Γ_h be the autocovariance matrix of the process Z_t with the lag $h \geq 0$, which is estimated by $\hat{\Gamma}_h = T^{-1} \sum_{t=h+1}^T (Z_t - \bar{Z})(Z_{t-h} - \bar{Z})^\top$. Let $Y = (Z_{p+1} - \mu, \dots, Z_T - \mu)$, $\Theta = (\Theta_1, \dots, \Theta_p)$, and $U = (U_{p+1}, \dots, U_T)$. Define $W_t = ((Z_t - \mu)^\top, \dots, (Z_{t-p+1} - \mu)^\top)^\top$ and $W = (W_p, \dots, W_{T-1})$. Then, the model (14) can be rewritten as $Y = \Theta W + U$ and the least squares estimator of Θ is given by $\hat{\Theta} = \hat{Y} \hat{W}^\top (\hat{W} \hat{W}^\top)^{-1}$, where \hat{Y} and \hat{W} are the same as Y and W , respectively, except that μ is replaced by \bar{Z} . Likewise, fitting a VAR(p) model with the estimated factor process \tilde{Z}_t yields $\hat{\Theta} = \hat{Y} \hat{W}^\top (\hat{W} \hat{W}^\top)^{-1}$, where \hat{Y} and \hat{W} are defined as \hat{Y} and \hat{W} with Z_t being replaced by \tilde{Z}_t . Both \hat{Y} and \hat{W} are matrices composed of $\hat{\Gamma}_h$ for various h . The matrices \hat{Y} and \hat{W} have the same forms as \hat{Y} and \hat{W} , respectively, but with $\hat{\Gamma}_h$ being replaced by $\hat{\Gamma}_h = T^{-1} \sum_{t=h+1}^T (\tilde{Z}_t - \bar{\tilde{Z}})(\tilde{Z}_{t-h} - \bar{\tilde{Z}})^\top$. It is well known that $\sqrt{T}(\hat{\Theta} - \Theta) = \mathcal{O}_P(1)$, see Lütkepohl (1993). By Theorem 3, we have $\sqrt{T}(\hat{\Theta} - \Theta) = o_P(1)$.

APPENDIX: PROOFS OF THEOREMS

A.1 Proof of Theorem 1

We use the Newton-Kantorovich theorem to prove the theorem. The statement of the theorem may be found in Kantorovich and Akilov (1982), for example.

Suppose that $\sum_{t=1}^T \|\mathcal{Z}_t^{(0)\top} \mathcal{A}^{(0)} - \hat{\mathcal{Z}}_t^\top \hat{\mathcal{A}}\|^2 \leq r$ for some $r > 0$, which will be chosen later. With the Frobenius norm $\|M\|$ for a matrix M , we get

$$\begin{aligned} \|\mathcal{A}^{(0)} - \hat{\mathcal{A}}\|^2 &\leq \left\| \left(\sum_{t=1}^T \mathcal{Z}_t^{(0)} \mathcal{Z}_t^{(0)\top} \right)^{-1} \right\|^2 \\ &\quad \left\| \sum_{t=1}^T \mathcal{Z}_t^{(0)} \mathcal{Z}_t^{(0)\top} (\mathcal{A}^{(0)} - \hat{\mathcal{A}}) \right\|^2 \\ &= \left\| \left(\sum_{t=1}^T \mathcal{Z}_t^{(0)} \mathcal{Z}_t^{(0)\top} \right)^{-1} \right\|^2 \cdot \left\| \sum_{t=1}^T \mathcal{Z}_t^{(0)} \mathcal{Z}_t^{(0)\top} \mathcal{A}^{(0)} \right. \\ &\quad \left. - \sum_{t=1}^T \mathcal{Z}_t^{(0)} \hat{\mathcal{Z}}_t^\top \hat{\mathcal{A}} \right\|^2 \leq \left\| \left(\sum_{t=1}^T \mathcal{Z}_t^{(0)} \mathcal{Z}_t^{(0)\top} \right)^{-1} \right\|^2 \\ &\quad \times \left(\sum_{t=1}^T \left\| \mathcal{Z}_t^{(0)} \mathcal{Z}_t^{(0)\top} \mathcal{A}^{(0)} - \mathcal{Z}_t^{(0)} \hat{\mathcal{Z}}_t^\top \hat{\mathcal{A}} \right\|^2 \right) \leq \\ &\quad r \left\| \left(\sum_{t=1}^T \mathcal{Z}_t^{(0)} \mathcal{Z}_t^{(0)\top} \right)^{-1} \right\|^2 \left(\sum_{t=1}^T \|\mathcal{Z}_t^{(0)}\|^2 \right) \\ &= rc_1^2. \end{aligned} \quad (A.1)$$

For a matrix M , define $\|M\|_2 = \sup_{\|x\|=1} \|Mx\|$. It is known that $\|M\|_2 \leq \|M\|$. We get

$$\begin{aligned} \|\hat{\mathcal{A}}^\top (\mathcal{Z}_t^{(0)} - \hat{\mathcal{Z}}_t)\| &\geq \|\hat{\mathcal{A}}\|_2^{-1} \\ &\quad \cdot \|(\hat{\mathcal{A}} \hat{\mathcal{A}}^\top)^{-1}\|^{-1} \cdot \|\mathcal{Z}_t^{(0)} - \hat{\mathcal{Z}}_t\|, \end{aligned} \quad (A.2)$$

$$\begin{aligned} \|(\mathcal{Z}_t^{(0)} - \hat{\mathcal{Z}}_t)^\top \hat{\mathcal{A}}\| &\leq \|\mathcal{Z}_t^{(0)\top} (\hat{\mathcal{A}} - \mathcal{A}^{(0)})\| + \|\mathcal{Z}_t^{(0)\top} \mathcal{A}^{(0)} - \\ \hat{\mathcal{Z}}_t^\top \hat{\mathcal{A}}\| &\leq \|\mathcal{Z}_t^{(0)}\| \cdot \|\hat{\mathcal{A}} - \mathcal{A}^{(0)}\| + \|\mathcal{Z}_t^{(0)\top} \mathcal{A}^{(0)} - \hat{\mathcal{Z}}_t^\top \hat{\mathcal{A}}\|. \end{aligned} \quad (A.3)$$

The two inequalities (A.2) and (A.3) together with (A.1) give

$$\begin{aligned} \|\mathcal{Z}^{(0)} - \hat{\mathcal{Z}}\|^2 &\leq 2r \|\hat{\mathcal{A}}\|_2^2 \cdot \|(\hat{\mathcal{A}} \hat{\mathcal{A}}^\top)^{-1}\|^2 \\ &\quad \times \left(1 + c_1 \sum_{t=1}^T \|\mathcal{Z}_t^{(0)}\|^2 \right) = rc_2^2. \end{aligned} \quad (A.4)$$

Because $F'(\alpha, z)$ is quadratic in (α, z) , there exists $0 < c_3 < \infty$ for any compact set D in $\mathbb{R}^{K(L+1)+TL}$ such that $\|F'(\alpha', z') - F'(\alpha, z)\|_2 \leq c_3 \|(\alpha'^\top, z'^\top)^\top - (\alpha^\top, z^\top)^\top\|$ for all $(\alpha'^\top, z'^\top)^\top, (\alpha^\top, z^\top)^\top \in D$. Let $c_4 = \|F'_*(\alpha^{(0)}, Z^{(0)})^{-1}\|_2 < \infty$. Because F is continuous and $F(\hat{\alpha}, \hat{Z}) = 0$, there exists $r' > 0$ such that, if $\|\alpha^{(0)} - \hat{\alpha}\| + \|Z^{(0)} - \hat{Z}\| \leq r'$, then

$$\|F'_*(\alpha^{(0)}, Z^{(0)})^{-1} F(\alpha^{(0)}, Z^{(0)})\| \leq \frac{\gamma}{2c_3 c_4}.$$

By the Newton-Kantorovich theorem,

$$\|\alpha^{(k)} - \hat{\alpha}\| + \|Z^{(k)} - \hat{Z}\| \leq C_1 2^{-(k-1)} \gamma^{2^{k-1}} \quad (A.5)$$

for some $C_1 > 0$. This gives that if $\|\alpha^{(0)} - \hat{\alpha}\| + \|Z^{(0)} - \hat{Z}\| \leq r'$, then

$$\begin{aligned} \sum_{t=1}^T \|\mathcal{Z}_t^{(k)\top} \mathcal{A}^{(k)} - \hat{\mathcal{Z}}_t^\top \hat{\mathcal{A}}\|^2 &\leq C_2 (\|\alpha^{(k)} - \hat{\alpha}\|^2 + \\ \|\mathcal{Z}^{(k)} - \hat{\mathcal{Z}}\|^2) &\leq C_2^{-2(k-1)} \gamma^{2^{2(k-1)}} \end{aligned}$$

for some $C, C_2 > 0$. We take $r = (c_1 + c_2)^{-2} r'^2$. Then, by (A.1) and (A.4), $\|\alpha^{(0)} - \hat{\alpha}\| + \|Z^{(0)} - \hat{Z}\| \leq r'$ if $\sum_{t=1}^T \|\mathcal{Z}_t^{(0)\top} \mathcal{A}^{(0)} - \hat{\mathcal{Z}}_t^\top \hat{\mathcal{A}}\|^2 \leq r$. This completes the proof of the theorem.

A.2 Proof of Theorem 2

For functions $g(t, x)$ we define the norms $\|g\|_1^2 = (1/TJ) \sum_{t=1}^T \sum_{j=1}^J g(t, X_{t,j})^2$, $\|g\|_2^2 = (1/T) \sum_{t=1}^T \int g(t, x)^2 f_t(x) dx$, and $\|g\|_3^2 = (1/T) \sum_{t=1}^T \int g(t, x)^2 dx$. Note that because of Assumption (A2) the last two norms are equivalent. Thus, for the statement of the theorem we have to show for $\Delta(t, x) = (\hat{\mathcal{Z}}_t^\top \hat{\mathcal{A}} - \mathcal{Z}_t^\top \mathcal{A}^*) \psi(x)$ that

$$\|\Delta\|_2^2 = \mathcal{O}_P(\rho^2 + \delta_K^2). \quad (A.6)$$

We start by showing that

$$\|\Delta\|_1^2 = \mathcal{O}_P([(K+T) \log(JTM_T)]/(JT) + \delta_K^2). \quad (A.7)$$

For this aim we apply Theorem 10.11 in Van de Geer (2000) that treats rates of convergence for least squares estimators on sieves. In our case we have the following sieve: $\mathcal{G}_T^* = \{g: \{1, \dots, T\} \times [0, 1]^d \rightarrow \mathbb{R}, g(t, x) = (1, z_t^\top) \mathcal{A} \psi(x) \text{ for an } (L+1) \times K \text{ matrix } \mathcal{A} \text{ and } z_t \in \mathbb{R}^L \text{ with the following properties: } |(1, z_t^\top) \mathcal{A} \psi(x)| \leq M_T \text{ for } 1 \leq t \leq T \text{ and } x \in [0, 1]^d\}$. With a constant C the δ -entropy $H_T(\delta, \mathcal{G}_T^*)$ of \mathcal{G}_T^* with respect to the empirical norm $\|g\|_1$ is bounded by

$$H_T(\delta, \mathcal{G}_T^*) \leq CT \log(M_T/\delta) + CK \log(KM_T/\delta). \quad (A.8)$$

For the proof of (A.8) note first that each element $g(t, x) = (1, z_t^\top) \mathcal{A}\psi(x)$ of \mathcal{G}_T^* can be chosen such that $T^{-1} \sum_{t=1}^T z_t z_t^\top$ is equal to the $L \times L$ identity matrix I_L . Then the bound $|(1, z_t^\top) \mathcal{A}\psi(x)| \leq M_T$ implies that $\|\mathcal{A}\psi(x)\| \leq M_T$. For the proof of (A.8) we use that the (δ/M_T) -entropy of a unit ball in \mathbb{R}^T is of order $\mathcal{O}(T \log(M_T/\delta))$ and that the δ -entropy with respect to the sup-norm for functions $\mathcal{A}\psi(x)$ with $\|\mathcal{A}\psi(x)\| \leq M_T$ is of order $\mathcal{O}(K \log(KM_T/\delta))$. In the last entropy bound we used that for each x it holds that $\|\psi(x)\| \leq K^{1/2}$. These two entropy bounds imply (A.8). Application of Theorem 10.11 in Van de Geer (2000) gives (A.7).

We now show that (A.7) implies (A.6). For this aim note first that by Bernstein's inequality for $a, d > 0, g \in \mathcal{G}_T^*$ with $\|g\|_2^2 \leq d$

$$P(|\|g\|_1^2 - \|g\|_2^2| \geq a) \leq 2 \exp\left(-\frac{a^2 JT}{2(a+d)M_T^2}\right). \quad (\text{A.9})$$

Furthermore, for $g, h \in \mathcal{G}_T^*$ it holds with constants C, C' that

$$\begin{aligned} \|g\|_1^2 - \|h\|_1^2 &\leq CK \left(T^{-1} \sum_{t=1}^T \|e_t - f_t\|^2 \right)^{1/2} \\ \left(T^{-1} \sum_{t=1}^T \|e_t + f_t\|^2 \right)^{1/2} &\leq C'K \|g - h\|_2 (\|g\|_2 + \|h\|_2), \end{aligned} \quad (\text{A.10})$$

where e_t and f_t are chosen such that $g(x, t) = e_t^\top \psi(x)$ and $h(x, t) = f_t^\top \psi(x)$. From (A.9) and (A.10) we get with a constant $C > 0$ that for $d = 1, 2, \dots$

$$\begin{aligned} P\left(\sup_{g \in \mathcal{G}_T^*, d\rho^2 \leq \|g\|_2^2 \leq (d+1)\rho^2} |\|g\|_1^2 - \|g\|_2^2| \geq d\rho^2/2\right) \\ \leq C \exp\left((C + K + T) \log(dKM_T) - d\rho^2 JT/[20M_T^2]\right). \end{aligned}$$

By summing these inequalities over $d \geq 1$ we get $\|\Delta\|_2^2 \leq \rho^2$ or

$$\|\Delta\|_2^2 \leq \|\|\Delta\|_1^2 - \|\Delta\|_2^2\| + \|\Delta\|_1^2 \leq \|\Delta\|_2^2/2 + \|\Delta\|_1^2$$

with probability tending to one. This shows Equation (A.6) and concludes the proof of Theorem 2.

A.3 Proof of Theorem 3

We will prove the first equation of the theorem for $h \neq 0$. The second equation follows from the first equation. We first prove that the matrix $T^{-1} \sum_{t=1}^T \mathcal{Z}_{c,t} \widehat{\mathcal{Z}}_{c,t}^\top$ is invertible, where $\mathcal{Z}_{c,t}^\top = (1, \mathcal{Z}_{c,t}^\top)$, $\widehat{\mathcal{Z}}_{c,t}^\top = (1, \widehat{\mathcal{Z}}_{c,t}^\top)$, and $\widehat{\mathcal{Z}}_{c,t} = \widehat{\mathcal{Z}}_t - T^{-1} \sum_{s=1}^T \widehat{\mathcal{Z}}_s$. This implies that $T^{-1} \sum_{t=1}^T \mathcal{Z}_{c,t} \widehat{\mathcal{Z}}_{c,t}^\top$ is invertible. Suppose that the assertion is not true. We can choose a random vector e such that $\|e\| = 1$ and $e^\top \sum_{t=1}^T \mathcal{Z}_{c,t} \widehat{\mathcal{Z}}_{c,t}^\top = 0$. Let \widehat{A} and A^* be the $L \times K$ matrices that are obtained by deleting the first rows of \widehat{A} and A^* , respectively. Let \widehat{A}_c and A_c^* be the matrices obtained from \widehat{A} and A^* by replacing their first rows by $\widehat{\alpha}_0^\top + (T^{-1} \sum_{t=1}^T \widehat{\mathcal{Z}}_t)^\top \widehat{A}$ and $\alpha_0^{\ast\top} + (T^{-1} \sum_{t=1}^T \mathcal{Z}_t)^\top A^*$, respectively. By definition, it follows that

$$\widehat{\mathcal{Z}}_t^\top \widehat{A} = \widehat{\mathcal{Z}}_{c,t}^\top \widehat{A}_c, \quad \mathcal{Z}_t^\top A^* = \mathcal{Z}_{c,t}^\top A_c^*. \quad (\text{A.11})$$

Note that

$$\begin{aligned} &\left\| T^{-1} \sum_{t=1}^T \mathcal{Z}_{c,t} \widehat{\mathcal{Z}}_{c,t}^\top \widehat{A}_c - T^{-1} \sum_{t=1}^T \mathcal{Z}_{c,t} \mathcal{Z}_{c,t}^\top A_c^* \right\| \\ &\leq T^{-1} \sum_{t=1}^T \left\| \mathcal{Z}_{c,t} \left(\widehat{\mathcal{Z}}_{c,t}^\top \widehat{A}_c - \mathcal{Z}_{c,t}^\top A_c^* \right) \right\| \\ &\leq \left(T^{-1} \sum_{t=1}^T \left\| \mathcal{Z}_{c,t} \right\|^2 \right)^{1/2} \left(T^{-1} \sum_{t=1}^T \left\| \widehat{\mathcal{Z}}_t^\top \widehat{A} - \mathcal{Z}_t^\top A^* \right\|^2 \right)^{1/2} \\ &= \mathcal{O}_P(\rho + \delta_K), \end{aligned} \quad (\text{A.12})$$

because of Assumption (A6) and Theorem 2. Thus with $f = T^{-1} \sum_{t=1}^T \mathcal{Z}_{c,t} \mathcal{Z}_{c,t}^\top e$, we obtain

$$\begin{aligned} \|f^\top m\| &= \|f^\top (A_c^* \psi)\| + \mathcal{O}_P(T^{-1/2} + \delta_K) \\ &= \left\| e^\top T^{-1} \sum_{t=1}^T \mathcal{Z}_{c,t} \widehat{\mathcal{Z}}_{c,t}^\top \widehat{A}_c \psi \right\| + \mathcal{O}_P(T^{-1/2} + \rho + \delta_K) \\ &= \mathcal{O}_P(T^{-1/2} + \rho + \delta_K). \end{aligned}$$

This implies that m_0, \dots, m_d are linearly dependent, contradicting to Assumption (A10).

Let \widetilde{B} be the matrix given at (8) with B defined as in (9). Define $\widetilde{\mathcal{Z}}_{c,t} = \widetilde{B}^\top \widehat{\mathcal{Z}}_{c,t}$ and $\widetilde{A}_c = \widetilde{B}^{-1} \widehat{A}_c$. Then $\widetilde{\mathcal{Z}}_{c,t}^\top \widetilde{A}_c = \widehat{\mathcal{Z}}_{c,t}^\top \widehat{A}_c$ and $T^{-1} \sum_{t=1}^T \mathcal{Z}_{c,t} \widetilde{\mathcal{Z}}_{c,t}^\top = T^{-1} \sum_{t=1}^T \mathcal{Z}_{c,t} \mathcal{Z}_{c,t}^\top$. This gives with (A.12)

$$\begin{aligned} \|\widetilde{A}_c - A_c^*\| &= \left\| T^{-1} \sum_{t=1}^T \mathcal{Z}_{c,t} \mathcal{Z}_{c,t}^\top (\widetilde{A}_c - A_c^*) \right\| \mathcal{O}_P(1) \\ &= \left\| T^{-1} \sum_{t=1}^T \mathcal{Z}_{c,t} \widetilde{\mathcal{Z}}_{c,t}^\top \widetilde{A}_c - T^{-1} \sum_{t=1}^T \mathcal{Z}_{c,t} \mathcal{Z}_{c,t}^\top A_c^* \right\| \mathcal{O}_P(1) \\ &= \mathcal{O}_P(\rho + \delta_K). \end{aligned} \quad (\text{A.13})$$

Because of Theorem 2 this implies

$$\|\widetilde{A} - A^*\| = \mathcal{O}_P(\rho + \delta_K). \quad (\text{A.14})$$

Define $\widetilde{\mathcal{Z}}_{c,t}$ by $\widetilde{\mathcal{Z}}_{c,t}^\top = (1, \widetilde{\mathcal{Z}}_{c,t}^\top)$. Note that $\widetilde{\mathcal{Z}}_{c,t} = B^\top \widehat{\mathcal{Z}}_{c,t}$. Also, define $\widetilde{A} = B^{-1} \widehat{A}$, which equals \widehat{A}_c without the first row. From (A10), (A5), (A.14), and Theorem 2, we get

$$\begin{aligned} T^{-1} \sum_{t=1}^T \|\widetilde{\mathcal{Z}}_t - \mathcal{Z}_t\|^2 &= T^{-1} \sum_{t=1}^T \|\widetilde{\mathcal{Z}}_t - \mathcal{Z}_t\|^2 \\ &= T^{-1} \sum_{t=1}^T \left\| \widetilde{\mathcal{Z}}_t^\top (m_0, \dots, m_L)^\top - \mathcal{Z}_t^\top (m_0, \dots, m_L)^\top \right\|^2 \mathcal{O}_P(1) \\ &= T^{-1} \sum_{t=1}^T \left\| \widetilde{\mathcal{Z}}_t^\top A^* - \mathcal{Z}_t^\top A^* \right\|^2 \mathcal{O}_P(1) + T^{-1} \sum_{t=1}^T \left\| \widetilde{\mathcal{Z}}_t^\top \widetilde{A} - \mathcal{Z}_t^\top A^* \right\|^2 \\ &\quad \times \mathcal{O}_P(1) + \mathcal{O}_P(\delta_K^2) \\ &\leq T^{-1} \sum_{t=1}^T \|\widetilde{\mathcal{Z}}_t - \mathcal{Z}_t\|^2 \|\widetilde{A} - A^*\|^2 \mathcal{O}_P(1) + T^{-1} \sum_{t=1}^T \|\mathcal{Z}_t\|^2 \\ &\quad \times \|\widetilde{A} - A^*\|^2 \mathcal{O}_P(1) \\ &+ T^{-1} \sum_{t=1}^T \|\widetilde{\mathcal{Z}}_t^\top \widetilde{A} - \mathcal{Z}_t^\top A^*\|^2 \mathcal{O}_P(1) + \mathcal{O}_P(\rho^2 + \delta_K^2) \\ &= \mathcal{O}_P(\rho^2 + \delta_K^2). \end{aligned} \quad (\text{A.15})$$

From Equation (A.15) one gets

$$T^{-1} \sum_{t=1}^T \|\tilde{Z}_{c,t} - Z_{c,t}\|^2 = \mathcal{O}_P(\rho^2 + \delta_K^2). \quad (\text{A.16})$$

We will show that for $h \neq 0$

$$T^{-1} \sum_{t=h+1}^T \{(\tilde{Z}_{c,t+h} - Z_{c,t+h}) - (\tilde{Z}_{c,t} - Z_{c,t})\} Z_{c,t}^\top = \mathcal{O}_P(T^{-1/2}). \quad (\text{A.17})$$

This implies the first statement of Theorem 3, because by (A.16)

$$\begin{aligned} T^{-1} \sum_{t=-h+1}^T (\tilde{Z}_{c,t} - Z_{c,t})(\tilde{Z}_{c,t+h}^\top - Z_{c,t+h}^\top) &= \mathcal{O}_P(\rho^2 + \delta_K^2) \\ &= \mathcal{O}_P(T^{-1/2}). \end{aligned}$$

For the proof of (A.17), let $\tilde{\alpha}_c$ be the stack form of $\tilde{\mathcal{A}}_c$ and $\tilde{\alpha}_{c,0}^\top$ be its first row. Using the representation (6) and the first identity of (A.11), it can be verified that

$$\tilde{Z}_{c,t} = \tilde{S}_{t,Z}^{-1} J^{-1} \sum_{j=1}^J \{Y_{t,j} \tilde{A} \psi(X_{t,j}) - \tilde{A} \psi(X_{t,j}) \psi(X_{t,j})^\top \tilde{\alpha}_{c,0}\}, \quad (\text{A.18})$$

$$\tilde{\alpha}_c = \tilde{S}_\alpha^{-1} T^{-1} J^{-1} \sum_{t=1}^T \sum_{j=1}^J \{\psi(X_{t,j}) \otimes \tilde{Z}_{c,t}\} Y_{t,j}, \quad (\text{A.19})$$

where $\tilde{S}_{t,Z} = J^{-1} \sum_{j=1}^J \tilde{A} \psi(X_{t,j}) \psi(X_{t,j})^\top \tilde{A}^\top$ and $\tilde{S}_\alpha = T^{-1} J^{-1} \sum_{t=1}^T \sum_{j=1}^J \{\psi(X_{t,j}) \otimes \tilde{Z}_{c,t}\} \{\psi(X_{t,j}) \otimes \tilde{Z}_{c,t}\}^\top$. Define $\tilde{S}_{t,Z}$ as $\tilde{S}_{t,Z}$ with $\tilde{\mathcal{A}}_c$ replacing \tilde{A} . Also, define $\mathcal{S}_{t,Z} = \mathcal{A}_c^* E\{\psi(X_{t,j}) \psi(X_{t,j})^\top\} \mathcal{A}_c^{*\top}$ and

$$\mathcal{S}_\alpha = T^{-1} \sum_{t=1}^T E\{\{\psi(X_{t,j}) \otimes Z_{c,t}\} \{\psi(X_{t,j}) \otimes Z_{c,t}\}^\top | Z_t\}.$$

Let $\gamma = T^{-1/2}(\rho + \delta_K)^{-1}$. We argue that

$$\sup_{1 \leq t \leq T} \|\tilde{S}_{t,Z} - \mathcal{S}_{t,Z}\| = \mathcal{O}_P(\gamma), \quad \|\tilde{S}_\alpha - \mathcal{S}_\alpha\| = \mathcal{O}_P(\gamma). \quad (\text{A.20})$$

We show the first part of (A.20). The second part can be shown similarly. To prove the first part it suffices to show that, uniformly for $1 \leq t \leq T$,

$$\begin{aligned} J^{-1} \sum_{j=1}^J \mathcal{A}_c^* [\psi(X_{t,j}) \psi(X_{t,j})^\top - E\{\psi(X_{t,j}) \psi(X_{t,j})^\top\}] (\tilde{\mathcal{A}}_c - \mathcal{A}_c^*)^\top \\ = \mathcal{O}_P(\gamma), \end{aligned} \quad (\text{A.21})$$

$$\begin{aligned} J^{-1} \sum_{j=1}^J (\tilde{\mathcal{A}}_c - \mathcal{A}_c^*) [\psi(X_{t,j}) \psi(X_{t,j})^\top - E\{\psi(X_{t,j}) \psi(X_{t,j})^\top\}] \\ (\tilde{\mathcal{A}}_c - \mathcal{A}_c^*)^\top = \mathcal{O}_P(\gamma), \end{aligned} \quad (\text{A.22})$$

$$\begin{aligned} J^{-1} \sum_{j=1}^J \mathcal{A}_c^* [\psi(X_{t,j}) \psi(X_{t,j})^\top - E\{\psi(X_{t,j}) \psi(X_{t,j})^\top\}] \mathcal{A}_c^{*\top} \\ = \mathcal{O}_P(\gamma), \end{aligned} \quad (\text{A.23})$$

$$\begin{aligned} J^{-1} \sum_{j=1}^J \mathcal{A}_c^* E\{\psi(X_{t,j}) \psi(X_{t,j})^\top\} (\tilde{\mathcal{A}}_c - \mathcal{A}_c^*)^\top &= \mathcal{O}_P(\gamma), \quad (\text{A.24}) \\ J^{-1} \sum_{j=1}^J (\tilde{\mathcal{A}}_c - \mathcal{A}_c^*) E\{\psi(X_{t,j}) \psi(X_{t,j})^\top\} (\tilde{\mathcal{A}}_c - \mathcal{A}_c^*)^\top &= \mathcal{O}_P(\gamma). \end{aligned} \quad (\text{A.25})$$

The proof of (A.23)–(A.25) follows by simple arguments. We now show (A.21). Claim (A.22) can be shown similarly. For the proof of (A.21) we use Bernstein's inequality for the following sum:

$$P\left(\left|\sum_{j=1}^J W_j\right| > x\right) \leq 2 \exp\left(-\frac{1}{2} \frac{x^2}{V + Mx/3}\right). \quad (\text{A.26})$$

Here for a value of t with $1 \leq t \leq T$, the random variable W_j is an element of the $(L+1) \times 1$ -matrix $S = J^{-1} \mathcal{A}_c^* [\psi(X_{t,j}) \psi(X_{t,j})^\top e - E\{\psi(X_{t,j}) \psi(X_{t,j})^\top e\}]$ where $e \in \mathbb{R}^K$ with $\|e\| = 1$. In (A.26), V is an upper bound for the variance of $\sum_{j=1}^J W_j$ and M is a bound for the absolute values of W_j (i.e. $|W_j| \leq M$ for $1 \leq j \leq J$, a.s.). With some constants C_1 and C_2 that do not depend on t and the row number we get $V \leq C_1 J^{-1}$ and $M \leq C_2 K^{1/2} J^{-1}$. Application of Bernstein's inequality gives that, uniformly for $1 \leq t \leq T$ and $e \in \mathbb{R}^K$ with $\|e\| = 1$, all $(L+1)$ elements of S are of order $\mathcal{O}_P(\gamma)$. This shows claim (A.21).

From (A.13), (A.15), (A.18), (A.19), and (A.20) it follows that uniformly for $1 \leq t \leq T$,

$$\begin{aligned} \tilde{Z}_{c,t} - Z_{c,t} &= S_{t,Z}^{-1} J^{-1} \sum_{j=1}^J \varepsilon_{t,j} \mathcal{A}_c^* \psi(X_{t,j}) + S_{t,Z}^{-1} J^{-1} \sum_{j=1}^J \varepsilon_{t,j} \\ &\quad \times (\tilde{\mathcal{A}}_c - \mathcal{A}_c^*) \psi(X_{t,j}) \\ &+ S_{t,Z}^{-1} J^{-1} \sum_{j=1}^J (\tilde{\mathcal{A}}_c - \mathcal{A}_c^*) \psi(X_{t,j}) \psi(X_{t,j})^\top \mathcal{A}_c^{*\top} Z_{c,t} + \mathcal{O}_P(T^{-1/2}) \\ &\equiv \Delta_{t,1,Z} + \Delta_{t,2,Z} + \Delta_{t,3,Z} + \mathcal{O}_P(T^{-1/2}). \end{aligned} \quad (\text{A.27})$$

For the proof of the theorem it remains to show that for $1 \leq j \leq 3$

$$T^{-1} \sum_{t=-h+1}^T (\Delta_{t+h,j,Z} - \Delta_{t,j,Z}) Z_{c,t}^\top = \mathcal{O}_P(T^{-1/2}). \quad (\text{A.28})$$

This can be easily checked for $j = 1$. For $j = 2$ it follows from $\|\tilde{\mathcal{A}}_c - \mathcal{A}_c^*\| = \mathcal{O}(\rho + \delta_K)$ and

$$E\left\{\left\|T^{-1} J^{-1} \sum_{t=1}^T \sum_{j=1}^J \varepsilon_{t,j} S_{t,Z}^{-1} \mathcal{M} \psi(X_{t,j})\right\|^2\right\} = \mathcal{O}(KJ^{-1}T^{-1})$$

for any $L \times K$ matrix \mathcal{M} with $\|\mathcal{M}\| = 1$. For the proof of (A.28) for $j = 3$, it suffices to show that

$$T^{-1} \sum_{t=1}^{T+h} \Delta_{t,j,Z} (Z_{c,t-h} - Z_{c,t})^\top = \mathcal{O}_P(T^{-1/2}). \quad (\text{A.29})$$

We note first that for $1 \leq l \leq L$

$$\begin{aligned} T^{-1} \sum_{t=1}^{T+h} \Delta_{t,3,Z} (Z_{c,t-h,l} - Z_{c,t,l}) \\ = T^{-1} J^{-1} \sum_{t=1}^{T+h} \sum_{j=1}^J \left\{ \left(V_{h,t}^\top \mathcal{A}_c^* \psi(X_{t,j}) \psi(X_{t,j})^\top \right) \otimes S_{t,Z}^{-1} \right\} (\tilde{\alpha} - \alpha^*), \end{aligned}$$

where $V_{h,t} = (Z_{c,t-h,t} - Z_{c,t,t})Z_{c,t,t}$, and $\tilde{\alpha}$ and α^* denote the stack forms of A and A^* , respectively. For the proof of (A.29) it suffices to show

$$T^{-1}J^{-1} \sum_{t=1}^{T+h} \sum_{j=1}^J \{ (E[V_{h,t}]^\top \mathcal{A}_c^* \psi(X_{t,j}) \psi(X_{t,j})^\top) \otimes S_{t,Z}^{-1} \} \times (\tilde{\alpha} - \alpha^*) = o_P(T^{-1/2}), \tag{A.30}$$

$$\left\| T^{-1}J^{-1} \sum_{t=1}^{T+h} \sum_{j=1}^J \{ (\{V_{h,t} - E[V_{h,t}]\}^\top \mathcal{A}_c^* \psi(X_{t,j}) \psi(X_{t,j})^\top) \otimes S_{t,Z}^{-1} \} \right\|^2 = O_P(KJ^{-1}T^{-1}). \tag{A.31}$$

Claim (A.31) can be easily shown by calculating the expectation of the left hand side of (A.31) and by using the mixing condition at Assumption (A9). For a proof of (A.30) we remark first that by construction

$$0 = T^{-1} \sum_{t=1}^T (\tilde{Z}_{c,t} - Z_{c,t}) Z_{c,t}^\top.$$

Using (A.27) and similar arguments as in the proof of (A.28) for $j = 1, 2$ we get that

$$T^{-1} \sum_{t=1}^T \Delta_{t,3,Z} Z_{c,t}^\top = T^{-1}J^{-1} \sum_{t=1}^T \sum_{j=1}^J \{ (Z_{c,t} Z_{c,t}^\top \mathcal{A}_c^* \psi(X_{t,j}) \psi(X_{t,j})^\top) \otimes S_{t,Z}^{-1} \} (\tilde{\alpha} - \alpha^*) = o_P(T^{-1/2}).$$

As in the proof of (A.31) one can show that

$$\left\| T^{-1}J^{-1} \sum_{t=1}^{T+h} \sum_{j=1}^J \{ (\{Z_{c,t} Z_{c,t}^\top - E[Z_{c,t} Z_{c,t}^\top]\} \mathcal{A}_c^* \psi(X_{t,j}) \psi(X_{t,j})^\top) \otimes S_{t,Z}^{-1} \} \right\|^2 = O_P(KJ^{-1}T^{-1}).$$

The last two equalities imply that

$$T^{-1}J^{-1} \sum_{t=1}^T \sum_{j=1}^J \{ (E[Z_{c,t} Z_{c,t}^\top] \mathcal{A}_c^* \psi(X_{t,j}) \psi(X_{t,j})^\top) \otimes S_{t,Z}^{-1} \} \times (\tilde{\alpha} - \alpha^*) = o_P(T^{-1/2}).$$

Because of Assumption (A9) this implies claim (A.29) and concludes the proof of Theorem 3.

[Received June 2007. Revised August 2008.]

REFERENCES

Biswal, B., Yetkin, F., Haughton, V., and Hyde, J. (1995), "Functional Connectivity in the Motor Cortex of Resting Human Brain Using Echo-Planar MRI," *Magnetic Resonance in Medicine*, 34, 537–541.
 Black, F., and Scholes, M. (1973), "The Pricing of Options and Corporate Liabilities," *The Journal of Political Economy*, 81, 637–654.

Brüggemann, R., Lütkepohl, H., and Saikkonen, P. (2006), "Residual Autocorrelation Testing for Vector Error Correction Models," *Journal of Econometrics*, 134, 579–604.
 Brumback, B., and Rice, J. A. (1998), "Smoothing Spline Models for the Analysis of Nested and Crossed Samples of Curves," *Journal of the American Statistical Association*, 93, 961–994.
 Connor, G., Hagmann, M., and Linton, O. (2007). *Efficient Semiparametric Estimation of the Fama-French Model and Extensions*, Preprint.
 Connor, G., and Linton, O. (2007), "Semiparametric Estimation of a Characteristic-based Factor Model of Stock Returns," *Journal of Empirical Finance*, 14, 694–717.
 Cont, R., and da Fonseca, J. (2002), "The Dynamics of Implied Volatility Surfaces," *Quantitative Finance*, 2, 45–60.
 de Boor, C. (2001). *A Practical Guide to Splines*, Berlin, Heidelberg: Springer-Verlag.
 Diebold, F. X., and Li, C. (2006), "Forecasting the Term Structure of Government Bond Yields," *Journal of Econometrics*, 130, 337–364.
 Fama, E. F., and French, K. R. (1992), "The Cross-Section of Expected Stock Returns," *The Journal of Finance*, 47, 427–465.
 Fan, J., Yao, Q., and Cai, Z. (2003), "Adaptive Varying-Coefficient Linear Models," *Journal of the Royal Statistical Society: Series B*, 65, 57–80.
 Fengler, M. R., Härdle, W., and Mammen, E. (2007), "A Semiparametric Factor Model for Implied Volatility Surface Dynamics," *Journal of Financial Econometrics*, 5, 189–218.
 Forni, M., Hallin, M., Lippi, M., and Reichlin, L. (2000), "The Generalized Dynamic Factor Model: Identification and Estimation," *The Review of Economics and Statistics*, 82, 540–554.
 Forni, M., and Lippi, M. (2001), "The Generalized Factor Model: Representation Theory," *Econometric Theory*, 17, 1113–1141.
 Gasser, T., Möcks, R., and Verleger, R. (1983), "Selavco: A Method to Deal With Trial-to-Trial Variability of Evoked Potential," *Electroencephalography and Clinical Neurophysiology*, 55, 717–723.
 Hafner, R. (2004). *Stochastic Implied Volatility*, Berlin: Springer.
 Hallin, M., and Liska, R. (2007), "Determining the Number of Factors in the Generalized Dynamic Factor Model," *Journal of the American Statistical Association*, 102, 603–617.
 Hansen, L.H., Nielsen, B., and Nielsen, J.P. (2004). "Two Sided Analysis of the Variance With a Latent Time Series," Nuffield College Economic Working Paper 2004-W25, University of Oxford.
 Hosking, J. R. M. (1980), "The Multivariate Portmanteau Statistic," *Journal of the American Statistical Association*, 75, 602–608.
 Hosking, J. R. M. (1981), "Lagrange-Multiplier Tests of Multivariate Time-Series Models," *Journal of the Royal Statistical Society, Series B*, 43, 219–230.
 Kantorovich, L. V., and Akilov, G. P. (1982). *Functional Analysis* (2nd ed.), Oxford, U.K.: Pergamon Press.
 Kauermann, G. (2000), "Modeling Longitudinal Data With Ordinal Response by Varying Coefficients," *Biometrics*, 56, 1692–1698.
 Lee, R. D., and Carter, L. (1992), "Modeling and Forecasting the Time Series of U.S. Mortality," *Journal of the American Statistical Association*, 87, 659–671.
 Logothetis, N., and Wandell, B. (2004), "Interpreting the Bold Signal," *Annual Review of Physiology*, 66, 735–769.
 Lütkepohl, H. (1993). *Introduction to Multiple Time Series Analysis*, Berlin, Heidelberg: Springer-Verlag.
 Martinussen, T., and Scheike, T. (2000), "A Nonparametric Dynamic Additive Regression Model for Longitudinal Data," *Annals of Statistics*, 28, 1000–1025.
 Nelson, C. R., and Siegel, A. F. (1987), "Parsimonious Modeling of Yield Curves," *Journal of Business*, 60, 473–489.
 Peña, D., and Box, E. P. (1987), "Identifying a Simplifying Structure in Time Series," *Journal of the American Statistical Association*, 82, 836–843.
 Stock, J. H., and Watson, M.W. (2005). "Implications of Dynamic Factor Models for VAR Analysis," NBER Working Papers 11467, National Bureau of Economic Research, Inc., available at <http://ideas.repec.org/p/nbr/nberwo/11467.html>.
 Van de Geer, S. (2000). *Empirical Processes in M-Estimation*, Cambridge, U.K.: Cambridge University Press.
 Vrahatis, M. N. (1989), "A Short Proof and a Generalization of Miranda's Existence Theorem," *Proceedings of the American Mathematical Society*, 107, 701–703.
 Worsley, K., Liao, C., Aston, J., Petre, V., Duncan, G., Morales, F., and Evans, A. (2002), "A General Statistical Analysis for fMRI Data," *NeuroImage*, 15, 1–15.
 Yang, L., Park, B. U., Xue, L., and Härdle, W. (2006), "Estimation and Testing for Varying Coefficients in Additive Models With Marginal Integration," *Journal of the American Statistical Association*, 101, 1212–1227.

A Generalized ARFIMA Process with Markov-Switching Fractional Differencing Parameter

Wen-Jen Tsay¹ and Wolfgang Karl Härdle^{2*}

April 24, 2007

¹ The Institute of Economics, Academia Sinica, Taiwan

² CASE – Center for Applied Statistics and Economics

Humboldt-Universität zu Berlin,

Spandauer Straße 1, 10178 Berlin, Germany

Abstract

We propose a general class of Markov-switching-ARFIMA processes in order to combine strands of long memory and Markov-switching literature. Although the coverage of this class of models is broad, we show that these models can be easily estimated with the DLV algorithm proposed. This algorithm combines the Durbin-Levinson and Viterbi procedures. A Monte Carlo experiment reveals that the finite sample performance of the proposed algorithm for a simple mixture model of Markov-switching mean and ARFIMA(1, d , 1) process is satisfactory. We apply the Markov-switching-ARFIMA models to the U.S. real interest rates, the Nile river level, and the U.S. unemployment rates, respectively. The results are all highly consistent with the conjectures made or empirical results found in the literature. Particularly, we confirm the conjecture in Beran and Terrin (1996) that the observations 1 to about 100 of the Nile river data seem to be more independent than the subsequent observations, and the value of differencing parameter is lower for the first 100 observations than for the subsequent data.

Key words: Markov chain; ARFIMA process; Viterbi algorithm; Long memory

JEL classification: C14, C22, C32, C52, C53, G12

*This research was supported by the Deutsche Forschungsgemeinschaft through the SFB 649 'Economic Risk'. Correspondence to Wen-Jen Tsay. The Institute of Economics, Academia Sinica, Taipei, Taiwan, R.O.C. Tel: (886-2) 2782-2791 ext. 296. Fax: (886-2) 2785-3946. E-Mail: wtsay@ieas.econ.sinica.edu.tw

1 Introduction

It is well known that many time series data exhibit long memory, or long-range dependence, including the Nile river level, *ex post* real interest rate, forward premium, and the dynamics of aggregate partisanship and macroideology. Among the many other examples that Beran (1994) gives the Nile river data has been known for its long memory behavior since ancient times, and this is one of the time series that led to the discovery of the Hurst effect (Hurst, 1951) and motivated Mandelbrot and his co-workers (Mandelbrot and van Ness, 1968; Mandelbrot and Wallis, 1969) to introduce fractional Gaussian noise to model long memory phenomenon.

Long range dependence also has been observed in financial data. As demonstrated by Ding et al. (1993), de Lima and Crato (1993) and Bollerslev and Mikkelsen (1996) that the volatility of most financial time series exhibits strong persistency and can be well described as a long memory process. Evidence of financial market volatility's strong persistency inspired Breidt et al. (1998) to propose a class of long memory stochastic volatility (LMSV) models. Deo et al. (2006) also show that the LMSV model is useful for forecasting realized volatility (RV) which is an important quantity in finance.

Figure 1 displays the yearly Nile river minima based on measurements at the Roda gauge near Cairo during the years 622-1284. Beran (1994, p.33) documents that "When one only looks at short time periods, then there seem to be cycles or local trend. However, looking at the whole series, there is no apparent persisting cycle." The changing pattern of the Nile river data leads Bhattacharya et al. (1983) to argue that the so-called Hurst effect can also be explained as if the observations are composed as the sum of a weakly dependent stationary process and a deterministic function. As a consequence it is important to distinguish between a long memory time series and a weakly dependent time series with change-points in the mean. This question has been intensively considered in the literature, including Künsch (1986) and Heyde and Dai (1996). Berkes et al. (2006) presents an overview about this strand of literature. Similarly, Diebold and Inoue (2001) shows that long memory also may be easily confused with a Markov-switching mean. Thus, most of the existing literature considers long memory as a competing modeling framework against the structural change and Markov-switching models.

The Nile river level time series is far more complicated than a pure long memory or




Figure 1: Yearly Nile river minima based on measurements at the Roda gauge near Cairo.

a weakly dependent time series with change-points in the mean to describe. Beran and Terrin (BT) (1996) suggest therefore that the Hurst parameter characterizing the yearly Nile river might change over time. When estimating the Nile river data with the autoregressive fractionally-integrated moving-average (ARFIMA) or $I(d)$ process introduced by Granger (1980), Granger and Joyeux (1980) and Hosking (1981), Beran and Terrin (1996, p.629) show that the data can be well fitted with an ARFIMA(0, d , 0) model with $d = 0.4$, where the fractional differencing parameter d of ARFIMA process acts like the Hurst parameter H of fractional Gaussian noise in characterizing the hyperbolic decay of the autocovariance function of a long memory process. BT further claim that the observations 1 to about 100 seem to be more independent than the subsequent observations, and the value of the fractional differencing parameter might be *lower* for the first 100 observations than for the subsequent data. If this claim is right, then there should be a structural change in the long range persistence of the Nile river data around the year 720, and the Nile river data neither can be described with a pure long memory nor a weakly dependent time series with change-points in the mean.

The possible change of the differencing parameter stimulate BT to propose a statistic for testing the stability of the fractional differencing parameter. This testing statistic has been further discussed and extended in Horváth and Shao (1999) and Horváth (2001). However,

their methods can not identify the change points of the fractional differencing parameter. A Bayesian random persistent-shift (RPS) method for detecting structural change in the differencing parameter and the process level has been considered in Ray and Tsay (2002). Nevertheless, the RPS method is not built on the Markov-switching framework, thus may not fully characterize the cycling behavior of the data series, i.e., “seven years of great abundance” and “seven years of famine” — the Joseph effect named by Mandelbrot and van Ness (1968) and Mandelbrot and Wallis (1969).

The above considerations lead us to combine the long memory and Markov-switching literature into a unified framework. We introduce a Markov-switching-ARFIMA (MS-ARFIMA) process by extending the hidden Markov model. Given that the hidden Markov model has become extremely popular in speech recognition as shown in Juang and Rabiner (1991) and Qian and Titterton (1991), and in econometrics, finance, genetics, and neurophysiology as outlined in Robert et al. (2000), the MS-ARFIMA model provides a flexible modeling framework for many applications to these fields. Moreover, the research conducted in this paper also solve the puzzle raised by Diebold and Inoue (2001) by estimating the differencing parameter allowing for the parameters of interest are Markov-switching.

The remaining parts of this paper are arranged as follows: Section 2 presents the MS-ARFIMA process and the algorithms for estimating the parameters of interest. In Section 3 we consider the finite sample performance of the proposed algorithm under the simple mixture of a Markov-switching mean and an ARFIMA(1, d , 1) process. We then apply the proposed methodology to the U.S. real interest rates, the Nile river data, and the U.S. unemployment rates in Section 4. Section 5 provides a conclusion.

2 Models and Main Results

The objective of this paper is to propose a general class of Markov-switching-ARFIMA processes in order to combine strands of long memory and Markov-switching literature. This class of models offers a rich dynamic mixture of a Markov chain and an $I(d)$ process.

Let $\{s_t\}_{t=1}^T$ be the latent sample path of an N -state Markov chain. At each time s_t can

assume only an integer value of $1, 2, \dots, N$, and its transition probability matrix is

$$\mathcal{P} \equiv \begin{bmatrix} p_{11} & p_{21} & \cdots & p_{N1} \\ p_{12} & p_{22} & \cdots & p_{N2} \\ \vdots & \vdots & \ddots & \vdots \\ p_{1N} & p_{2N} & \cdots & p_{NN} \end{bmatrix},$$

where $p_{ij} = P(s_t = j \mid s_{t-1} = i)$ and $\sum_{j=1}^N p_{ij} = 1$ for all i .

An $I(d)$ process, x_t , is defined as:

$$(1 - L)^d x_t = h_t,$$

where L is the lag operator ($Lk_t = k_{t-1}$) and h_t is a short memory process. When $d > 0$, the $I(d)$ process is often called the long memory process, because its autocovariance function is not summable so as to capture the long range dependence of a time series. In addition, the $I(d)$ process is nonstationary when $d \geq \frac{1}{2}$, otherwise, it is covariance stationary.

Combining the defining feature of a Markov chain and that of an $I(d)$ process, we propose the following MS-ARFIMA(p, d, q) process:

$$w_t = \mu_{s_t} I\{t \geq 1\} + (1 - L)^{-d_{s_t}} \sigma_{s_t} z_t I\{t \geq 1\} = \mu_{s_t} I\{t \geq 1\} + y_{s_t}, \quad (1)$$

where $I\{\cdot\}$ is the indicator function and z_t is stationary process with mean zero and bounded positive spectral density $f_u(\lambda) \sim G_0$ as $\lambda \rightarrow 0$ at each possible regime, thus including stationary and invertible ARMA process as its special case. The most distinguished feature of the process is that the fractional differencing parameter d_{s_t} well known in the long memory literature is allowed to be a Markov chain satisfying the following Assumption A:

Assumption A. s_t is independent of z_τ for all t and τ .

The model in (1) subsumes many interesting models in the literature. When $N = 1$, w_t reduces to the specification in (7) of Shimotsu and Phillips (2005):

$$w_t = \mu_0 + (1 - L)^{-d_0} \sigma_0 z_t I\{t \geq 1\} \quad (2)$$

which also can be represented as:

$$w_t = \mu_0 + \sum_{k=0}^{t-1} \frac{(d_0)_k}{k} \sigma_0 z_{t-k}, \quad (3)$$

where

$$(d_0)_k = \frac{\Gamma(d_0 + k)}{\Gamma(d_0)} = (d_0)(d_0 + 1) \dots (d_0 + k - 1) \quad (4)$$

is Pochhammer's symbol for the forward factorial and $\Gamma(\cdot)$ is the gamma function. Moreover, under the model in (1) and $d_{s_t} = 0$, w_t still includes the Markov-switching AR model considered in Hamilton (1989) as one of its special cases. We will show that the estimation of the model in (1) can be easily implemented with the algorithm proposed in this paper, even though the parameter estimation from a noisy version of realizations of Markov models is extremely difficult in all but very simple examples as well documented in Qian and Titterton (1991).

Let the total sample size be T , and denote $\mathcal{W}_t \equiv (w_1, w_2, \dots, w_t)^\top$ the column vector containing the observations from time 1 to time t , while $\mathcal{S}_t = (s_1, s_2, \dots, s_t)^\top$ represents the corresponding states, and $\mathcal{Y}_t = (y_1, y_2, \dots, y_t)^\top$ in (1) is similarly defined. The column vector $\alpha = (\mu_1, \dots, \mu_N, \sigma_1, \dots, \sigma_N, \phi_{11}, \dots, \phi_{1p}, \phi_{21}, \dots, \phi_{Np}, d_1, \dots, d_N, \theta_{11}, \dots, \theta_{Nq})^\top$ and p_{ij} (transition probabilities) consist of the parameters characterizing the conditional density function (cdf) of w_t . After stacking the parameter vector α and the transition probabilities p_{ij} into one column vector ξ , we can represent the cdf of w_t as $f(w_t | \mathcal{S}_t, \mathcal{W}_{t-1}; \xi)$, clearly showing that the cdf of w_t depends on the entire past routes of states (in general). Indeed, there are N^T possible paths of states running throughout the observations \mathcal{W}_T .

To illustrate the proposed algorithm for the model in (1), we first consider the simplest case where w_t in (1) is generated as:

$$w_t = \mu_{s_t} I\{t \geq 1\} + (1 - L)^{-d_0} \sigma_0 \varepsilon_t I\{t \geq 1\} = \mu_{s_t} I\{t \geq 1\} + y_t, \quad (5)$$

where $d < \frac{1}{2}$ and ε_t is a zero mean normally, independently and identically distributed white noise (i.i.d.) with $E(\varepsilon_t^2) = 1$. That is, w_t in (5) is a special type of MS-ARFIMA(0, d , 0) process whose differencing parameter is fixed across different regimes. Under Assumption A and $\varepsilon_t \sim N(0, 1)$ i.i.d. process, the likelihood function of \mathcal{W}_T , $L(\mathcal{S}_T, \mathcal{W}_T; \xi)$ hereafter, for the hidden Markov model in (5) equals

$$L(\mathcal{S}_T, \mathcal{W}_T; \xi) = (2\pi)^{-T/2} |\Lambda|^{-1/2} \exp\left(-\frac{1}{2} \mathcal{Y}_T^\top \Lambda^{-1} \mathcal{Y}_T\right) \prod_{t=1}^T \Pr(s_t | s_{t-1}), \quad (6)$$

where $\Lambda = E(\mathcal{Y}_T \mathcal{Y}_T^\top)$, and $\Pr(s_1 | s_0)$ is evaluated with the unconditional probability that the process will be in regime s_1 . Given that y_t in (5) is a simple ARFIMA(0, d , 0) process,

we can use the Durbin-Levinson algorithm to derive

$$(2\pi)^{-T/2} |\Lambda|^{-1/2} \exp\left(-\frac{1}{2} \mathcal{Y}_T^\top \Lambda^{-1} \mathcal{Y}_T\right) = \prod_{t=1}^T (2\pi)^{-1/2} v_{t-1}^{-1/2} \exp\left\{-\frac{(y_t - \hat{y}_t)^2}{2v_{t-1}}\right\}, \quad (7)$$

where \hat{y}_t denotes the one-step ahead predictor of y_t with the observation \mathcal{Y}_{t-1} as $j \geq 2$, and v_{t-1} is the corresponding one-step ahead prediction variance. Deriche and Tewfik (1993) also have employed the Durbin-Levinson algorithm to estimate a univariate ARFIMA(0, d , 0) processes without Markov-switching characteristic. Note that as $t = 1$, $\hat{y}_1 = 0$, and $v_0 = \gamma_0$ corresponds to the variance of y_t . As a result, the likelihood function in (6) can be rewritten as:

$$L(\mathcal{S}_T, \mathcal{W}_T; \xi) = \prod_{t=1}^T (2\pi)^{-1/2} v_{t-1}^{-1/2} \exp\left\{-\frac{(y_t - \hat{y}_t)^2}{2v_{t-1}}\right\} \Pr(s_t | s_{t-1}), \quad (8)$$

indicating that the *unconditional* likelihood function of the mixture model in (5) can be exactly and recursively evaluated provided that we can identify the true path of s_t , \mathcal{S}_T^* .

We do not know in reality the value of \mathcal{S}_T^* . However, the recursive structure shown in (8) is especially suitable for implementing the Viterbi (1967) algorithm in the digital communication literature to identify the most likely path of states among the N^T possible routes within \mathcal{W}_T . We thus combine the Durbin-Levinson algorithm and the Viterbi algorithm to suggest a *Durbin-Levinson-Viterbi* (DLV) algorithm for the model in (5). When compared to the original Viterbi algorithm designed for solving the problem of maximum a posteriori probability estimate of the state sequence of a finite-state discrete-time Markov process observed in white noise, the DLV algorithm proposed in this paper is concerned with the hidden Markov process observed in a much more general ARFIMA noise. Since the DLV algorithm can estimate the differencing parameter of a time series allowing for the presence of a Markov-switching mean, the puzzle raised by Diebold and Inoue (2001) that long memory can be easily confused with a Markov-switching mean is thus resolved by using this DLV algorithm.

To locate the most likely path running through the data \mathcal{W}_T with the idea of Viterbi (1967), we note first that, for each time t , there are N possible states ending at time t , i.e., $(s_t = i)$, $i = 1, \dots, N$. For a particular node of these N end points at time t , say $(s_t = j)$, there exists a corresponding most likely path:

$$(\mathcal{S}_{t-1}(s_t = j), s_t = j) = (s_1(s_t = j), s_2(s_t = j), \dots, s_{t-1}(s_t = j), s_t = j), \quad (9)$$

which ends at this particular node ($s_t = j$). We refer to the path $(\mathcal{S}_{t-1}(s_t = j), s_t = j)$ in (9) as the *survivor* associated with the node ($s_t = j$). Note that, with little loss of clarity, we do not explicitly specify that the path depends on the parameter ξ and the observations \mathcal{W}_t in order to simplify the notation. The likelihood function generated from this survivor $(\mathcal{S}_{t-1}(s_t = j), s_t = j)$ and the formula in (8) is recorded as $L(\mathcal{S}_{t-1}(s_t = j), s_t = j, \mathcal{W}_t; \xi)$ and is crucial for locating the most likely path running from time 1 to time T . In short, for each node ($s_t = j$) at time t , there exists a most likely path, survivor $(\mathcal{S}_{t-1}(s_t = j), s_t = j)$, and its associated likelihood function $L(\mathcal{S}_{t-1}(s_t = j), s_t = j, \mathcal{W}_t; \xi)$. Most importantly, the number of survivors at each time t is always equal to N .

Given the N survivors at time t and in order to locate the survivor $(\mathcal{S}_t(s_{t+1} = i), s_{t+1} = i)$ for a particular node ($s_{t+1} = i$) at time $t + 1$, among the N segments connecting the node ($s_{t+1} = i$) and the N time- t survivors $(\mathcal{S}_{t-1}(s_t = j), s_t = j)$ recorded at time t , we select the one producing the largest likelihood function $L(\mathcal{S}_t(s_{t+1} = i), s_{t+1} = i, \mathcal{W}_{t+1}; \xi)$ among these N possible candidates, and name it as the survivor $(\mathcal{S}_t(s_{t+1} = i), s_{t+1} = i)$ for this particular node ($s_{t+1} = i$). The computation of the aforementioned likelihoods is simple, because we record the likelihood functions of the N time- t survivors at each time t .

This recursive updating process proceeds from time 1 to time T and results in N time- T survivors $(\mathcal{S}_{T-1}(s_T = i), s_T = i)$ and their associated likelihood function $L(\mathcal{S}_{T-1}(s_T = i), s_T = i, \mathcal{W}_T; \xi)$, for each $i = 1, \dots, N$. From these N time- T survivors we select the one producing the largest likelihood function, say $L(\mathcal{S}_{T-1}(s_T = g), s_T = g, \mathcal{W}_T; \xi)$, as the most likely path running from time 1 to time T . Combining a numerical optimization procedure and this chosen likelihood function $L(\mathcal{S}_{T-1}(s_T = g), s_T = g, \mathcal{W}_T; \xi)$ generated from the Viterbi algorithm and the Durbin-Levinson algorithm displayed in (7), we can estimate the parameters ξ and identify the states \mathcal{S}_T hidden in the observations \mathcal{W}_T .

We now consider another special type of MS-ARFIMA(p, d, q) process:

$$w_t = \mu_{s_t} I\{t \geq 1\} + y_t = \mu_{s_t} I\{t \geq 1\} + (1 - L)^{-d_0} \sigma_0 z_t I\{t \geq 1\}, \quad \phi(L)z_t = \theta(L)\varepsilon_t, \quad (10)$$

where

$$\phi(L) = 1 - \phi_1 L - \dots - \phi_p L^p, \quad \theta(L) = 1 + \theta_1 L + \dots + \theta_q L^q, \quad (11)$$

and the roots of the polynomial $\phi(L)$ and those of $\theta(L)$ in (11) are all outside the unit circle and share no common roots. The model in (10) is much more general than that in (5), but

still can be estimated with the preceding Viterbi algorithm after some modifications. Please note that the value of fractional differencing parameter is unchanged across different regimes as that imposed in (5).

Note that the term y_t in (10) can be rearranged as

$$y_t = (1 - L)^{-d_0} \sigma_0 \phi(L)^{-1} \theta(L) \varepsilon_t, \quad t = 1, 2, \dots \quad (12)$$

We then have

$$\phi(L)y_t = (1 - L)^{-d_0} \sigma_0 \theta(L) \varepsilon_t = \sigma_0 \theta(L) (1 - L)^{-d_0} \varepsilon_t = \sigma_0 \theta(L) \tilde{y}_t, \quad t = 1, 2, \dots, \quad (13)$$

where $\tilde{y}_t = (1 - L)^{-d_0} \varepsilon_t$ is an ARFIMA(0, d , 0) process. Dueker and Serletis (2000) use the same transformation method for estimating an ARFIMA(p , d , q) process. Conditional on a set of $\phi(L)$ and $\theta(L)$ and a suitable starting value, the *conditional* likelihood function of y_t in (12) can still be evaluated exactly with the transformed ARFIMA(0, d , 0) \tilde{y}_t in (13) and the Durbin-Levinson algorithm defined in (7). For example, conditional on y_0 being equal to 0, we can extract an ARFIMA(0, d , 0) process from an ARFIMA(1, d , 1) process as follows:

$$\sigma_0 \tilde{y}_t = y_t - \phi_1 y_{t-1} - \sigma_0 \theta_1 \tilde{y}_{t-1}, \quad t = 1, \dots, T. \quad (14)$$

Conditional on a set of $\phi(L)$ and $\theta(L)$ and a suitable starting value for the parameter ξ , we can recursively and exactly evaluate the conditional likelihood function of the hidden Markov model using the DLV algorithm proposed previously.

The same idea also applies to the class of MS-ARFIMA(p , d , q) processes in (1) where d can be Markov-switching. However, we cannot use the Durbin-Levinson algorithm when the fractional differencing parameter is allowed to be Markov-switching. Nevertheless, the Viterbi algorithm is still powerful enough to locate the most likely path under this circumstance. That is, conditional on a suitable starting value for the parameter ξ , we employ the recursive structure inherent in Viterbi algorithms to identify the most likely path running through the data set.

3 Monte Carlo Experiment

In this section we consider a Monte Carlo experiment to demonstrate the finite sample performance of the proposed DLV algorithm on a special version of the model in (1):

$$w_t = \mu_{s_t} I\{t \geq 1\} + (1 - L)^{-d_0} \sigma_0 (1 - \phi_1 L)^{-1} (1 + \theta_1 L) \varepsilon_t I\{t \geq 1\}. \quad (15)$$

We employ three different values of the fractional differencing parameter:

$$d_0 = \{0.2, 0.3, 0.4\}, \tag{16}$$

along with the following parameters:

$$\mu_1 = 4, \mu_2 = 1, \phi_1 = 0.5, \theta_1 = 0.5, p_{11} = p_{22} = 0.95, \tag{17}$$

and σ_0 is chosen to ensure that the variance of the ARFIMA(1, d , 1) noise in (15) is equal to 1 across different configurations. Note that the positive values of d_0 in (16) are chosen to reflect the variations used in the long memory literature.

All the computations are performed with GAUSS. Two hundred replications are conducted for each specification at 3 different sample sizes ($T = 100, 200, 400$) usually encountered in the empirical applications. For each sample size T , 200 additional values are generated in order to obtain random starting values. The optimization algorithm used to implement the DLV algorithm is the quasi-Newton algorithm of Broyden, Fletcher, Goldfarb, and Shanno (BFGS) contained in the GAUSS MAXLIK library. The maximum number of iterations for each replication is 100.

Table 1 contains the simulation results when the true value of parameters are used as the initial values for estimation procedure. The results reveal that the bias performance from the DLV algorithm is satisfactory (especially when the sample size is larger) for all configurations considered. Moreover, the associated root-mean-squared error (RMSE) almost always decreases with the increasing sample size. We find only two cases where the pattern of RMSE change is not what we expect, i.e., when $d_0 = 0.4$, the RMSE of estimating the parameters μ_1 and μ_2 as $T = 400$ is found to be a little higher than that of estimating the parameters μ_1 and μ_2 as $T = 200$. These two observations demonstrate the ability of the DLV algorithm to deal with the mixture model considered in this section. The performance of DLV algorithm for estimating the fractional differencing parameter is particularly displayed with the box-plots in Figure 2. The above-mentioned observations are clearly borne out in this figure.

We also check the robustness of the preceding simulation results by changing the choice of initial values for estimation. The simulations in Table 1 are replicated by setting the initial values for parameters at the true values except that of d_0 is set at zero. The results

**Table 1. Finite sample performance of the DLV algorithm:
Initial values of parameters are set at the true values of parameters**

Parameter		μ_1	μ_2	p_{11}	p_{22}	σ_0	d_0	ϕ_1	θ_1
$d_0 = 0.4$									
$T = 100$	Bias	-0.010	-0.106	0.010	0.016	0.008	0.183	-0.128	-0.040
	RMSE	1.008	0.991	0.039	0.060	0.022	0.294	0.233	0.138
$T = 200$	Bias	-0.094	-0.098	0.006	0.006	0.003	0.135	-0.101	-0.026
	RMSE	0.978	0.978	0.028	0.025	0.015	0.233	0.191	0.086
$T = 400$	Bias	-0.074	-0.076	0.004	0.004	0.001	0.096	-0.073	-0.013
	RMSE	0.990	0.990	0.019	0.017	0.010	0.192	0.163	0.060
$d_0 = 0.3$									
$T = 100$	Bias	-0.057	-0.070	0.009	0.017	0.012	0.175	-0.109	-0.041
	RMSE	1.042	1.024	0.037	0.060	0.030	0.319	0.245	0.131
$T = 200$	Bias	-0.058	-0.055	0.006	0.006	0.005	0.122	-0.079	-0.030
	RMSE	0.947	0.944	0.027	0.025	0.020	0.260	0.212	0.086
$T = 400$	Bias	-0.038	-0.043	0.004	0.004	0.001	0.090	-0.061	-0.016
	RMSE	0.885	0.883	0.019	0.017	0.015	0.217	0.185	0.061
$d_0 = 0.2$									
$T = 100$	Bias	-0.017	-0.041	0.009	0.017	0.014	0.201	-0.115	-0.044
	RMSE	0.874	0.853	0.037	0.060	0.037	0.341	0.258	0.128
$T = 200$	Bias	-0.042	-0.047	0.006	0.006	0.006	0.167	-0.106	-0.037
	RMSE	0.795	0.792	0.028	0.025	0.024	0.297	0.239	0.088
$T = 400$	Bias	-0.038	-0.046	0.004	0.004	0.002	0.122	-0.085	-0.019
	RMSE	0.670	0.669	0.019	0.017	0.018	0.239	0.203	0.061

Notes: Simulations are based on 200 replications. The data is generated from the mixture model defined in (15), (16) and (17). DLV algorithm is the Durbin-Levinson-Viterbi algorithm proposed in this paper. Bias is computed as the true parameter minus the corresponding average estimated values.




Figure 2: Box-plots of the estimated d from the model defined in (15), (16) and (17) with 200 realizations. The initial values of parameters are set at the true values of parameters. The value $f(g)$ denotes the model specification where $d = f$ and $T = 100 \times g$.

contained Table 2 and Figure 3 indicate that the finite sample performance of our procedure is not sensitive to the initial values used for estimation.

4 Empirical Applications

The methodology developed in this paper is motivated by the dynamic pattern of long memory behavior. Evidence has been given by many methods for such a changing covariance behavior of the Nile river. The applications of the proposed MS-ARFIMA model to actual data are far reaching. For that reason, we consider three data set. The first one is the U.S. real interest rates, the second one is the Nile river data, and the third one is the U.S. unemployment rates.

4.1 Example with real interest rates

In this subsection we first consider the U.S. ex post monthly real interest rate constructed from monthly inflation and Treasury bill rates from January 1953 to December 1990 in Mishkin (1990). The reason we use the original dataset of Mishkin (1990) is to employ it as a benchmark for a clear comparison between the results from the MS-ARFIMA model and

**Table 2. Finite sample performance of the DLV algorithm:
Initial values of parameters are set at the true values of parameters
except that of d_0 is set at zero**

Parameter		μ_1	μ_2	p_{11}	p_{22}	σ_0	d_0	ϕ_1	θ_1
$d_0 = 0.4$									
$T = 100$	Bias	-0.116	-0.122	0.010	0.017	0.009	0.188	-0.130	-0.041
	RMSE	1.030	1.017	0.039	0.060	0.021	0.298	0.235	0.137
$T = 200$	Bias	-0.093	-0.096	0.006	0.006	0.003	0.138	-0.103	-0.027
	RMSE	0.979	0.979	0.028	0.025	0.015	0.238	0.193	0.087
$T = 400$	Bias	-0.074	-0.076	0.004	0.004	0.001	0.096	-0.073	-0.013
	RMSE	0.990	0.990	0.019	0.017	0.010	0.192	0.163	0.060
$d_0 = 0.3$									
$T = 100$	Bias	-0.021	-0.034	0.010	0.017	0.012	0.186	-0.115	-0.040
	RMSE	0.972	0.949	0.039	0.060	0.030	0.325	0.241	0.127
$T = 200$	Bias	-0.046	-0.049	0.006	0.006	0.005	0.126	-0.081	-0.030
	RMSE	0.936	0.937	0.028	0.025	0.021	0.261	0.212	0.086
$T = 400$	Bias	-0.040	-0.044	0.004	0.004	0.002	0.088	-0.059	-0.016
	RMSE	0.912	0.912	0.019	0.017	0.015	0.217	0.184	0.060
$d_0 = 0.2$									
$T = 100$	Bias	-0.018	-0.038	0.009	0.016	0.014	0.195	-0.110	-0.045
	RMSE	0.892	0.864	0.037	0.060	0.037	0.340	0.260	0.130
$T = 200$	Bias	-0.036	-0.040	0.006	0.006	0.006	0.160	-0.100	-0.037
	RMSE	0.804	0.801	0.028	0.025	0.024	0.294	0.238	0.087
$T = 400$	Bias	-0.044	-0.051	0.004	0.004	0.002	0.117	-0.082	-0.018
	RMSE	0.674	0.673	0.019	0.017	0.018	0.235	0.200	0.060

Notes: Simulations are based on 200 replications. The data is generated from the mixture model defined in (15), (16) and (17). DLV algorithm is the Durbin-Levinson-Viterbi algorithm proposed in this paper. Bias is computed as the true parameter minus the corresponding average estimated values.




Figure 3: Box-plots of the estimated d from the model defined in (15), (16) and (17) with 200 realizations. The initial values of parameters are set at the true values except that of d_0 is set at zero. The value $f(g)$ denotes the model specification where $d = f$ and $T = 100 \times g$.

those generated from the methodology employed in earlier papers.

The main feature of the real interest rate is that the whole dataset can be split into three subperiods, January 1953-October 1979, November 1979-October 1982, and November 1982-December 1990, because the operating procedure of the monetary authority changed in October 1979 and October 1982 as argued in Mishkin (1990). Another interesting feature of the real interest rate is that the data of these three subperiods can be well described with the ARFIMA models as shown in Tsay (2000). The simultaneous presence of structural break and long memory within the real interest rate allows itself to be an ideal subject to be investigated with the MS-ARFIMA model.

Allowing the break points to be endogenously determined, Table 3 contains the parameter estimates from the following mixture model with a 2-state Markov chain and an ARFIMA(1, d , 1) noise:

$$w_t = \mu_{s_t} I\{t \geq 1\} + (1 - L)^{-d_0} \sigma_0 z_t I\{t \geq 1\}, \quad (1 - \phi_1 L) z_t = (1 + \theta_1 L) \varepsilon_t, \quad (18)$$

where ϕ_1 or θ_1 is assumed to be zero depending on the noise specification. Following Hamilton (1989), asymptotic standard errors are calculated numerically.

Table 3 shows that the estimates of μ_1 , μ_2 , p_{11} , p_{22} , σ_0 , and d_0 from the DLV algorithm

Table 3. Estimates of Parameters Based on Data for U.S. Monthly Real Interest Rate and the DLV Algorithm

	ARFIMA(0, d , 0)		ARFIMA(0, d , 1)		ARFIMA(1, d , 0)		ARFIMA(1, d , 1)	
	Estimate	S.E.	Estimate	S.E.	Estimate	S.E.	Estimate	S.E.
μ_1	5.3455	0.7494	5.3168	0.7162	5.3116	0.7124	5.3626	0.7706
μ_2	0.7226	0.4814	0.7194	0.4383	0.7184	0.4322	0.7352	0.4958
p_{11}	0.9833	0.0150	0.9833	0.0150	0.9833	0.0150	0.9833	0.0150
p_{22}	0.9977	0.0023	0.9977	0.0023	0.9977	0.0023	0.9977	0.0023
σ_0	2.5094	0.0831	2.5091	0.0831	2.5091	0.0831	2.4979	0.0827
d_0	0.2225	0.0367	0.2062	0.0520	0.2034	0.0653	0.2337	0.0376
ϕ_1	-	-	-	-	0.0324	0.0946	-0.9847	0.0155
θ_1	-	-	0.0279	0.0663	-	-	0.9675	0.0200
L^*	1079.0875		1079.0009		1078.9918		1077.0173	

Notes: The results are based on the MS-ARFIMA model defined in (18). S.E. stands for the standard error of the estimate. L^* represents the negative of the log-likelihood function of the switching model. DLV algorithm is the Durbin-Levinson-Viterbi algorithm proposed in this paper.




Figure 4: US monthly ex post real interest rates, January 1953-December 1990. Solid line denotes the path of estimated switching means from the specification $ARFIMA(0, d, 0)$ in Table 3, while dotted line denotes the observed monthly ex post real interest rates.

are quite robust across all 4 different configurations. More importantly, two identical break points are identified with these four models, thus divide the whole data into three subperiods as suggested in Mishkin (1990). The endogenous break points identified are November 1980 and May 1986, respectively.

Figure 4 displays the U.S. monthly ex post real interest rates and the path of estimated switching means generated from the DLV algorithm. Without loss of generality, only the path of the estimated switching means from the specification $ARFIMA(0, d, 0)$ in Table 3 is reported. Figure 4 shows that the model in (18) provides a satisfactory fitting of the U.S. monthly real interest rates. Although the endogenously identified break points are later than the well-known monetary operating procedure change points (October 1979 and October 1982), this finding is quite reasonable, because it takes some time for the ex post real interest rate to adjust its path after new information arrives. This argument is buttressed with the findings in Figure 4 that the endogenously identified break points are more closely connected to the observed path of the U.S. monthly ex post real interest rates than the monetary operating procedure change points are.

Table 3 also shows that a long memory phenomenon is found in the real interest rate as has been documented in Tsay (2000). Nevertheless, the estimate of the fractional differencing parameter in Table 3 is much lower than that of 0.666 in Table 3 of Tsay (2000) where the

change points are exogenously determined, and it is more in line with the estimates of 0.204, 0.275, and 0.193 from the individual subperiod data presented in Table 3 of Tsay (2000). This implies that the persistence of long memory in the real interest rate is much more mitigated, once we take the potentially switching mean of the data into account, thus confirming the arguments of Diebold and Inoue (2001) that the presence of Markov-switching level might increase the persistence of the data under investigation.

4.2 Example with Nile river data

In this subsection we apply the Viterbi algorithm to the Nile river data with the following model:

$$w_t = \mu_{s_t} I\{t \geq 1\} + (1 - L)^{-d_{s_t}} \sigma_{s_t} \varepsilon_t I\{t \geq 1\}, \quad (19)$$

where N is assumed to be 2. For the purpose of comparison, we estimate a fixed regime ARFIMA(0, d , 0) model for the Nile river data, i.e., $N = 1$ is imposed on this model. The estimated value of d from such a fixed regime ARFIMA(0, d , 0) model is 0.3986 and is almost identical to the finding in Beran and Terrin (1996).

When estimating the model in (19) with the Viterbi algorithm, we find that the value of the differencing parameter in Table 4 is 0.5770 (nonstationary) for one state, and is 0.2143 (stationary) for the other one. In addition, we identify 5 transitions within the Nile river data in the year 720, 805, 815, 878, and 1070. The estimated path of d_{s_t} from the MS-ARFIMA(0, d , 0) model in Table 4 is graphed in Figure 5.

Most impressively, the first transition data occurs in the year of 720, and the associated estimated value of d_{s_t} within the period 622 to 719 is 0.2143 which is lower than the 0.5770 observed in the other regime. These two findings correspond closely to the conjectures in Beran and Terrin (1996) that the observations 1 to about 100 seem to be more independent than the subsequent observations and the value of differencing parameter might be lower for the first 100 observations than for the subsequent data.

In Figures 6 and 7 we present the observations and the fitted values generated from the estimated parameters displayed in Table 4. It is clear that the fitted value from the MS-ARFIMA(0, d , 0) model is much closer to the real data than that generated from the model whose differencing parameter is not Markov switching. Combining the findings of the likelihood values in Table 4, we find that the MS-ARFIMA(0, d , 0) model is a promising

Table 4. Estimates of MS-ARFIMA(0, d , 0) Model based on the Nile River Data

	MS-ARFIMA(0, d , 0)		ARFIMA(0, d , 0)	
	Estimate	S.E.	Estimate	S.E.
μ_1	10.8593	0.6903	11.4847	0.2607
μ_2	11.4939	0.0917	-	-
p_{11}	0.9930	0.0042	-	-
p_{22}	0.9918	0.0050	-	-
σ_1	0.5430	0.0202	0.6995	0.0192
σ_2	0.8143	0.0332	-	-
d_1	0.5770	0.0430	0.3986	0.0309
d_2	0.2143	0.0510	-	-
L^*	687.5642		703.8541	

Notes: The MS-ARFIMA(0, d , 0) model is defined in (19). S.E. stands for the standard error of the estimate based on numerical derivative. L^* represents the negative of the log-likelihood function of the estimated model.




Figure 5: Estimated d_{st} from the MS-ARFIMA(0, d , 0) model in Table 4.




Figure 6: Solid line denotes the Nile river water level divided by 100, while dotted line denotes the corresponding fitted values from the MS-ARFIMA(0, d , 0) model in Table 4.




Figure 7: Solid line denotes the Nile river water level divided by 100, while dotted line denotes the corresponding fitted values from the ARFIMA(0, d , 0) model in Table 4.

alternative to describe the Nile river data.

4.3 Example with unemployment rates

In this subsection we apply the Viterbi algorithm to the U.S. quarterly unemployment rates from 1948 to 2006. This data is based on the monthly unemployment rates contained in *Bureau of Labour Statistics* as those employed in van Dijk et al. (2002) for estimating a fractionally integrated smooth transition autoregressive (FI-STAR) model. However, van Dijk et al (2002) employ the original monthly unemployment rates ranging from July 1986 to December 1999, while we use all the data contained in *Bureau of Labour Statistics*, but focusing on the quarterly frequency usually considered in the business cycle related studies.

As clearly argued in van Dijk et al. (2002) and shown in Figure 8, there are two important empirical features of U.S. unemployment rates, i.e., the shocks to the series is quite persistent and the series seem to rise faster during recessions than it falls during expansions. van Dijk et al. (2002) find that the estimated d is 0.43 from a FI-STAR model presented in their Table 1. This implies that a time series model describing long memory and nonlinearity simultaneously may be useful for modeling U.S. unemployment rates and many other applications.

The aforementioned two features contained in U.S. unemployment also provide another good opportunity to test the applicability of the MS-ARFIMA model. As a consequence we estimate the U.S. quarterly unemployment rates with the following MS-ARFIMA($p, d, 0$)




Figure 8: U.S. quarterly seasonally adjusted unemployment rates, 1948-2006.

model:

$$w_t = \mu_{s_t} I\{t \geq 1\} + (1 - L)^{-d_{s_t}} \sigma_{s_t} \phi(B)^{-1} \varepsilon_t I\{t \geq 1\}, \quad (20)$$

where N is assumed to be 2, and $p = \{3, 4\}$. The choice of $p = 4$ is adopted by following the model specification in (30) of van Dijk et al (2002), while $p = 3$ is chosen to check the robustness of the estimation results from the specification $p = 4$. The major objective of this subsection is to investigate whether the long memory observed in van Dijk et al. (2002) can also be retained from the MS-ARFIMA methodology.

When estimating the model in (20) with the Viterbi algorithm, we find that the values of the estimated fractional differencing parameter from both MS-ARFIMA(3, d , 0) and MS-ARFIMA(4, d , 0) models in Table 5 are very close to that found in van Dijk et al. (2002), thus confirming that long memory phenomenon seems to be present in the U.S. unemployment rates. For clarity of exposition, the estimated path of d_{s_t} from the MS-ARFIMA(3, d , 0) model and that of d_{s_t} from the MS-ARFIMA(4, d , 0) one are graphed in Figure 9 and Figure 10, respectively. These figures clearly show that d_{s_t} are around 0.4-0.5 for both regimes estimated in each MS-ARFIMA(p , d , 0) model in Table 5.

We also check to what extent the fitted values generated from the models in Table 5 can capture the feature of U.S. unemployment rates. This task is not taken in van Dijk et al. (2002) when estimating their FI-STAR model for the U.S. monthly unemployment rates. It is interesting to find in Figure 11 and Figure 12 that the MS-ARFIMA(p , d , 0) model in (20)

Table 5. Estimates of MS-ARFIMA($p, d, 0$) Model based on the U.S. quarterly unemployment rates

	MS-ARFIMA(3, d , 0)		MS-ARFIMA(4, d , 0)	
	Estimate	S.E.	Estimate	S.E.
μ_1	3.8080	0.1552	3.4572	0.3711
μ_2	5.1358	0.4403	3.8254	0.3334
p_{11}	0.9939	0.0067	0.9877	0.0093
p_{22}	0.9896	0.0083	0.9867	0.0101
σ_1	0.1973	0.0135	0.1535	0.0101
σ_2	0.3380	0.0206	0.3921	0.0274
d_1	0.4919	0.1215	0.4429	0.0987
d_2	0.4143	0.1337	0.4342	0.1058
ϕ_1	1.2570	0.1415	1.1325	0.1215
ϕ_2	-0.3822	0.1510	-0.2301	0.1239
ϕ_3	-0.0666	0.0788	-0.0141	0.1053
ϕ_4	-	-	-0.0495	0.0712
L^*	36.1377		14.7662	

Notes: The results are based on the MS-ARFIMA($p, d, 0$) model defined in (20). S.E. stands for the standard error of the estimate based on numerical derivative. L^* represents the negative of the log-likelihood function of the estimated model.




Figure 9: Estimated d_{st} from the MS-ARFIMA(3, d , 0) model in Table 5.




Figure 10: Estimated d_{st} from the MS-ARFIMA(4, d , 0) model in Table 5.




Figure 11: Solid line denotes the U.S. quarterly seasonally adjusted unemployment rates (1948-2006), while dotted line denotes the corresponding fitted values from the MS-ARFIMA(3, d , 0) model in Table 5.

provides a reasonable fit to the data, even though we do not include some seasonal control variables, like seasonal difference operator, as van Dijk et al. (2002) have done for their empirical studies.

5 Conclusions

A general class of MS-ARFIMA processes is suggested to combine long memory and Markov-switching models into one unified framework. The coverage of this class of MS-ARFIMA models is far-reaching, but we show that they still can be easily estimated with the original Viterbi algorithm or the DLV algorithm proposed in this paper. In addition, the simulation reveals that the finite sample performance of the DLV algorithm for a simple mixture model of Markov-switching mean and ARFIMA(1, d , 1) process is satisfactory. When applying the MS-ARFIMA models to the U.S. real interest rates, the Nile river level, and the U.S. unemployment rates, the estimation results are both highly compatible with the conjectures made in the literature. Accordingly, the MS-ARFIMA model considered in this paper not only can be used for solving the puzzle raised by Diebold and Inoue (2001), but can also find many potential applications in several scientific research fields.




Figure 12: Solid line denotes the U.S. quarterly seasonally adjusted unemployment rates (1948-2006), while dotted line denotes the corresponding fitted values from the MS-ARFIMA(4, d , 0) model in Table 5.

References

- Beran, J. (1994), *Statistics for Long-Memory Processes*. New York: Chapman and Hall.
- Beran, J. and Terrin, N. (1996), “Testing for a Change of the Long-Memory Parameter”, *Biometrika*, 83, 627-638.
- Berkes, I., Horváth, L., Kokoszka, P. and Shao, Q.M. (2006), “On Discriminating between Long-Range Dependence and Changes in Mean”, *The Annals of Statistics*, 34, 1140-1165.
- Bhattacharya, R.N., Gupta, V.K. and Waymire, E. (1983), “The Hurst Effect under Trends”, *Journal of Applied Probability*, 20, 649-662.
- Bollerslev, T. and Mikkelsen, H.O.A. (1996), “Modeling and Pricing Long-Memory in Stock Market Volatility”, *Journal of Econometrics*, 73, 151-184.
- Breidt, F.J., Crato, N. and de Lima, P. (1998), “The Detection and Estimation of Long Memory in Stochastic Volatility”, *Journal of Econometrics*, 83, 325-348.
- de Lima, P. and Crato, N. (1993), “Long-Range Dependence in the Conditional Variance of Stock Returns”, *Proceedings of the Business and Economic Statistics Section*, August 1993 Joint Statistical Meetings, San Francisco.

- Deo, R., Hurvich, C. and Lu, Y. (2006), "Forecasting Realized Volatility Using a Long-Memory Stochastic Volatility Model: Estimation, Prediction and Seasonal Adjustment", *Journal of Econometrics* 131, 29-58.
- Deriche, J.A. and Tewfik, A.H. (1993), "Maximum Likelihood Estimation of the Parameters of Discrete Fractionally Differenced Gaussian Noise Process," *IEEE Transactions on Signal Processing*, 41, 2977-2989.
- Diebold, F.X. and Inoue, A. (2001), "Long Memory and Regime Switching", *Journal of Econometrics*, 105, 131-159.
- Ding, Z., Granger, C.W.J. and Engle, R.F. (1993), "A Long Memory Property of Stock Market Returns and a New Model", *Journal of Empirical Finance*, 1, 83-106.
- Dueker, M. and Serletis, A. (2000), "Do Real Exchange Rates Have Autoregressive Unit Roots? A Test under the Alternative of Long Memory and Breaks", Working Paper 2000-016A, Federal Reserve Bank of St. Louis.
- Granger, C.W.J. (1980), "Long Memory Relationships and the Aggregation of Dynamic Models", *Journal of Econometrics*, 14, 227-238.
- Granger, C.W.J. and Joyeux, R. (1980), "An Introduction to Long-Memory Time Series Models and Fractional Differencing", *Journal of Time Series Analysis*, 1, 15-29.
- Hamilton, J.D. (1989), "A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle", *Econometrica*, 57, 357-384.
- Heyde, C.C. and Dai, W. (1996), "On the Robustness to Small Trends of Estimation based on the Smoothed Periodogram", *Journal of Time Series Analysis*, 17, 141-150.
- Horváth, L. (2001), "Change-Point Detection in Long-Memory Processes. *Journal of Multivariate Analysis*, 78, 218-234.
- Horváth, L. and Shao, Q.M. (1999), "Limit Theorems for the Union-Intersection Test", *Journal of Statistical Planning and Inference*, 44, 133-148.
- Hosking, J.R.M. (1981), "Fractional Differencing", *Biometrika*, 68, 165-176.

- Hurst, H.E. (1951), “Long-Term Storage Capacity of Reservoirs”, *Transactions of the American Society of Civil Engineers*, 116, 770-799.
- Juang, B.H. and Rabiner, L.R. (1991), “Hidden Markov Models for Speech Recognition”, *Technometrics*, 33, 251-272.
- Künsch, H. (1986), “Discrimination between Monotonic Trends and Long-Range Dependence”, *Journal of Applied Probability*, 23, 1025-1030.
- Mandelbrot, B.B. and van Ness, J.W. (1968), “Fractional Brownian Motions, Fractional Noises and Applications”, *SIAM Review*, 10, 422-437.
- Mandelbrot, B.B. and Wallis, J.R. (1969), “Some Long-Run Properties of Geophysical Records”, *Water Resources Research*, 5, 321-340.
- Mishkin, F.S. (1990), “What does the Term Structure of Interest Rate Tell Us about Future Inflation? *Journal of Monetary Economics*, 25, 77-95.
- Qian, W. and Titterton, D.M. (1991), “Estimation of Parameters in Hidden Markov Models”, *Philosophical Transactions: Physical Sciences and Engineering*, 337, 407-428.
- Ray, B.K. and Tsay, R.S. (2002), “Bayesian Methods for Change-Point Detection in Long-Range Dependent Process”, *Journal of Time Series Analysis*, 23, 687-705.
- Robert, C.P., Rydén, R. and Titterton, D.M. (2000), “Bayesian Inference in Hidden Markov Models through the Reversible Jump Markov Chain Monte Carlo Method”, *Journal of the Royal Statistical Society B*, 62, 57-75.
- Shimotsu, K. and Phillips, P.C.B. (2005), “Exact Local Whittle Estimation of Fractional Integration”, *The Annals of Statistics*, 33, 1890-1933.
- Tsay, W.J. (2000), “Long Memory Story of the Real Interest Rate”, *Economic Letters*, 67, 325-330.
- van Dijk, D., Franses, P.H. and Raap, R. (2002), “A Nonlinear Long Memory Model, with an Application to US Unemployment”, *Journal of Econometrics*, 110, 135-165.
- Viterbi, A.J. (1967), “Error Bounds for Convolutional Codes and an Asymptotic Optimum Decoding Algorithm”, *IEEE Transactions on Signal Processing*, IT-13, 260-269.

SFB 649 Discussion Paper Series 2007


For a complete list of Discussion Papers published by the SFB 649, please visit <http://sfb649.wiwi.hu-berlin.de>.

- 001 "Trade Liberalisation, Process and Product Innovation, and Relative Skill Demand" by Sebastian Braun, January 2007.
- 002 "Robust Risk Management. Accounting for Nonstationarity and Heavy Tails" by Ying Chen and Vladimir Spokoiny, January 2007.
- 003 "Explaining Asset Prices with External Habits and Wage Rigidities in a DSGE Model." by Harald Uhlig, January 2007.
- 004 "Volatility and Causality in Asia Pacific Financial Markets" by Enzo Weber, January 2007.
- 005 "Quantile Sieve Estimates For Time Series" by Jürgen Franke, Jean-Pierre Stockis and Joseph Tadjuidje, February 2007.
- 006 "Real Origins of the Great Depression: Monopolistic Competition, Union Power, and the American Business Cycle in the 1920s" by Monique Ebell and Albrecht Ritschl, February 2007.
- 007 "Rules, Discretion or Reputation? Monetary Policies and the Efficiency of Financial Markets in Germany, 14th to 16th Centuries" by Oliver Volckart, February 2007.
- 008 "Sectoral Transformation, Turbulence, and Labour Market Dynamics in Germany" by Ronald Bachmann and Michael C. Burda, February 2007.
- 009 "Union Wage Compression in a Right-to-Manage Model" by Thorsten Vogel, February 2007.
- 010 "On σ -additive robust representation of convex risk measures for unbounded financial positions in the presence of uncertainty about the market model" by Volker Krätschmer, March 2007.
- 011 "Media Coverage and Macroeconomic Information Processing" by Alexandra Niessen, March 2007.
- 012 "Are Correlations Constant Over Time? Application of the CC-TRIG_t-test to Return Series from Different Asset Classes." by Matthias Fischer, March 2007.
- 013 "Uncertain Paternity, Mating Market Failure, and the Institution of Marriage" by Dirk Bethmann and Michael Kvasnicka, March 2007.
- 014 "What Happened to the Transatlantic Capital Market Relations?" by Enzo Weber, March 2007.
- 015 "Who Leads Financial Markets?" by Enzo Weber, April 2007.
- 016 "Fiscal Policy Rules in Practice" by Andreas Thams, April 2007.
- 017 "Empirical Pricing Kernels and Investor Preferences" by Kai Detlefsen, Wolfgang Härdle and Rouslan Moro, April 2007.
- 018 "Simultaneous Causality in International Trade" by Enzo Weber, April 2007.
- 019 "Regional and Outward Economic Integration in South-East Asia" by Enzo Weber, April 2007.
- 020 "Computational Statistics and Data Visualization" by Antony Unwin, Chun-houh Chen and Wolfgang Härdle, April 2007.
- 021 "Ideology Without Ideologists" by Lydia Mechtenberg, April 2007.
- 022 "A Generalized ARFIMA Process with Markov-Switching Fractional Differencing Parameter" by Wen-Jen Tsay and Wolfgang Härdle, April 2007.

SFB 649, Spandauer Straße 1, D-10178 Berlin
<http://sfb649.wiwi.hu-berlin.de>

This research was supported by the Deutsche
Forschungsgemeinschaft through the SFB 649 "Economic Risk".

SFB 649, Spandauer Straße 1, D-10178
<http://sfb649.wiwi.hu-berlin.de>



This research was supported by the Deutsche
Forschungsgemeinschaft through the SFB 649 "Economic Risk".

COMMON FUNCTIONAL PRINCIPAL COMPONENTS¹

BY MICHAL BENKO, WOLFGANG HÄRDLE AND ALOIS KNEIP

Humboldt-Universität, Humboldt-Universität and Bonn Universität

Functional principal component analysis (FPCA) based on the Karhunen–Loève decomposition has been successfully applied in many applications, mainly for one sample problems. In this paper we consider common functional principal components for two sample problems. Our research is motivated not only by the theoretical challenge of this data situation, but also by the actual question of dynamics of implied volatility (IV) functions. For different maturities the log-returns of IVs are samples of (smooth) random functions and the methods proposed here study the similarities of their stochastic behavior. First we present a new method for estimation of functional principal components from discrete noisy data. Next we present the two sample inference for FPCA and develop the two sample theory. We propose bootstrap tests for testing the equality of eigenvalues, eigenfunctions, and mean functions of two functional samples, illustrate the test-properties by simulation study and apply the method to the IV analysis.

1. Introduction. In many applications in biometrics, chemometrics, econometrics, etc., the data come from the observation of continuous phenomena of time or space and can be assumed to represent a sample of i.i.d. smooth random functions $X_1(t), \dots, X_n(t) \in L^2[0, 1]$. Functional data analysis has received considerable attention in the statistical literature during the last decade. In this context functional principal component analysis (FPCA) has proved to be a key technique. An early reference is Rao (1958), and important methodological contributions have been given by various authors. Case studies and references, as well as methodological and algorithmical details, can be found in the books by Ramsay and Silverman (2002, 2005) or Ferraty and Vieu (2006).

Received January 2006; revised February 2007.

¹Supported by the Deutsche Forschungsgemeinschaft and the Sonderforschungsbereich 649 “Ökonomisches Risiko.”

AMS 2000 subject classifications. Primary 62H25, 62G08; secondary 62P05.

Key words and phrases. Functional principal components, nonparametric regression, bootstrap, two sample problem.

<p>This is an electronic reprint of the original article published by the Institute of Mathematical Statistics in <i>The Annals of Statistics</i>, 2009, Vol. 37, No. 1, 1–34. This reprint differs from the original in pagination and typographic detail.</p>

The well-known Karhunen–Loève (KL) expansion provides a basic tool to describe the distribution of the random functions X_i and can be seen as the theoretical basis of FPCA. For $v, w \in L^2[0, 1]$, let $\langle v, w \rangle = \int_0^1 v(t)w(t) dt$, and let $\|\cdot\| = \langle \cdot, \cdot \rangle^{1/2}$ denote the usual L^2 -norm. With $\lambda_1 \geq \lambda_2 \geq \dots$ and $\gamma_1, \gamma_2, \dots$ denoting eigenvalues and corresponding orthonormal eigenfunctions of the covariance operator Γ of X_i , we obtain $X_i = \mu + \sum_{r=1}^{\infty} \beta_{ri} \gamma_r$, $i = 1, \dots, n$, where $\mu = E(X_i)$ is the mean function and $\beta_{ri} = \langle X_i - \mu, \gamma_r \rangle$ are (scalar) factor loadings with $E(\beta_{ri}^2) = \lambda_r$. Structure and dynamics of the random functions can be assessed by analyzing the “functional principal components” γ_r , as well as the distribution of the factor loadings. For a given functional sample, the unknown characteristics λ_r, γ_r are estimated by the eigenvalues and eigenfunctions of the empirical covariance operator $\hat{\Gamma}_n$ of X_1, \dots, X_n . Note that an eigenfunction γ_r is identified (up to sign) only if the corresponding eigenvalue λ_r has multiplicity one. This therefore establishes a necessary regularity condition for any inference based on an estimated functional principal component $\hat{\gamma}_r$ in FPCA. Signs are arbitrary (γ_r and β_{ri} can be replaced by $-\gamma_r$ and $-\beta_{ri}$) and may be fixed by a suitable standardization. More detailed discussion on this topic and precise assumptions can be found in Section 2.

In many important applications a small number of functional principal components will suffice to approximate the functions X_i with a high degree of accuracy. Indeed, FPCA plays a much more central role in functional data analysis than its well-known analogue in multivariate analysis. There are two major reasons. First, distributions on function spaces are complex objects, and the Karhunen–Loève expansion seems to be the only practically feasible way to access their structure. Second, in multivariate analysis a substantial interpretation of principal components is often difficult and has to be based on vague arguments concerning the correlation of principal components with original variables. Such a problem does not at all exist in the functional context, where $\gamma_1(t), \gamma_2(t), \dots$ are functions representing the major modes of variation of $X_i(t)$ over t .

In this paper we consider inference and tests of hypotheses on the structure of functional principal components. Motivated by an application to implied volatility analysis, we will concentrate on the two sample case. A central point is the use of bootstrap procedures. We will show that the bootstrap methodology can also be applied to functional data.

In Section 2 we start by discussing one-sample inference for FPCA. Basic results on asymptotic distributions have already been derived by [Dauxois, Pousse and Romain \(1982\)](#) in situations where the functions are directly observable. [Hall and Hosseini-Nasab \(2006\)](#) develop asymptotic Taylor expansions of estimated eigenfunctions in terms of the difference $\hat{\Gamma}_n - \Gamma$.

Without deriving rigorous theoretical results, they also provide some qualitative arguments as well as simulation results motivating the use of bootstrap in order to construct confidence regions for principal components.

In practice, the functions of interest are often not directly observed, but are regression curves which have to be reconstructed from discrete, noisy data. In this context the standard approach is to first estimate individual functions nonparametrically (e.g., by B-splines) and then to determine principal components of the resulting estimated empirical covariance operator—see [Besse and Ramsay \(1986\)](#), [Ramsay and Dalzell \(1991\)](#), among others. Approaches incorporating a smoothing step into the eigenanalysis have been proposed by [Rice and Silverman \(1991\)](#), [Pezzulli and Silverman \(1993\)](#) or [Silverman \(1996\)](#). Robust estimation of principal components has been considered by [Lacontore et al. \(1999\)](#). [Yao, Müller and Wang \(2005\)](#) and [Hall, Müller and Wang \(2006\)](#) propose techniques based on nonparametric estimation of the covariance function $E[\{X_i(t) - \mu(t)\}\{X_i(s) - \mu(s)\}]$ which can also be applied if there are only a few scattered observations per curve.

Section 2.1 presents a new method for estimation of functional principal components. It consists in an adaptation of a technique introduced by [Kneip and Utikal \(2001\)](#) for the case of density functions. The key-idea is to represent the components of the Karhunen–Loève expansion in terms of an (L^2) scalar-product matrix of the sample. We investigate the asymptotic properties of the proposed method. It is shown that under mild conditions the additional error caused by estimation from discrete, noisy data is first-order asymptotically negligible, and inference may proceed “as if” the functions were directly observed. Generalizing the results of [Dauxois, Pousse and Romain \(1982\)](#), we then present a theorem on the asymptotic distributions of the empirical eigenvalues and eigenfunctions. The structure of the asymptotic expansion derived in the theorem provides a basis to show consistency of bootstrap procedures.

Section 3 deals with two-sample inference. We consider two independent samples of functions $\{X_i^{(1)}\}_{i=1}^{n_1}$ and $\{X_i^{(2)}\}_{i=1}^{n_2}$. The problem of interest is to test in how far the distributions of these random functions coincide. The structure of the different distributions in function space can be accessed by means of the respective Karhunen–Loève expansions

$$X_i^{(p)} = \mu^{(p)} + \sum_{r=1}^{\infty} \beta_{ri}^{(p)} \gamma_r^{(p)}, \quad p = 1, 2.$$

Differences in the distribution of these random functions will correspond to differences in the components of the respective KL expansions above. Without restriction, one may require that signs are such that $\langle \gamma_r^{(1)}, \gamma_r^{(2)} \rangle \geq 0$. Two sample inference for FPCA in general has not been considered in the literature so far. In Section 3 we define bootstrap procedures for testing

the equality of mean functions, eigenvalues, eigenfunctions and eigenspaces. Consistency of the bootstrap is derived in Section 3.1, while Section 3.2 contains a simulation study providing insight into the finite sample performance of our tests.

It is of particular interest to compare the functional components characterizing the two samples. If these factors are “common,” this means $\gamma_r := \gamma_r^{(1)} = \gamma_r^{(2)}$, then only the factor loadings $\beta_{ri}^{(p)}$ may vary across samples. This situation may be seen as a functional generalization of the concept of “common principal components” as introduced by Flury (1988) in multivariate analysis. A weaker hypothesis may only require equality of the eigenspaces spanned by the first $L \in \mathbb{N}$ functional principal components. [\mathbb{N} denotes the set of all natural numbers $1, 2, \dots$ ($0 \notin \mathbb{N}$)]. If for both samples the common L -dimensional eigenspaces suffice to approximate the functions with high accuracy, then the distributions in function space are well represented by a low-dimensional factor model, and subsequent analysis may rely on comparing the multivariate distributions of the random vectors $(\beta_{r1}^{(p)}, \dots, \beta_{rL}^{(p)})^\top$.

The idea of “common functional principal components” is of considerable importance in implied volatility (IV) dynamics. This application is discussed in detail in Section 4. Implied volatility is obtained from the pricing model proposed by Black and Scholes (1973) and is a key parameter for quoting options prices. Our aim is to construct low-dimensional factor models for the log-returns of the IV functions of options with different maturities. In our application the first group of functional observations— $\{X_i^{(1)}\}_{i=1}^{n_1}$, are log-returns on the maturity “1 month” (1M group) and second group— $\{X_i^{(2)}\}_{i=1}^{n_2}$, are log-returns on the maturity “3 months” (3M group).

The first three eigenfunctions (ordered with respect to the corresponding eigenvalues), estimated by the method described in Section 2.1, are plotted in Figure 1. The estimated eigenfunctions for both groups are of similar structure, which motivates a common FPCA approach. Based on discretized vectors of functional values, a (multivariate) common principal components analysis of implied volatilities has already been considered by Fengler, Härdle and Villa (2003). They rely on the methodology introduced by Flury (1988) which is based on maximum likelihood estimation under the assumption of multivariate normality. Our analysis overcomes the limitations of this approach by providing specific hypothesis tests in a fully functional setup. It will be shown in Section 4 that for both groups $L = 3$ components suffice to explain 98.2% of the variability of the sample functions. An application of the tests developed in Section 3 does not reject the equality of the corresponding eigenspaces.

2. Functional principal components and one sample inference. In this section we will focus on one sample of i.i.d. smooth random functions $X_1, \dots,$

$X_n \in L^2[0, 1]$. We will assume a well-defined mean function $\mu = E(X_i)$, as well as the existence of a continuous covariance function $\sigma(t, s) = E[\{X_i(t) - \mu(t)\}\{X_i(s) - \mu(s)\}]$. Then $E(\|X_i - \mu\|^2) = \int \sigma(t, t) dt < \infty$, and the covariance operator Γ of X_i is given by

$$(\Gamma v)(t) = \int \sigma(t, s)v(s) ds, \quad v \in L^2[0, 1].$$

The Karhunen–Loève decomposition provides a basic tool to describe the distribution of the random functions X_i . With $\lambda_1 \geq \lambda_2 \geq \dots$ and $\gamma_1, \gamma_2, \dots$ denoting eigenvalues and a corresponding complete orthonormal basis of eigenfunctions of Γ , we obtain

$$(1) \quad X_i = \mu + \sum_{r=1}^{\infty} \beta_{ri} \gamma_r, \quad i = 1, \dots, n,$$

where $\beta_{ri} = \langle X_i - \mu, \gamma_r \rangle$ are uncorrelated (scalar) factor loadings with $E(\beta_{ri}) = 0$, $E(\beta_{ri}^2) = \lambda_r$ and $E(\beta_{ri}\beta_{ki}) = 0$ for $r \neq k$. Structure and dynamics of the random functions can be assessed by analyzing the “functional principal components” γ_r , as well as the distribution of the factor loadings.

A discussion of basic properties of (1) can, for example, be found in [Gihman and Skorohod \(1973\)](#). Under our assumptions, the infinite sums in (1) converge with probability 1, and $\sum_{r=1}^{\infty} \lambda_r = E(\|X_i - \mu\|^2) < \infty$. Smoothness of X_i carries over to a corresponding degree of smoothness of $\sigma(t, s)$ and γ_r . If, with probability 1, $X_i(t)$ is twice continuously differentiable, then σ as well as γ_r are also twice continuously differentiable. The particular case of a Gaussian random function X_i implies that the β_{ri} are independent $N(0, \lambda_r)$ -distributed random variables.




FIG. 1. *Estimated eigenfunctions for 1M group in the left plot and 3M group in the right plot: solid—first function, dashed—second function, finely dashed—third function.*

An important property of (1) consists in the known fact that the first L principal components provide a “best basis” for approximating the sample functions in terms of the integrated square error; see [Ramsay and Silverman \(2005\)](#), Section 6.2.3, among others. For any choice of L orthonormal basis functions v_1, \dots, v_L , the mean integrated square error

$$(2) \quad \rho(v_1, \dots, v_L) = \mathbb{E} \left(\left\| X_i - \mu - \sum_{r=1}^L \langle X_i - \mu, v_r \rangle v_r \right\|^2 \right)$$

is minimized by $v_r = \gamma_r$.

2.1. Estimation of functional principal components. For a given sample an empirical analog of (1) can be constructed by using eigenvalues $\hat{\lambda}_1 \geq \hat{\lambda}_2 \geq \dots$ and orthonormal eigenfunctions $\hat{\gamma}_1, \hat{\gamma}_2, \dots$ of the empirical covariance operator $\hat{\Gamma}_n$, where

$$(\hat{\Gamma}_n v)(t) = \int \hat{\sigma}(t, s) v(s) ds,$$

with $\bar{X} = n^{-1} \sum_{i=1}^n X_i$ and $\hat{\sigma}(t, s) = n^{-1} \sum_{i=1}^n \{X_i(t) - \bar{X}(t)\} \{X_i(s) - \bar{X}(s)\}$ denoting sample mean and covariance function. Then

$$(3) \quad X_i = \bar{X} + \sum_{r=1}^n \hat{\beta}_{ri} \hat{\gamma}_r, \quad i = 1, \dots, n,$$

where $\hat{\beta}_{ri} = \langle \hat{\gamma}_r, X_i - \bar{X} \rangle$. We necessarily obtain $n^{-1} \sum_i \hat{\beta}_{ri} = 0$, $n^{-1} \sum_i \hat{\beta}_{ri} \hat{\beta}_{si} = 0$ for $r \neq s$, and $n^{-1} \sum_i \hat{\beta}_{ri}^2 = \hat{\lambda}_r$.

Analysis will have to concentrate on the leading principal components explaining the major part of the variance. In the following we will assume that $\lambda_1 > \lambda_2 > \dots > \lambda_{r_0} > \lambda_{r_0+1}$, where r_0 denotes the maximal number of components to be considered. For all $r = 1, \dots, r_0$, the corresponding eigenfunction γ_r is then uniquely defined up to sign. Signs are arbitrary, decompositions (1) or (3) may just as well be written in terms of $-\gamma_r, -\beta_{ri}$ or $-\hat{\gamma}_r, -\hat{\beta}_{ri}$, and any suitable standardization may be applied by the statistician. In order to ensure that $\hat{\gamma}_r$ may be viewed as an estimator of γ_r rather than of $-\gamma_r$, we will in the following only assume that signs are such that $\langle \gamma_r, \hat{\gamma}_r \rangle \geq 0$. More generally, any subsequent statement concerning differences of two eigenfunctions will be based on the condition of a nonnegative inner product. This does not impose any restriction and will go without saying.

The results of [Dauxois, Pousse and Romain \(1982\)](#) imply that, under regularity conditions, $\|\hat{\gamma}_r - \gamma_r\| = \mathcal{O}_p(n^{-1/2})$, $|\hat{\lambda}_r - \lambda_r| = \mathcal{O}_p(n^{-1/2})$, as well as $|\hat{\beta}_{ri} - \beta_{ri}| = \mathcal{O}_p(n^{-1/2})$ for all $r \leq r_0$.

However, in practice, the sample functions X_i are often not directly observed, but have to be reconstructed from noisy observations Y_{ij} at discrete

design points t_{ik} :

$$(4) \quad Y_{ik} = X_i(t_{ik}) + \varepsilon_{ik}, \quad k = 1, \dots, T_i,$$

where ε_{ik} are independent noise terms with $E(\varepsilon_{ik}) = 0$, $\text{Var}(\varepsilon_{ik}) = \sigma_i^2$.

Our approach for estimating principal components is motivated by the well-known duality relation between row and column spaces of a data matrix; see [Härdle and Simar \(2003\)](#), Chapter 8, among others. In a first step this approach relies on estimating the elements of the matrix:

$$(5) \quad M_{lk} = \langle X_l - \bar{X}, X_k - \bar{X} \rangle, \quad l, k = 1, \dots, n.$$

Some simple linear algebra shows that all nonzero eigenvalues $\hat{\lambda}_1 \geq \hat{\lambda}_2 \cdots$ of $\hat{\Gamma}_n$ and $l_1 \geq l_2 \cdots$ of M are related by $\hat{\lambda}_r = l_r/n$, $r = 1, 2, \dots$. When using the corresponding orthonormal eigenvectors p_1, p_2, \dots of M , the empirical scores $\hat{\beta}_{ri}$, as well as the empirical eigenfunctions $\hat{\gamma}_r$, are obtained by $\hat{\beta}_{ri} = \sqrt{l_r} p_{ir}$ and

$$(6) \quad \hat{\gamma}_r = \frac{1}{\sqrt{l_r}} \sum_{i=1}^n p_{ir} (X_i - \bar{X}) = \frac{1}{\sqrt{l_r}} \sum_{i=1}^n p_{ir} X_i.$$

The elements of M are functionals which can be estimated with asymptotically negligible bias and a parametric rate of convergence $T_i^{-1/2}$. If the data in (4) is generated from a balanced, equidistant design, then it is easily seen that for $i \neq j$ this rate of convergence is achieved by the estimator

$$\widehat{M}_{ij} = T^{-1} \sum_{k=1}^T (Y_{ik} - \bar{Y}_{\cdot k})(Y_{jk} - \bar{Y}_{\cdot k}), \quad i \neq j,$$

and

$$\widehat{M}_{ii} = T^{-1} \sum_{k=1}^T (Y_{ik} - \bar{Y}_{\cdot k})^2 - \hat{\sigma}_i^2,$$

where $\hat{\sigma}_i^2$ denotes some nonparametric estimator of variance and $\bar{Y}_{\cdot k} = n^{-1} \times \sum_{j=1}^n Y_{jk}$.

In the case of a random design some adjustment is necessary: Define the ordered sample $t_{i(1)} \leq t_{i(2)} \leq \dots \leq t_{i(T_i)}$ of design points, and for $j = 1, \dots, T_i$, let $Y_{i(j)}$ denote the observation belonging to $t_{i(j)}$. With $t_{i(0)} = -t_{i(1)}$ and $t_{i(T_i+1)} = 2 - t_{i(T_i)}$, set

$$\chi_i(t) = \sum_{j=1}^{T_i} Y_{i(j)} I\left(t \in \left[\frac{t_{i(j-1)} + t_{i(j)}}{2}, \frac{t_{i(j)} + t_{i(j+1)}}{2}\right)\right), \quad t \in [0, 1],$$

where $I(\cdot)$ denotes the indicator function, and for $i \neq j$, define the estimate of M_{ij} by

$$\widehat{M}_{ij} = \int_0^1 \{\chi_i(t) - \bar{\chi}(t)\} \{\chi_j(t) - \bar{\chi}(t)\} dt,$$

where $\bar{\chi}(t) = n^{-1} \sum_{i=1}^n \chi_i(t)$. Finally, by redefining $t_{i(1)} = -t_{i(2)}$ and $t_{i(T_i+1)} = 2 - t_{i(T_i)}$, set $\chi_i^*(t) = \sum_{j=2}^{T_i} Y_{i(j-1)} I(t \in [\frac{t_{i(j-1)}+t_{i(j)}}{2}, \frac{t_{i(j)}+t_{i(j+1)}}{2}])$, $t \in [0, 1]$. Then construct estimators of the diagonal terms M_{ii} by

$$(7) \quad \widehat{M}_{ii} = \int_0^1 \{\chi_i(t) - \bar{\chi}(t)\} \{\chi_i^*(t) - \bar{\chi}(t)\} dt.$$

The aim of using the estimator (7) for the diagonal terms is to avoid the additional bias implied by $E_\varepsilon(Y_{ik}^2) = X_i(t_{ij})^2 + \sigma_i^2$. Here E_ε denotes conditional expectation given t_{ij} , X_i . Alternatively, we can construct a bias corrected estimator using some nonparametric estimation of variance σ_i^2 , for example, the difference based model-free variance estimators studied in Hall, Kay and Titterton (1990) can be employed.

The eigenvalues $\hat{l}_1 \geq \hat{l}_2 \cdots$ and eigenvectors $\hat{p}_1, \hat{p}_2, \dots$ of the resulting matrix \widehat{M} then provide estimates $\hat{\lambda}_{r;T} = \hat{l}_r/n$ and $\hat{\beta}_{ri;T} = \sqrt{\hat{l}_r} \hat{p}_{ir}$ of λ_r and β_{ri} . Estimates $\hat{\gamma}_{r;T}$ of the empirical functional principal component $\hat{\gamma}_r$ can be determined from (6) when replacing the unknown true functions X_i by nonparametric estimates \hat{X}_i (as, for example, local polynomial estimates) with smoothing parameter (bandwidth) b :

$$(8) \quad \hat{\gamma}_{r;T} = \frac{1}{\sqrt{\hat{l}_r}} \sum_{i=1}^n \hat{p}_{ir} \hat{X}_i.$$

When considering (8), it is important to note that $\hat{\gamma}_{r;T}$ is defined as a *weighted average* of all estimated sample functions. Averaging reduces variance, and efficient estimation of $\hat{\gamma}_r$ therefore requires *undersmoothing* of individual function estimates \hat{X}_i . Theoretical results are given in Theorem 1 below. Indeed, if, for example, n and $T = \min_i T_i$ are of the same order of magnitude, then under suitable additional regularity conditions it will be shown that for an optimal choice of a smoothing parameter $b \sim (nT)^{-1/5}$ and twice continuously differentiable X_i , we obtain the rate of convergence $\|\hat{\gamma}_r - \hat{\gamma}_{r;T}\| = \mathcal{O}_p\{(nT)^{-2/5}\}$. Note, however, that the bias corrected estimator (7) may yield negative eigenvalues. In practice, these values will be small and will have to be interpreted as zero. Furthermore, the eigenfunctions determined by (8) may not be exactly orthogonal. Again, when using reasonable bandwidths, this effect will be small, but of course (8) may be followed by a suitable orthogonalization procedure.

It is of interest to compare our procedure to more standard methods for estimating λ_r and $\hat{\gamma}_r$ as mentioned above. When evaluating eigenvalues and eigenfunctions of the empirical covariance operator of nonparametrically estimated curves \hat{X}_i , then for fixed $r \leq r_0$ the above rate of convergence for the estimated eigenfunctions may well be achieved for a suitable choice of smoothing parameters (e.g., number of basis functions). But as will be seen

from Theorem 1, our approach also implies that $|\hat{\lambda}_r - \frac{\hat{t}_r}{n}| = \mathcal{O}_p(T^{-1} + n^{-1})$. When using standard methods it does not seem to be possible to obtain a corresponding rate of convergence, since any smoothing bias $|\mathbb{E}[\hat{X}_i(t)] - X_i(t)|$ will invariably affect the quality of the corresponding estimate of $\hat{\lambda}_r$.

We want to emphasize that any finite sample interpretation will require that T is sufficiently large such that our nonparametric reconstructions of individual curves can be assumed to possess a fairly small bias. The above arguments do not apply to extremely sparse designs with very few observations per curve [see Hall, Müller and Wang (2006) for an FPCA methodology focusing on sparse data].

Note that, in addition to (8), our final estimate of the empirical mean function $\hat{\mu} = \bar{X}$ will be given by $\hat{\mu}_T = n^{-1} \sum_i \hat{X}_i$. A straightforward approach to determine a suitable bandwidth b consists in a “leave-one-individual-out” cross-validation. For the maximal number r_0 of components to be considered, let $\hat{\mu}_{T,-i}$ and $\hat{\gamma}_{r;T,-i}$, $r = 1, \dots, r_0$, denote the estimates of $\hat{\mu}$ and $\hat{\gamma}_r$ obtained from the data (Y_{lj}, t_{lj}) , $l = 1, \dots, i-1, i+1, \dots, n$, $j = 1, \dots, T_k$. By (8), these estimates depend on b , and one may approximate an optimal smoothing parameter by minimizing

$$\sum_i \sum_j \left\{ Y_{ij} - \hat{\mu}_{T,-i}(t_{ij}) - \sum_{r=1}^{r_0} \hat{\vartheta}_{ri} \hat{\gamma}_{r;T,-i}(t_{ij}) \right\}^2$$

over b , where $\hat{\vartheta}_{ri}$ denote ordinary least squares estimates of $\hat{\beta}_{ri}$. A more sophisticated version of this method may even allow to select different bandwidths b_r when estimating different functional principal components by (8). Although, under certain regularity conditions, the same qualitative rates of convergence hold for any arbitrary fixed $r \leq r_0$, the quality of estimates decreases when r becomes large. Due to $\langle \gamma_s, \gamma_r \rangle = 0$ for $s < r$, the number of zero crossings, peaks and valleys of γ_r has to increase with r . Hence, in tendency γ_r will be less and less smooth as r increases. At the same time, $\lambda_r \rightarrow 0$, which means that for large r the r th eigenfunctions will only possess a very small influence on the structure of X_i . This in turn means that the relative importance of the error terms ε_{ik} in (4) on the structure of $\hat{\gamma}_{r;T}$ will increase with r .

2.2. One sample inference. Clearly, in the framework described by (1)–(4) we are faced with two sources of variability of estimated functional principal components. Due to sampling variation, $\hat{\gamma}_r$ will differ from the true component γ_r , and due to (4), there will exist an additional estimation error when approximating $\hat{\gamma}_r$ by $\hat{\gamma}_{r;T}$.

The following theorems quantify the order of magnitude of these different types of error. Our theoretical results are based on the following assumptions on the structure of the random functions X_i .

ASSUMPTION 1. $X_1, \dots, X_n \in L^2[0, 1]$ is an i.i.d. sample of random functions with mean μ and continuous covariance function $\sigma(t, s)$, and (1) holds for a system of eigenfunctions satisfying $\sup_{s \in \mathbb{N}} \sup_{t \in [0, 1]} |\gamma_s(t)| < \infty$. Furthermore, $\sum_{r=1}^{\infty} \sum_{s=1}^{\infty} \mathbb{E}[\beta_{ri}^2 \beta_{si}^2] < \infty$ and $\sum_{q=1}^{\infty} \sum_{s=1}^{\infty} \mathbb{E}[\beta_{ri}^2 \beta_{qi} \beta_{si}] < \infty$ for all $r \in \mathbb{N}$.

Recall that $\mathbb{E}[\beta_{ri}] = 0$ and $\mathbb{E}[\beta_{ri} \beta_{si}] = 0$ for $r \neq s$. Note that the assumption on the factor loadings is necessarily fulfilled if X_i are Gaussian random functions. Then β_{ri} and β_{si} are independent for $r \neq s$, all moments of β_{ri} are finite, and hence $\mathbb{E}[\beta_{ri}^2 \beta_{qi} \beta_{si}] = 0$ for $q \neq s$, as well as $\mathbb{E}[\beta_{ri}^2 \beta_{si}^2] = \lambda_r \lambda_s$ for $r \neq s$; see [Gihman and Skorohod \(1973\)](#).

We need some further assumptions concerning smoothness of X_i and the structure of the discrete model (4).

ASSUMPTION 2. (a) X_i is a.s. twice continuously differentiable. There exists a constant $D_1 < \infty$ such that the derivatives are bounded by $\sup_t \mathbb{E}[X_i'(t)^4] \leq D_1$, as well as $\sup_t \mathbb{E}[X_i''(t)^4] \leq D_1$.

(b) The design points t_{ik} , $i = 1, \dots, n$, $k = 1, \dots, T_i$, are i.i.d. random variables which are independent of X_i and ε_{ik} . The corresponding design density f is continuous on $[0, 1]$ and satisfies $\inf_{t \in [0, 1]} f(t) > 0$.

(c) For any i , the error terms ε_{ik} are i.i.d. zero mean random variables with $\text{Var}(\varepsilon_{ik}) = \sigma_i^2$. Furthermore, ε_{ik} is independent of X_i , and there exists a constant D_2 such that $\mathbb{E}(\varepsilon_{ik}^8) < D_2$ for all i, k .

(d) The estimates \hat{X}_i used in (8) are determined by either a local linear or a Nadaraya–Watson kernel estimator with smoothing parameter b and kernel function K . K is a continuous probability density which is symmetric at 0.

The following theorems provide asymptotic results as $n, T \rightarrow \infty$, where $T = \min_{i=1}^n \{T_i\}$.

THEOREM 1. *In addition to Assumptions 1 and 2, assume that $\inf_{s \neq r} |\lambda_r - \lambda_s| > 0$ holds for some $r = 1, 2, \dots$. Then we have the following:*

(i) $n^{-1} \sum_{i=1}^n (\hat{\beta}_{ri} - \hat{\beta}_{ri;T})^2 = \mathcal{O}_p(T^{-1})$ and

$$(9) \quad \left| \hat{\lambda}_r - \frac{\hat{l}_r}{n} \right| = \mathcal{O}_p(T^{-1} + n^{-1}).$$

(ii) *If additionally $b \rightarrow 0$ and $(Tb)^{-1} \rightarrow 0$ as $n, T \rightarrow \infty$, then for all $t \in [0, 1]$,*

$$(10) \quad |\hat{\gamma}_r(t) - \hat{\gamma}_{r;T}(t)| = \mathcal{O}_p\{b^2 + (nTb)^{-1/2} + (Tb^{1/2})^{-1} + n^{-1}\}.$$

A proof is given in the [Appendix](#).

THEOREM 2. Under Assumption 1 we obtain the following:

(i) For all $t \in [0, 1]$,

$$\sqrt{n}\{\bar{X}(t) - \mu(t)\} = \sum_r \left\{ \frac{1}{\sqrt{n}} \sum_{i=1}^n \beta_{ri} \right\} \gamma_r(t) \xrightarrow{\mathcal{L}} N\left(0, \sum_r \lambda_r \gamma_r(t)^2\right).$$

If, furthermore, $\lambda_{r-1} > \lambda_r > \lambda_{r+1}$ holds for some fixed $r \in \{1, 2, \dots\}$, then

(ii)

$$(11) \quad \sqrt{n}(\hat{\lambda}_r - \lambda_r) = \frac{1}{\sqrt{n}} \sum_{i=1}^n (\beta_{ri}^2 - \lambda_r) + \mathcal{O}_p(n^{-1/2}) \xrightarrow{\mathcal{L}} N(0, \Lambda_r),$$

where $\Lambda_r = \mathbb{E}[(\beta_{ri}^2 - \lambda_r)^2]$,

(iii) and for all $t \in [0, 1]$

$$(12) \quad \hat{\gamma}_r(t) - \gamma_r(t) = \sum_{s \neq r} \left\{ \frac{1}{n(\lambda_r - \lambda_s)} \sum_{i=1}^n \beta_{si} \beta_{ri} \right\} \gamma_s(t) + R_r(t),$$

where $\|R_r\| = \mathcal{O}_p(n^{-1})$.

Moreover,

$$\begin{aligned} & \sqrt{n} \sum_{s \neq r} \left\{ \frac{1}{n(\lambda_r - \lambda_s)} \sum_{i=1}^n \beta_{si} \beta_{ri} \right\} \gamma_s(t) \\ & \xrightarrow{\mathcal{L}} N\left(0, \sum_{q \neq r} \sum_{s \neq r} \frac{\mathbb{E}[\beta_{ri}^2 \beta_{qi} \beta_{si}]}{(\lambda_q - \lambda_r)(\lambda_s - \lambda_r)} \gamma_q(t) \gamma_s(t)\right). \end{aligned}$$

A proof can be found in the [Appendix](#). The theorem provides a generalization of the results of [Dauxois, Pousse and Romain \(1982\)](#) who derive explicit asymptotic distributions by assuming Gaussian random functions X_i . Note that in this case $\Lambda_r = 2\lambda_r^2$ and $\sum_{q \neq r} \sum_{s \neq r} \frac{\mathbb{E}[\beta_{ri}^2 \beta_{qi} \beta_{si}]}{(\lambda_q - \lambda_r)(\lambda_s - \lambda_r)} \gamma_q(t) \gamma_s(t) = \sum_{s \neq r} \frac{\lambda_r \lambda_s}{(\lambda_s - \lambda_r)^2} \gamma_s(t)^2$.

When evaluating the bandwidth-dependent terms in (10), best rates of convergence $|\hat{\gamma}_r(t) - \hat{\gamma}_{r;T}(t)| = \mathcal{O}_p\{(nT)^{-2/5} + T^{-4/5} + n^{-1}\}$ are achieved when choosing an undersmoothing bandwidth $b \sim \max\{(nT)^{-1/5}, T^{-2/5}\}$. Theoretical work in functional data analysis is usually based on the implicit assumption that the additional error due to (4) is negligible, and that one can proceed ‘‘as if’’ the functions X_i were directly observed. In view of Theorems 1 and 2, this approach is justified in the following situations:

(1) T is much larger than n , that is, $n/T^{4/5} \rightarrow 0$, and the smoothing parameter b in (8) is of order $T^{-1/5}$ (optimal smoothing of individual functions).

(2) An undersmoothing bandwidth $b \sim \max\{(nT)^{-1/5}, T^{-2/5}\}$ is used and $n/T^{8/5} \rightarrow 0$. This means that T may be smaller than n , but T must be at least of order of magnitude larger than $n^{5/8}$.

In both cases (1) and (2) the above theorems imply that $|\hat{\lambda}_r - \frac{\hat{L}_r}{n}| = \mathcal{O}_p(|\hat{\lambda}_r - \lambda_r|)$, as well as $\|\hat{\gamma}_r - \hat{\gamma}_{r;T}\| = \mathcal{O}_p(\|\hat{\gamma}_r - \gamma_r\|)$. Inference about functional principal components will then be first-order equivalent to an inference based on known functions X_i .

In such situations Theorem 2 suggests bootstrap procedures as tools for one sample inference. For example, the distribution of $\|\hat{\gamma}_r - \gamma_r\|$ may be approximated by the bootstrap distribution of $\|\hat{\gamma}_r^* - \hat{\gamma}_r\|$, where $\hat{\gamma}_r^*$ are estimates to be obtained from i.i.d. bootstrap resamples $X_1^*, X_2^*, \dots, X_n^*$ of $\{X_1, X_2, \dots, X_n\}$. This means that $X_1^* = X_{i_1}, \dots, X_n^* = X_{i_n}$ for some indices i_1, \dots, i_n drawn independently and with replacement from $\{1, \dots, n\}$ and, in practice, $\hat{\gamma}_r^*$ may thus be approximated from corresponding discrete data $(Y_{i_1 j}, t_{i_1 j})_{j=1, \dots, T_{i_1}}, \dots, (Y_{i_n j}, t_{i_n j})_{j=1, \dots, T_{i_n}}$. The additional error is negligible if either (1) or (2) is satisfied.

One may wonder about the validity of such a bootstrap. Functions are complex objects and there is no established result in bootstrap theory which readily generalizes to samples of random functions. But by (1), i.i.d. bootstrap resamples $\{X_i^*\}_{i=1, \dots, n}$ may be equivalently represented by corresponding, i.i.d. resamples $\{\beta_{1i}^*, \beta_{2i}^*, \dots\}_{i=1, \dots, n}$ of factor loadings. Standard multivariate bootstrap theorems imply that for any $q \in \mathbb{N}$ the distribution of moments of the random vectors $(\beta_{1i}, \dots, \beta_{qi})$ may be consistently approximated by the bootstrap distribution of corresponding moments of $(\beta_{1i}^*, \dots, \beta_{qi}^*)$. Together with some straightforward limit arguments as $q \rightarrow \infty$, the structure of the first-order terms in the asymptotic expansions (11) and (12) then allows to establish consistency of the functional bootstrap. These arguments will be made precise in the proof of Theorem 3 below, which concerns related bootstrap statistics in two sample problems.

REMARK. Theorem 2(iii) implies that the variance of $\hat{\gamma}_r$ is large if one of the differences $\lambda_{r-1} - \lambda_r$ or $\lambda_r - \lambda_{r+1}$ is small. In the limit case of eigenvalues of multiplicity $m > 1$ our theory does not apply. Note that then only the m -dimensional eigenspace is identified, but not a particular basis (eigenfunctions). In multivariate PCA Tyler (1981) provides some inference results on corresponding projection matrices assuming that $\lambda_r > \lambda_{r+1} \geq \dots \geq \lambda_{r+m} > \lambda_{r+m+1}$ for known values of r and m .

Although the existence of eigenvalues λ_r , $r \leq r_0$, with multiplicity $m > 1$ may be considered as a degenerate case, it is immediately seen that $\lambda_r \rightarrow 0$ and, hence, $\lambda_r - \lambda_{r+1} \rightarrow 0$ as r increases. Even in the case of fully observed

functions X_i , estimates of eigenfunctions corresponding to very small eigenvalues will thus be poor. The problem of determining a sensible upper limit of the number r_0 of principal components to be analyzed is addressed in [Hall and Hosseini-Nasab \(2006\)](#).

3. Two sample inference. The comparison of functional components across groups leads naturally to two sample problems. Thus, let

$$X_1^{(1)}, X_2^{(1)}, \dots, X_{n_1}^{(1)} \quad \text{and} \quad X_1^{(2)}, X_2^{(2)}, \dots, X_{n_2}^{(2)}$$

denote two independent samples of smooth functions. The problem of interest is to test in how far the distributions of these random functions coincide. The structure of the different distributions in function space can be accessed by means of the respective Karhunen–Loève decompositions. The problem to be considered then translates into testing equality of the different components of these decompositions given by

$$(13) \quad X_i^{(p)} = \mu^{(p)} + \sum_{r=1}^{\infty} \beta_{ri}^{(p)} \gamma_r^{(p)}, \quad p = 1, 2,$$

where again $\gamma_r^{(p)}$ are the eigenfunctions of the respective covariance operator $\Gamma^{(p)}$ corresponding to the eigenvalues $\lambda_1^{(p)} = \text{E}\{(\beta_{1i}^{(p)})^2\} \geq \lambda_2^{(p)} = \text{E}\{(\beta_{2i}^{(p)})^2\} \geq \dots$. We will again suppose that $\lambda_{r-1}^{(p)} > \lambda_r^{(p)} > \lambda_{r+1}^{(p)}$, $p = 1, 2$, for all $r \leq r_0$ components to be considered. Without restriction, we will additionally assume that signs are such that $\langle \gamma_r^{(1)}, \gamma_r^{(2)} \rangle \geq 0$, as well as $\langle \hat{\gamma}_r^{(1)}, \hat{\gamma}_r^{(2)} \rangle \geq 0$.

It is of great interest to detect possible variations in the functional components characterizing the two samples in (13). Significant difference may give rise to substantial interpretation. Important hypotheses to be considered thus are as follows:

$$H_{01} : \mu^{(1)} = \mu^{(2)} \quad \text{and} \quad H_{02,r} : \gamma_r^{(1)} = \gamma_r^{(2)}, \quad r \leq r_0.$$

Hypothesis $H_{02,r}$ is of particular importance. Then $\gamma_r^{(1)} = \gamma_r^{(2)}$ and only the factor loadings β_{ri} may vary across samples. If, for example, $H_{02,r}$ is accepted, one may additionally want to test hypotheses about the distributions of $\beta_{ri}^{(p)}$, $p = 1, 2$. Recall that necessarily $\text{E}\{\beta_{ri}^{(p)}\} = 0$, $\text{E}\{\beta_{ri}^{(p)}\}^2 = \lambda_r^{(p)}$, and $\beta_{si}^{(p)}$ is uncorrelated with $\beta_{ri}^{(p)}$ if $r \neq s$. If the $X_i^{(p)}$ are Gaussian random variables, the $\beta_{ri}^{(p)}$ are independent $N(0, \lambda_r)$ random variables. A natural hypothesis to be tested then refers to the equality of variances:

$$H_{03,r} : \lambda_r^{(1)} = \lambda_r^{(2)}, \quad r = 1, 2, \dots$$

Let $\hat{\mu}^{(p)}(t) = \frac{1}{n_p} \sum_i X_i^{(p)}(t)$, and let $\hat{\lambda}_1^{(p)} \geq \hat{\lambda}_2^{(p)} \geq \dots$ and $\hat{\gamma}_1^{(p)}, \hat{\gamma}_2^{(p)}, \dots$ denote eigenvalues and corresponding eigenfunctions of the empirical covariance operator $\hat{\Gamma}_{n_p}^{(p)}$ of $X_1^{(p)}, X_2^{(p)}(t), \dots, X_{n_p}^{(p)}$. The following test statistics are

defined in terms of $\hat{\mu}^{(p)}$, $\hat{\lambda}_r^{(p)}$ and $\hat{\gamma}_r^{(p)}$. As discussed in the proceeding section, all curves in both samples are usually not directly observed, but have to be reconstructed from noisy observations according to (4). In this situation, the “true” empirical eigenvalues and eigenfunctions have to be replaced by their discrete sample estimates. Bootstrap estimates are obtained by resampling the observations corresponding to the unknown curves $X_i^{(p)}$. As discussed in Section 2.2, the validity of our test procedures is then based on the assumption that T is sufficiently large such that the additional estimation error is asymptotically negligible.

Our tests of the hypotheses H_{0_1} , $H_{0_{2,r}}$ and $H_{0_{3,r}}$ rely on the statistics

$$\begin{aligned} D_1 &\stackrel{\text{def}}{=} \|\hat{\mu}^{(1)} - \hat{\mu}^{(2)}\|^2, \\ D_{2,r} &\stackrel{\text{def}}{=} \|\hat{\gamma}_r^{(1)} - \hat{\gamma}_r^{(2)}\|^2, \\ D_{3,r} &\stackrel{\text{def}}{=} |\hat{\lambda}_r^{(1)} - \hat{\lambda}_r^{(2)}|^2. \end{aligned}$$

The respective null-hypothesis has to be rejected if $D_1 \geq \Delta_{1;1-\alpha}$, $D_{2,r} \geq \Delta_{2,r;1-\alpha}$ or $D_{3,r} \geq \Delta_{3,r;1-\alpha}$, where $\Delta_{1;1-\alpha}$, $\Delta_{2,r;1-\alpha}$ and $\Delta_{3,r;1-\alpha}$ denote the critical values of the distributions of

$$\begin{aligned} \Delta_1 &\stackrel{\text{def}}{=} \|\hat{\mu}^{(1)} - \mu^{(1)} - (\hat{\mu}^{(2)} - \mu^{(2)})\|^2, \\ \Delta_{2,r} &\stackrel{\text{def}}{=} \|\hat{\gamma}_r^{(1)} - \gamma_r^{(1)} - (\hat{\gamma}_r^{(2)} - \gamma_r^{(2)})\|^2, \\ \Delta_{3,r} &\stackrel{\text{def}}{=} |\hat{\lambda}_r^{(1)} - \lambda_r^{(1)} - (\hat{\lambda}_r^{(2)} - \lambda_r^{(2)})|^2. \end{aligned}$$

Of course, the distributions of the different Δ 's cannot be accessed directly, since they depend on the unknown true population mean, eigenvalues and eigenfunctions. However, it will be shown below that these distributions and, hence, their critical values are approximated by the bootstrap distribution of

$$\begin{aligned} \Delta_1^* &\stackrel{\text{def}}{=} \|\hat{\mu}^{(1)*} - \hat{\mu}^{(1)} - (\hat{\mu}^{(2)*} - \hat{\mu}^{(2)})\|^2, \\ \Delta_{2,r}^* &\stackrel{\text{def}}{=} \|\hat{\gamma}_r^{(1)*} - \hat{\gamma}_r^{(1)} - (\hat{\gamma}_r^{(2)*} - \hat{\gamma}_r^{(2)})\|^2, \\ \Delta_{3,r}^* &\stackrel{\text{def}}{=} |\hat{\lambda}_r^{(1)*} - \hat{\lambda}_r^{(1)} - (\hat{\lambda}_r^{(2)*} - \hat{\lambda}_r^{(2)})|^2, \end{aligned}$$

where $\hat{\mu}^{(1)*}$, $\hat{\gamma}_r^{(1)*}$, $\hat{\lambda}_r^{(1)*}$, as well as $\hat{\mu}^{(2)*}$, $\hat{\gamma}_r^{(2)*}$, $\hat{\lambda}_r^{(2)*}$, are estimates to be obtained from independent bootstrap samples $X_1^{1*}(t)$, $X_2^{1*}(t)$, \dots , $X_{n_1}^{1*}(t)$, as well as $X_1^{2*}(t)$, $X_2^{2*}(t)$, \dots , $X_{n_2}^{2*}(t)$.

This test procedure is motivated by the following insights:

(1) Under each of our null-hypotheses the respective test statistics D is equal to the corresponding Δ . The test will thus asymptotically possess the correct level: $P(D > \Delta_{1-\alpha}) \approx \alpha$.

(2) If the null hypothesis is false, then $D \neq \Delta$. Compared to the distribution of Δ , the distribution of D is shifted by the difference in the true means, eigenfunctions or eigenvalues. In tendency D will be larger than $\Delta_{1-\alpha}$.

Let $1 < L \leq r_0$. Even if for $r \leq L$ the equality of eigenfunctions is rejected, we may be interested in the question of whether at least the L -dimensional eigenspaces generated by the first L eigenfunctions are identical. Therefore, let $\mathcal{E}_L^{(1)}$, as well as $\mathcal{E}_L^{(2)}$, denote the L -dimensional linear function spaces generated by the eigenfunctions $\gamma_1^{(1)}, \dots, \gamma_L^{(1)}$ and $\gamma_1^{(2)}, \dots, \gamma_L^{(2)}$, respectively. We then aim to test the null hypothesis:

$$H_{0_{4,L}} : \mathcal{E}_L^{(1)} = \mathcal{E}_L^{(2)}.$$

Of course, $H_{0_{4,L}}$ corresponds to the hypothesis that the operators projecting into $\mathcal{E}_L^{(1)}$ and $\mathcal{E}_L^{(2)}$ are identical. This in turn translates into the condition that

$$\sum_{r=1}^L \gamma_r^{(1)}(t)\gamma_r^{(1)}(s) = \sum_{r=1}^L \gamma_r^{(2)}(t)\gamma_r^{(2)}(s) \quad \text{for all } t, s \in [0, 1].$$

Similar to above, a suitable test statistic is given by

$$D_{4,L} \stackrel{\text{def}}{=} \iint \left\{ \sum_{r=1}^L \hat{\gamma}_r^{(1)}(t)\hat{\gamma}_r^{(1)}(s) - \sum_{r=1}^L \hat{\gamma}_r^{(2)}(t)\hat{\gamma}_r^{(2)}(s) \right\}^2 dt ds$$

and the null hypothesis is rejected if $D_{4,L} \geq \Delta_{4,L;1-\alpha}$, where $\Delta_{4,L;1-\alpha}$ denotes the critical value of the distribution of

$$\Delta_{4,L} \stackrel{\text{def}}{=} \iint \left[\sum_{r=1}^L \{ \hat{\gamma}_r^{(1)}(t)\hat{\gamma}_r^{(1)}(s) - \gamma_r^{(1)}(t)\gamma_r^{(1)}(s) \} - \sum_{r=1}^L \{ \hat{\gamma}_r^{(2)}(t)\hat{\gamma}_r^{(2)}(s) - \gamma_r^{(2)}(t)\gamma_r^{(2)}(s) \} \right]^2 dt ds.$$

The distribution of $\Delta_{4,L}$ and, hence, its critical values are approximated by the bootstrap distribution of

$$\Delta_{4,L}^* \stackrel{\text{def}}{=} \iint \left[\sum_{r=1}^L \{ \hat{\gamma}_r^{(1)*}(t)\hat{\gamma}_r^{(1)*}(s) - \hat{\gamma}_r^{(1)}(t)\hat{\gamma}_r^{(1)}(s) \} - \sum_{r=1}^L \{ \hat{\gamma}_r^{(2)*}(t)\hat{\gamma}_r^{(2)*}(s) - \hat{\gamma}_r^{(2)}(t)\hat{\gamma}_r^{(2)}(s) \} \right]^2 dt ds.$$

It will be shown in Theorem 3 below that under the null hypothesis, as well as under the alternative, the distributions of $n\Delta_1, n\Delta_{2,r}, n\Delta_{3,r}, n\Delta_{4,L}$ converge to continuous limit distributions which can be consistently approximated by the bootstrap distributions of $n\Delta_1^*, n\Delta_{2,r}^*, n\Delta_{3,r}^*, n\Delta_{4,L}^*$.

3.1. *Theoretical results.* Let $n = (n_1 + n_2)/2$. We will assume that asymptotically $n_1 = n \cdot q_1$ and $n_2 = n \cdot q_2$ for some fixed proportions q_1 and q_2 . We will then study the asymptotic behavior of our statistics as $n \rightarrow \infty$.

We will use $\mathcal{X}_1 = \{X_1^{(1)}, \dots, X_{n_1}^{(1)}\}$ and $\mathcal{X}_2 = \{X_1^{(2)}, \dots, X_{n_2}^{(2)}\}$ to denote the observed samples of random functions.

THEOREM 3. *Assume that $\{X_1^{(1)}, \dots, X_{n_1}^{(1)}\}$ and $\{X_1^{(2)}, \dots, X_{n_2}^{(2)}\}$ are two independent samples of random functions, each of which satisfies Assumption 1. As $n \rightarrow \infty$ we then obtain the following:*

(i) *There exists a nondegenerated, continuous probability distribution F_1 such that $n\Delta_1 \xrightarrow{\mathcal{L}} F_1$, and for any $\delta > 0$,*

$$|P(n\Delta_1 \geq \delta) - P(n\Delta_1^* \geq \delta | \mathcal{X}_1, \mathcal{X}_2)| = o_p(1).$$

(ii) *If, furthermore, $\lambda_{r-1}^{(1)} > \lambda_r^{(1)} > \lambda_{r+1}^{(1)}$ and $\lambda_{r-1}^{(2)} > \lambda_r^{(2)} > \lambda_{r+1}^{(2)}$ hold for some fixed $r = 1, 2, \dots$, there exist a nondegenerated, continuous probability distributions $F_{k,r}$ such that $n\Delta_{k,r} \xrightarrow{\mathcal{L}} F_{k,r}$, $k = 2, 3$, and for any $\delta > 0$,*

$$|P(n\Delta_{k,r} \geq \delta) - P(n\Delta_{k,r}^* \geq \delta | \mathcal{X}_1, \mathcal{X}_2)| = o_p(1), \quad k = 2, 3.$$

(iii) *If $\lambda_r^{(1)} > \lambda_{r+1}^{(1)} > 0$ and $\lambda_r^{(2)} > \lambda_{r+1}^{(2)} > 0$ hold for all $r = 1, \dots, L$, there exists a nondegenerated, continuous probability distribution $F_{4,L}$ such that $n\Delta_{4,L} \xrightarrow{\mathcal{L}} F_{4,L}$, and for any $\delta > 0$,*

$$|P(n\Delta_{4,L} \geq \delta) - P(n\Delta_{4,L}^* \geq \delta | \mathcal{X}_1, \mathcal{X}_2)| = o_p(1).$$

The structures of the distributions $F_1, F_{2,r}, F_{3,r}, F_{4,L}$ are derived in the proof of the theorem which can be found in the [Appendix](#). They are obtained as limits of distributions of quadratic forms.

3.2. *Simulation study.* In this paragraph we illustrate the finite behavior of the proposed test. The basic simulation-setup (setup ‘‘a’’) is established as follows: the first sample is generated by the random combination of orthonormalized sine and cosine functions (Fourier functions) and the second sample is generated by the random combination of the same but shifted factor functions:

$$\begin{aligned} X_i^{(1)}(t_{ik}) &= \beta_{1i}^{(1)} \sqrt{2} \sin(2\pi t_{ik}) + \beta_{2i}^{(1)} \sqrt{2} \cos(2\pi t_{ik}), \\ X_i^{(2)}(t_{ik}) &= \beta_{1i}^{(2)} \sqrt{2} \sin\{2\pi(t_{ik} + \delta)\} + \beta_{2i}^{(2)} \sqrt{2} \cos\{2\pi(t_{ik} + \delta)\}. \end{aligned}$$

The factor loadings are i.i.d. random variables with $\beta_{1i}^{(p)} \sim N(0, \lambda_1^{(p)})$ and $\beta_{2i}^{(p)} \sim N(0, \lambda_2^{(p)})$. The functions are generated on the equidistant grid $t_{ik} = t_k = k/T$, $k = 1, \dots, T = 100$, $i = 1, \dots, n = 70$. The simulation setup is based

TABLE 1

The results of the simulations for $\alpha = 0.1$, $n = 70$, $T = 100$, number of simulations 250

Setup/shift	0	0.05	0.1	0.15	0.2	0.25
(a) 10, 5, 8, 4	0.13	0.41	0.85	0.96	1	1
(a) 4, 2, 2, 1	0.12	0.48	0.87	0.96	1	1
(a) 2, 1, 1.5, 2	0.14	0.372	0.704	0.872	0.92	0.9
(b) 10, 5, 8, 4 D_1	0.10	0.44	0.86	0.95	1	1
(b) 10, 5, 8, 4 D_2	1	1	1	1	1	1

on the fact that the error of the estimation of the eigenfunctions simulated by sine and cosine functions is, in particular, manifested by some shift of the estimated eigenfunctions. The focus of this simulation study is the test of common eigenfunctions.

For the presentation of results in Table 1, we use the following notation: “(a) $\lambda_1^{(1)}, \lambda_2^{(1)}, \lambda_2^{(2)}, \lambda_2^{(2)}$.” The shift parameter δ is changing from 0 to 0.25 with the step 0.05. It should be mentioned that the shift $\delta = 0$ yields the simulation of level and setup with shift $\delta = 0.25$ yields the simulation of the alternative, where the two factor functions are exchanged.

In the second setup (setup “b”) the first factor functions are the same and the second factor functions differ:

$$X_i^{(1)}(t_{ik}) = \beta_{1i}^{(1)} \sqrt{2} \sin(2\pi t_{ik}) + \beta_{2i}^{(1)} \sqrt{2} \cos(2\pi t_{ik}),$$

$$X_i^{(2)}(t_{ik}) = \beta_{1i}^{(2)} \sqrt{2} \sin\{2\pi(t_{ik} + \delta)\} + \beta_{2i}^{(2)} \sqrt{2} \sin\{4\pi(t_{ik} + \delta)\}.$$

In Table 1 we use the notation “(b) $\lambda_1^{(1)}, \lambda_2^{(1)}, \lambda_2^{(2)}, \lambda_2^{(2)}, D_r$.” D_r means the test for the equality of the r th eigenfunction. In the bootstrap tests we used 500 bootstrap replications. The critical level in this simulation is $\alpha = 0.1$. The number of simulations is 250.

We can interpret Table 1 in the following way: In power simulations ($\delta \neq 0$) test behaves as expected: less powerful if the functions are “hardly distinguishable” (small shift, small difference in eigenvalues). The level approximation seems to be less precise if the difference in the eigenvalues ($\lambda_1^{(p)} - \lambda_2^{(p)}$) becomes smaller. This can be explained by relative small sample-size n , small number of bootstrap-replications and increasing estimation-error as argued in Theorem 2, assertion (iii).

In comparison to our general setup (4), we used an equidistant and common design for all functions. This simplification is necessary, it simplifies and speeds-up the simulations, in particular, using general random and observation-specific design makes the simulation computationally untractable.

Second, we omitted the additional observation error, this corresponds to the standard assumptions in the functional principal components theory. As

TABLE 2
The results of the simulation for $\alpha = 0.1$, $n = 70$, $T = 100$ with additional error in observation

Setup/shift	0	0.05	0.1	0.15	0.2	0.25
(a) 10, 5, 8, 4	0.09	0.35	0.64	0.92	0.94	0.97

argued in Section 2.2, the inference based on the directly observed functions and estimated functions X_i is first-order equivalent under mild conditions implied by Theorems 1 and 2. In order to illustrate this theoretical result in the simulation, we used the following setup:

$$X_i^{(1)}(t_{ik}) = \beta_{1i}^{(1)} \sqrt{2} \sin(2\pi t_{ik}) + \beta_{2i}^{(1)} \sqrt{2} \cos(2\pi t_{ik}) + \varepsilon_{ik}^{(1)},$$

$$X_i^{(2)}(t_{ik}) = \beta_{1i}^{(2)} \sqrt{2} \sin\{2\pi(t_{ik} + \delta)\} + \beta_{2i}^{(2)} \sqrt{2} \cos\{2\pi(t_{ik} + \delta)\} + \varepsilon_{ik}^{(2)},$$

where $\varepsilon_{ik}^{(p)} \sim N(0, 0.25)$, $p = 1, 2$, all other parameters remain the same as in the simulation setup “a.” Using this setup, we recalculate the simulation presented in the second “row” of Table 1, for estimation of the functions $X_i^{(p)}$, $p = 1, 2$, we used the Nadaraya–Watson estimation with Epanechnikov kernel and bandwidth $b = 0.05$. We run the simulations with various bandwidths, the choice of the bandwidth does not have a strong influence on results except by oversmoothing (large bandwidths). The results are printed in Table 2. As we can see, the difference of the simulation results using estimated functions is not significant in comparison to the results printed in the second line of Table 1—directly observed functional values.

The last limitation of this simulation study is the choice of a particular alternative. A more general setup of this simulation study might be based on the following model: $X_i^{(1)}(t) = \beta_{1i}^{(1)} \gamma_1^{(1)}(t) + \beta_{2i}^{(1)} \gamma_2^{(1)}(t)$, $X_i^{(2)}(t) = \beta_{1i}^{(2)} \gamma_1^{(2)}(t) + \beta_{2i}^{(2)} \gamma_2^{(2)}(t)$, where $\gamma_1^{(1)}$, $\gamma_1^{(2)}$, $\gamma_2^{(1)}$ and g are mutually orthogonal functions on $L^2[0, 1]$ and $\gamma_2^{(2)} = (1 + v^2)^{-1/2} \{\gamma_2^{(1)} + vg\}$. Basically we create the alternative by the contamination of one of the “eigenfunctions” (in our case the second one) in the direction g and ensure $\|\gamma_2^{(2)}\| = 1$. The amount of the contamination is controlled by the parameter v . Note that the exact squared integral difference $\|\gamma_2^{(1)} - \gamma_2^{(2)}\|^2$ does not depend on function g . Thus, in the “functional sense” particular “direction of the alternative hypothesis” represented by the function g has no impact on the power of the test. However, since we are using a nonparametric estimation technique, we might expect that rough (highly fluctuating) functions g will yield higher error of estimation and, hence, decrease the precision (and power) of the test. Finally, a higher number of factor functions (L) in simulation may cause less precise approximation of critical values and more bootstrap replications and

larger sample-size may be needed. This can also be expected from Theorem 2 in Section 2.2—the variance of the estimated eigenfunctions depends on all eigenfunctions corresponding to nonzero eigenvalues.

4. Implied volatility analysis. In this section we present an application of the method discussed in previous sections to the implied volatilities of European options on the German stock index (ODAX). Implied volatilities are derived from the Black–Scholes (BS) pricing formula for European options; see Black and Scholes (1973). European call and put options are derivatives written on an underlying asset with price process S_i , which yield the pay-off $\max(S_I - K, 0)$ and $\max(K - S_I, 0)$, respectively. Here i denotes the current day, I the expiration day and K the strike price. Time to maturity is defined as $\tau = I - i$. The BS pricing formula for a Call option is

$$(14) \quad C_i(S_i, K, \tau, r, \sigma) = S_i \Phi(d_1) - K e^{-r\tau} \Phi(d_2),$$

where $d_1 = \frac{\ln(S_i/K) + (r + \sigma^2/2)\tau}{\sigma\sqrt{\tau}}$, $d_2 = d_1 - \sigma\sqrt{\tau}$, r is the risk-free interest rate, σ is the (unknown and constant) volatility parameter, and Φ denotes the c.d.f. of a standard normal distributed random variable. In (14) we assume the zero-dividend case. The Put option price P_i can be obtained from the put–call parity $P_i = C_i - S_i + e^{-r\tau}K$.

The implied volatility $\tilde{\sigma}$ is defined as the volatility σ , for which the BS price C_i in (14) equals the price \tilde{C}_i observed on the market. For a single asset, we obtain at each time point (day i) and for each maturity τ a IV function $\tilde{\sigma}_i^\tau(K)$. Practitioners often rescale the strike dimension by plotting this surface in terms of (futures) moneyness $\kappa = K/F_i(\tau)$, where $F_i(\tau) = S_i e^{r\tau}$.

Clearly, for given parameters S_i, r, K, τ the mapping from prices to IVs is a one-to-one mapping. The IV is often used for quoting the European options in financial practice, since it reflects the “uncertainty” of the financial market better than the option prices. It is also known that if the stock price drops, the IV raises (so-called leverage effect), motivates hedging strategies based on IVs. Consequently, for the purpose of this application, we will regard the BS–IV as an individual financial variable. The practical relevance of such an approach is justified by the volatility based financial products such as VDAX, which are commonly traded on the option markets.

The goal of this analysis is to study the dynamics of the IV functions for different maturities. More specifically, our aim is to construct low dimensional factor models based on the truncated Karhunen–Loève expansions (1) for the log-returns of the IV functions of options with different maturities and compare these factor models using the methodology presented in the previous sections. Analysis of IVs based on a low-dimensional factor model gives directly a descriptive insight into the structure of distribution

of the log-IV-returns—structure of the factors and empirical distribution of the factor loadings may be a good starting point for further pricing models. In practice, such a factor model can also be used in Monte Carlo based pricing methods and for risk-management (hedging) purposes. For comprehensive monographs on IV and IV-factor models, see [Hafner \(2004\)](#) or [Fengler \(2005b\)](#).

The idea of constructing and analyzing the factor models for log-IV-returns for different maturities was originally proposed by [Fengler, Härdle and Villa \(2003\)](#), who studied the dynamics of the IV via PCA on discretized IV functions for different maturity groups and tested the Common Principal Components (CPC) hypotheses (equality of eigenvectors and eigenspaces for different groups). [Fengler, Härdle and Villa \(2003\)](#) proposed a PCA-based factor model for log-IV-returns on (short) maturities 1, 2 and 3 months and grid of moneyness $[0.85, 0.9, 0.95, 1, 1.05, 1.1]$. They showed that the factor functions do not significantly differ and only the factor loadings differ across maturity groups. Their method relies on the CPC methodology introduced by [Flury \(1988\)](#) which is based on maximum likelihood estimation under the assumption of multivariate normality. The log-IV-returns are extracted by the two-dimensional Nadaraya–Watson estimate.

The main aim of this application is to reconsider their results in a functional sense. Doing so, we overcome two basic weaknesses of their approach. First, the factor model proposed by [Fengler, Härdle and Villa \(2003\)](#) is performed only on a sparse design of moneyness. However, in practice (e.g., in Monte Carlo pricing methods), evaluation of the model on a fine grid is needed. Using the functional PCA approach, we may overcome this difficulty and evaluate the factor model on an arbitrary fine grid. The second difficulty of the procedure proposed by [Fengler, Härdle and Villa \(2003\)](#) stems from the data design—on the exchange we cannot observe options with desired maturity on each day and we need to estimate them from the IV-functions with maturities observed on the particular day. Consequently, the two-dimensional Nadaraya–Watson estimator proposed by [Fengler, Härdle and Villa \(2003\)](#) results essentially in the (weighted) average of the IVs (with closest maturities) observed on a particular day, which may affect the test of the common eigenfunction hypothesis. We use the linear interpolation scheme in the *total variance* $\sigma_{\text{TOT},i}^2(\kappa, \tau) \stackrel{\text{def}}{=} (\sigma_i^\tau(\kappa))^2 \tau$, in order to recover the IV functions with fixed maturity (on day i). This interpolation scheme is based on the arbitrage arguments originally proposed by [Kahalé \(2004\)](#) for zero-dividend and zero-interest rate case and generalized for deterministic interest rate by [Fengler \(2005a\)](#). More precisely, having IVs with maturities observed on a particular day i : $\tilde{\sigma}_i^{\tau_{j_i}}(\kappa)$, $j_i = 1, \dots, p_{\tau_i}$, we calculate the corresponding total variance $\tilde{\sigma}_{\text{TOT},i}(\kappa, \tau_{j_i})$. From these total variances

we linearly interpolate the total variance with the desired maturity from the nearest maturities observed on day i . The total variance can be easily transformed to corresponding IV $\tilde{\sigma}_i^\tau(\kappa)$. As the last step, we calculate the log-returns $\Delta \log \tilde{\sigma}_i^\tau(\kappa) \stackrel{\text{def}}{=} \log \tilde{\sigma}_{i+1}^\tau(\kappa) - \log \tilde{\sigma}_i^\tau(\kappa)$. The log-IV-returns are observed for each maturity τ on a discrete grid κ_{ik}^τ . We assume that observed log-IV-return $\Delta \log \tilde{\sigma}_i^\tau(\kappa_{ik}^\tau)$ consists of true log-return of the IV function denoted by $\Delta \log \sigma_i^\tau(\kappa_{ik}^\tau)$ and possibly of some additional error ε_{ik}^τ . By setting $Y_{ik}^\tau := \Delta \log \tilde{\sigma}_i^\tau(\kappa_{ik}^\tau)$, $X_i^\tau(\kappa) := \Delta \log \sigma_i^\tau(\kappa)$, we obtain an analogue of the model (4) with the argument κ :

$$(15) \quad Y_{ik}^\tau = X_i^\tau(\kappa_{ik}) + \varepsilon_{ik}^\tau, \quad i = 1, \dots, n_\tau.$$

In order to simplify the notation and make the connection with the theoretical part clear, we will use the notation of (15).

For our analysis we use a recent data set containing daily data from January 2004 to June 2004 from the German–Swiss exchange (EUREX). Violations of the arbitrage-free assumptions (“obvious” errors in data) were corrected using the procedure proposed by [Fengler \(2005a\)](#). Similarly to [Fengler, Härdle and Villa \(2003\)](#), we excluded options with maturity smaller than 10 days, since these option-prices are known to be very noisy, partially because of a special and arbitrary setup in the pricing systems of the dealers. Using the interpolation scheme described above, we calculate the log-IV-returns for two maturity groups: “1M” group with maturity $\tau = 0.12$ (measured in years) and “3M” group with maturity $\tau = 0.36$. The observed log-IV-returns are denoted by Y_{ik}^{1M} , $k = 1, \dots, K_i^{1M}$, Y_{ik}^{3M} , $k = 1, \dots, K_i^{3M}$. Since we ensured that for no i , the interpolation procedure uses data with the same maturity for both groups, this procedure has no impact on the independence of both samples.

The underlying models based on the truncated version of (3) are as follows:

$$(16) \quad X_i^{1M}(\kappa) = \bar{X}^{1M}(\kappa) + \sum_{r=1}^{L_{1M}} \hat{\beta}_{ri}^{1M} \widehat{\gamma}_r^{1M}(\kappa), \quad i = 1, \dots, n_{1M},$$

$$(17) \quad X_i^{3M}(\kappa) = \bar{X}^{3M}(\kappa) + \sum_{r=1}^{L_{3M}} \hat{\beta}_{ri}^{3M} \widehat{\gamma}_r^{3M}(\kappa), \quad i = 1, \dots, n_{3M}.$$

Models (16) and (17) can serve, for example, in a Monte Carlo pricing tool in the risk management for pricing exotic options where the whole path of implied volatilities is needed to determine the price. Estimating the factor functions in (16) and (17) by eigenfunctions displayed in [Figure 1](#), we only need to fit the (estimated) factor loadings $\hat{\beta}_{ji}^{1M}$ and $\hat{\beta}_{ji}^{3M}$. The pillar of the model is the dimension reduction. Keeping the factor function fixed for a certain time period, we need to analyze (two) multivariate random processes

of the factor loadings. For the purposes of this paper we will focus on the comparison of factors from models (16) and (17) and the technical details of the factor loading analysis will not be discussed here, since in this respect we refer to [Fengler, Härdle and Villa \(2003\)](#), who proposed to fit the factor loadings by centered normal distributions with diagonal variance matrix containing the corresponding eigenvalues. For a deeper discussion of the fitting of factor loadings using a more sophisticated approach, basically based on (possibly multivariate) GARCH models; see [Fengler \(2005b\)](#).

From our data set we obtained 88 functional observations for the 1M group (n_{1M}) and 125 observations for the 3M group (n_{3M}). We will estimate the model on the interval for futures moneyness $\kappa \in [0.8, 1.1]$. In comparison to [Fengler, Härdle and Villa \(2003\)](#), we may estimate models (16) and (17) on an arbitrary fine grid (we used an equidistant grid of 500 points on the interval $[0.8, 1.1]$). For illustration, the Nadaraya–Watson (NW) estimator of resulting log-returns is plotted in [Figure 2](#). The smoothing parameters have been chosen in accordance with the requirements in [Section 2.2](#). As argued in [Section 2.2](#), we should use small smoothing parameters in order to avoid a possible bias in the estimated eigenfunctions. Thus, we use for each i essentially the smallest bandwidth b_i that guarantees that estimator \hat{X}_i is defined on the entire support $[0.8, 1.1]$.

Using the procedures described in [Section 2.1](#), we first estimate the eigenfunctions of both maturity groups. The estimated eigenfunctions are plotted in [Figure 1](#). The structure of the eigenfunctions is in accordance with other empirical studies on IV-surfaces. For a deeper discussion and economical interpretation, see, for example, [Fengler, Härdle and Mammen \(2007\)](#) or [Fengler, Härdle and Villa \(2003\)](#).

Clearly, the ratio of the variance explained by the k th factor function is given by the quantity $\hat{\nu}_k^{1M} = \hat{\lambda}_k^{1M} / \sum_{j=1}^{r_{1M}} \hat{\lambda}_j^{1M}$ for the 1M group and, correspondingly, by $\hat{\nu}_k^{3M}$ for the 3M group. In [Table 3](#) we list the contributions of the factor functions. Looking at [Table 3](#), we can see that 4th factor functions explain less than 1% of the variation. This number was the “threshold” for the choice of L_{1M} and L_{2M} .

We can observe (see [Figure 1](#)) that the factor functions for both groups are similar. Thus, in the next step we use the bootstrap test for testing the

TABLE 3
Variance explained by the eigenfunctions

	Var. explained 1M	Var. explained 3M
$\hat{\nu}_1^T$	89.9%	93.0%
$\hat{\nu}_2^T$	7.7%	4.2%
$\hat{\nu}_3^T$	1.7%	1.0%
$\hat{\nu}_4^T$	0.6%	0.4%




FIG. 2. *Nadaraya–Watson estimate of the log-IV-returns for maturity 1M (left figure) and 3M (right figure). The bold line is the sample mean of the corresponding group.*

equality of the factor functions. We use 2000 bootstrap replications. The test of equality of the eigenfunctions was rejected for the first eigenfunction for the analyzed time period (January 2004–June 2004) at a significance level $\alpha = 0.05$ (P-value 0.01). We may conclude that the (first) factor functions are not identical in the factor model for both maturity groups. However, from a practical point of view, we are more interested in checking the appropriateness of the entire models for a fixed number of factors: $L = 2$ or $L = 3$ in (16) and (17). This requirement translates into the testing of the equality of eigenspaces. Thus, in the next step we use the same setup (2000 bootstrap replications) to test the hypotheses that the first two and first three eigenfunctions span the same eigenspaces \mathcal{E}_L^{1M} and \mathcal{E}_L^{3M} . None of the hypotheses for $L = 2$ and $L = 3$ is rejected at significance level $\alpha = 0.05$ (P-value is 0.61 for $L = 2$ and 0.09 for $L = 3$). Summarizing, even in the functional sense we have no significant reason to reject the hypothesis of common eigenspaces for these two maturity groups. Using this hypothesis, the factors governing the movement of the returns of IV surface are invariant to time to maturity, only their relative importance can vary. This leads to the common factor model: $X_i^\tau(\kappa) = \bar{X}^\tau(\kappa) + \sum_{r=1}^{L_\tau} \hat{\beta}_{ri}^\tau \widehat{\gamma}_r(\kappa)$, $i = 1, \dots, n_\tau$, $\tau = 1M, 3M$, where $\gamma_r := \gamma_r^{1M} = \gamma_r^{3M}$. Beside contributing to the understanding of the structure of the IV function dynamics, the common factor model helps us to reduce the number of functional factors by half compared to models (16) and (17). Furthermore, from the technical point of view, we also obtain an additional dimension reduction and higher estimation precision, since under this hypothesis we may estimate the eigenfunctions from the (individually centered) pooled sample $X_i(\kappa)^{1M}$, $i = 1, \dots, n_{1M}$, $X_i^{3M}(\kappa)$, $i =$

$1, \dots, n_{3M}$. The main improvement compared to the multivariate study by [Fengler, Härdle and Villa \(2003\)](#) is that our test is performed in the functional sense – it does not depend on particular discretization and our factor model can be evaluated on an arbitrary fine grid.

APPENDIX: MATHEMATICAL PROOFS

In the following, $\|v\| = (\int_0^1 v(t)^2 dt)^{1/2}$ will denote the L^2 -norm for any square integrable function v . At the same time, $\|a\| = (\frac{1}{k} \sum_{i=1}^k a_i^2)^{1/2}$ will indicate the Euclidean norm, whenever $a \in \mathbb{R}^k$ is a k -vector for some $k \in \mathbb{N}$.

In the proof of [Theorem 1](#), E_ε and Var_ε denote expectation and variance with respect to ε only (i.e., conditional on t_{ij} and X_i).

PROOF OF THEOREM 1. Recall the definition of the $\chi_i(t)$ and note that $\chi_i(t) = \chi_i^X(t) + \chi_i^\varepsilon(t)$, where

$$\chi_i^\varepsilon(t) = \sum_{j=1}^{T_i} \varepsilon_{i(j)} I\left(t \in \left[\frac{t_{i(j-1)} + t_{i(j)}}{2}, \frac{t_{i(j)} + t_{i(j+1)}}{2}\right)\right),$$

as well as

$$\chi_i^X(t) = \sum_{j=1}^{T_i} X_i(t_{i(j)}) I\left(t \in \left[\frac{t_{i(j-1)} + t_{i(j)}}{2}, \frac{t_{i(j)} + t_{i(j+1)}}{2}\right)\right)$$

for $t \in [0, 1]$, $t_{i(0)} = -t_{i(1)}$ and $t_{i(T_i+1)} = 2 - t_{i(T_i)}$. Similarly, $\chi_i^*(t) = \chi_i^{X^*}(t) + \chi_i^{\varepsilon^*}(t)$.

By [Assumption 2](#), $E(|t_{i(j)} - t_{i(j-1)}|^s) = \mathcal{O}(T^{-s})$ for $s = 1, \dots, 4$, and the convergence is uniform in $j < n$. Our assumptions on the structure of X_i together with some straightforward Taylor expansions then lead to

$$\langle \chi_i, \chi_j \rangle = \langle X_i, X_j \rangle + \mathcal{O}_p(1/T)$$

and

$$\langle \chi_i, \chi_i^* \rangle = \|X_i\|^2 + \mathcal{O}_p(1/T).$$

Moreover,

$$\begin{aligned} E_\varepsilon(\langle \chi_i^\varepsilon, \chi_j^X \rangle) &= 0, & E_\varepsilon(\|\chi_i^\varepsilon\|^2) &= \sigma_i^2, \\ E_\varepsilon(\langle \chi_i^\varepsilon, \chi_i^{\varepsilon^*} \rangle) &= 0, & E_\varepsilon(\langle \chi_i^\varepsilon, \chi_i^{\varepsilon^*} \rangle^2) &= \mathcal{O}_p(1/T), \\ E_\varepsilon(\langle \chi_i^\varepsilon, \chi_j^X \rangle^2) &= \mathcal{O}_p(1/T), & E_\varepsilon(\langle \chi_i^\varepsilon, \chi_j^X \rangle \langle \chi_k^\varepsilon, \chi_l^X \rangle) &= 0 \quad \text{for } i \neq k, \end{aligned}$$

$$E_\varepsilon(\langle \chi_i^\varepsilon, \chi_j^\varepsilon \rangle \langle \chi_i^\varepsilon, \chi_k^\varepsilon \rangle) = 0 \quad \text{for } j \neq k \text{ and } E_\varepsilon(\|\chi_i^\varepsilon\|^4) = \mathcal{O}_p(1)$$

hold (uniformly) for all $i, j = 1, \dots, n$.

Consequently, $E_\varepsilon(\|\bar{\chi}\|^2 - \|\bar{X}\|^2) = \mathcal{O}_p(T^{-1} + n^{-1})$.

When using these relations, it is easily seen that for all $i, j = 1, \dots, n$

$$(18) \quad \begin{aligned} \widehat{M}_{ij} - M_{ij} &= \mathcal{O}_p(T^{-1/2} + n^{-1}) \quad \text{and} \\ \text{tr}\{(\widehat{M} - M)^2\}^{1/2} &= \mathcal{O}_p(1 + nT^{-1/2}). \end{aligned}$$

Since the orthonormal eigenvectors p_q of M satisfy $\|p_q\| = 1$, we furthermore obtain for any $i = 1, \dots, n$ and all $q = 1, 2, \dots$

$$(19) \quad \sum_{j=1}^n p_{jq} \left\{ \widehat{M}_{ij} - M_{ij} - \int_0^1 \chi_i^\varepsilon(t) \chi_j^X(t) dt \right\} = \mathcal{O}_p(T^{-1/2} + n^{-1/2}),$$

as well as

$$(20) \quad \sum_{j=1}^n p_{jq} \int_0^1 \chi_i^\varepsilon(t) \chi_j^X(t) dt = \mathcal{O}_p\left(\frac{n^{1/2}}{T^{1/2}}\right)$$

and

$$(21) \quad \sum_{i=1}^n a_i \sum_{j=1}^n p_{jq} \int_0^1 \chi_i^\varepsilon(t) \chi_j^X(t) dt = \mathcal{O}_p\left(\frac{n^{1/2}}{T^{1/2}}\right)$$

for any further vector a with $\|a\| = 1$.

Recall that the j th largest eigenvalue l_j satisfies $n\hat{\lambda}_j = l_j$. Since by assumption $\inf_{s \neq r} |\lambda_r - \lambda_s| > 0$, the results of [Dauxois, Pousse and Romain \(1982\)](#) imply that $\hat{\lambda}_r$ converges to λ_r as $n \rightarrow \infty$, and $\sup_{s \neq r} \frac{1}{|\hat{\lambda}_r - \hat{\lambda}_s|} = \mathcal{O}_p(1)$, which leads to $\sup_{s \neq r} \frac{1}{|l_r - l_s|} = \mathcal{O}_p(1/n)$. Assertion (a) of Lemma A of [Kneip and Utikal \(2001\)](#) together with (18)–(21) then implies that

$$(22) \quad \begin{aligned} \left| \hat{\lambda}_r - \frac{\hat{l}_r}{n} \right| &= n^{-1} |l_r - \hat{l}_r| = n^{-1} |p_r^\top (\widehat{M} - M) p_r| + \mathcal{O}_p(T^{-1} + n^{-1}) \\ &= \mathcal{O}_p\{(nT)^{-1/2} + T^{-1} + n^{-1}\}. \end{aligned}$$

When analyzing the difference between the estimated and true eigenvectors \hat{p}_r and p_r , assertion (b) of Lemma A of [Kneip and Utikal \(2001\)](#) together with (18) lead to

$$(23) \quad \hat{p}_r - p_r = -\mathcal{S}_r(\widehat{M} - M)p_r + \mathcal{R}_r, \quad \text{with } \|\mathcal{R}_r\| = \mathcal{O}_p(T^{-1} + n^{-1})$$

and $\mathcal{S}_r = \sum_{s \neq r} \frac{1}{l_s - l_r} p_s p_s^\top$. Since $\sup_{\|a\|=1} a^\top \mathcal{S}_r a \leq \sup_{s \neq r} \frac{1}{|l_r - l_s|} = \mathcal{O}_p(1/n)$, we can conclude that

$$(24) \quad \|\hat{p}_r - p_r\| = \mathcal{O}_p(T^{-1/2} + n^{-1}),$$

and our assertion on the sequence $n^{-1} \sum_i (\hat{\beta}_{ri} - \hat{\beta}_{ri;T})^2$ is an immediate consequence.

Let us now consider assertion (ii). The well-known properties of local linear estimators imply that $|\mathbb{E}_\varepsilon\{\hat{X}_i(t) - X_i(t)\}| = \mathcal{O}_p(b^2)$, as well as $\text{Var}_\varepsilon\{\hat{X}_i(t)\} = \mathcal{O}_p\{Tb\}$, and the convergence is uniform for all i, n . Furthermore, due to the independence of the error term ε_{ij} , $\text{Cov}_\varepsilon\{\hat{X}_i(t), \hat{X}_j(t)\} = 0$ for $i \neq j$. Therefore,

$$\left| \hat{\gamma}_r(t) - \frac{1}{\sqrt{l_r}} \sum_{i=1}^n p_{ir} \hat{X}_i(t) \right| = \mathcal{O}_p\left(b^2 + \frac{1}{\sqrt{nTb}}\right).$$

On the other hand, (18)–(24) imply that with $\hat{X}(t) = (\hat{X}_1(t), \dots, \hat{X}_n(t))^\top$

$$\begin{aligned} & \left| \hat{\gamma}_{r;T}(t) - \frac{1}{\sqrt{l_r}} \sum_{i=1}^n p_{ir} \hat{X}_i(t) \right| \\ &= \left| \frac{1}{\sqrt{l_r}} \sum_{i=1}^n (\hat{p}_{ir} - p_{ir}) X_i(t) + \frac{1}{\sqrt{l_r}} \sum_{i=1}^n (\hat{p}_{ir} - p_{ir}) \{\hat{X}_i(t) - X_i(t)\} \right| \\ & \quad + \mathcal{O}_p(T^{-1} + n^{-1}) \\ &= \frac{\|\mathcal{S}_r X(t)\|}{\sqrt{l_r}} \left| p_r^\top (\hat{M} - M) \mathcal{S}_r \frac{X(t)}{\|\mathcal{S}_r X(t)\|} \right| \\ & \quad + \mathcal{O}_p(b^2 T^{-1/2} + T^{-1} b^{-1/2} + n^{-1}) \\ &= \mathcal{O}_p(n^{-1/2} T^{-1/2} + b^2 T^{-1/2} + T^{-1} b^{-1/2} + n^{-1}). \end{aligned}$$

This proves the theorem. \square

PROOF OF THEOREM 2. First consider assertion (i). By definition,

$$\bar{X}(t) - \mu(t) = n^{-1} \sum_{i=1}^n \{X_i(t) - \mu(t)\} = \sum_r \left(n^{-1} \sum_{i=1}^n \beta_{ri} \right) \gamma_r(t).$$

Recall that, by assumption, β_{ri} are independent, zero mean random variables with variance λ_r , and that the above series converges with probability 1. When defining the truncated series

$$V(q) = \sum_{r=1}^q \left(n^{-1} \sum_{i=1}^n \beta_{ri} \right) \gamma_r(t),$$

standard central limit theorems therefore imply that $\sqrt{n}V(q)$ is asymptotically $N(0, \sum_{r=1}^q \lambda_r \gamma_r(t)^2)$ distributed for any possible $q \in \mathbb{N}$.

The assertion of a $N(0, \sum_{r=1}^\infty \lambda_r \gamma_r(t)^2)$ limiting distribution now is a consequence of the fact that for all $\delta_1, \delta_2 > 0$ there exists a q_δ such that $P\{|\sqrt{n}V(q) - \sqrt{n}\sum_r(n^{-1}\sum_{i=1}^n\beta_{ri})\gamma_r(t)| > \delta_1\} < \delta_2$ for all $q \geq q_\delta$ and all n sufficiently large.

In order to prove assertions (i) and (ii), consider some fixed $r \in \{1, 2, \dots\}$ with $\lambda_{r-1} > \lambda_r > \lambda_{r+1}$. Note that Γ as well as $\hat{\Gamma}_n$ are nuclear, self-adjoint and non-negative linear operators with $\Gamma v = \int \sigma(t, s)v(s) ds$ and $\hat{\Gamma}_n v = \int \hat{\sigma}(t, s)v(s) ds$, $v \in L^2[0, 1]$. For $m \in \mathbb{N}$, let Π_m denote the orthogonal projector from $L^2[0, 1]$ into the m -dimensional linear space spanned by $\{\gamma_1, \dots, \gamma_m\}$, that is, $\Pi_m v = \sum_{j=1}^m \langle v, \gamma_j \rangle \gamma_j$, $v \in L^2[0, 1]$. Now consider the operator $\Pi_m \hat{\Gamma}_n \Pi_m$, as well as its eigenvalues and corresponding eigenfunctions denoted by $\hat{\lambda}_{1,m} \geq \hat{\lambda}_{2,m} \geq \dots$ and $\hat{\gamma}_{1,m}, \hat{\gamma}_{2,m}, \dots$, respectively. It follows from well-known results in the Hilbert space theory that $\Pi_m \hat{\Gamma}_n \Pi_m$ converges strongly to $\hat{\Gamma}_n$ as $m \rightarrow \infty$. Furthermore, we obtain (Rayleigh–Ritz theorem)

$$(25) \quad \lim_{m \rightarrow \infty} \hat{\lambda}_{r,m} = \lambda_r \quad \text{and} \quad \lim_{m \rightarrow \infty} \|\hat{\gamma}_r - \hat{\gamma}_{r,m}\| = 0 \quad \text{if } \hat{\lambda}_{r-1} > \hat{\lambda}_r > \hat{\lambda}_{r+1}.$$

Note that under the above condition $\hat{\gamma}_r$ is uniquely determined up to sign, and recall that we always assume that the right “versions” (with respect to sign) are used so that $\langle \hat{\gamma}_r, \hat{\gamma}_{r,m} \rangle \geq 0$. By definition, $\beta_{ji} = \int \gamma_j(t) \{X_i(t) - \mu(t)\} dt$, and therefore, $\int \gamma_j(t) \{X_i(t) - \bar{X}(t)\} dt = \beta_{ji} - \bar{\beta}_j$, as well as $X_i - \bar{X} = \sum_j (\beta_{ji} - \bar{\beta}_j) \gamma_j$, where $\bar{\beta}_j = \frac{1}{n} \sum_{i=1}^n \beta_{ji}$. When analyzing the structure of $\Pi_m \hat{\Gamma}_n \Pi_m$ more deeply, we can verify that $\Pi_m \hat{\Gamma}_n \Pi_m v = \int \hat{\sigma}_m(t, s)v(s) ds$, $v \in L^2[0, 1]$, with

$$\hat{\sigma}_m(t, s) = g_m(t)^\top \hat{\Sigma}_m g_m(s),$$

where $g_m(t) = (\gamma_1(t), \dots, \gamma_m(t))^\top$, and where $\hat{\Sigma}_m$ is the $m \times m$ matrix with elements $\{\frac{1}{n} \sum_{i=1}^n (\beta_{ji} - \bar{\beta}_j)(\beta_{ki} - \bar{\beta}_k)\}_{j,k=1, \dots, m}$. Let $\lambda_1(\hat{\Sigma}_m) \geq \lambda_2(\hat{\Sigma}_m) \geq \dots \geq \lambda_m(\hat{\Sigma}_m)$ and $\hat{\zeta}_{1,m}, \dots, \hat{\zeta}_{m,m}$ denote eigenvalues and corresponding eigenvectors of $\hat{\Sigma}_m$. Some straightforward algebra then shows that

$$(26) \quad \hat{\lambda}_{r,m} = \lambda_r(\hat{\Sigma}_m), \quad \hat{\gamma}_{r,m} = g_m(t)^\top \hat{\zeta}_{r,m}.$$

We will use Σ_m to represent the $m \times m$ diagonal matrix with diagonal entries $\lambda_1 \geq \dots \geq \lambda_m$. Obviously, the corresponding eigenvectors are given by the m -dimensional unit vectors denoted by $e_{1,m}, \dots, e_{m,m}$. Lemma A of [Kneip and Utikal \(2001\)](#) now implies that the differences between eigenvalues and eigenvectors of Σ_m and $\hat{\Sigma}_m$ can be bounded by

$$(27) \quad \hat{\lambda}_{r,m} - \lambda_r = \text{tr}\{e_{r,m} e_{r,m}^\top (\hat{\Sigma}_m - \Sigma_m)\} + \tilde{R}_{r,m},$$

$$\text{with } \tilde{R}_{r,m} \leq \frac{6 \sup_{\|a\|=1} a^\top (\hat{\Sigma}_m - \Sigma_m)^2 a}{\min_s |\lambda_s - \lambda_r|},$$

$$(28) \quad \hat{\zeta}_{r,m} - e_{r,m} = -S_{r,m} (\hat{\Sigma}_m - \Sigma_m) e_{r,m} + R_{r,m}^*,$$

$$\text{with } \|R_{r,m}^*\| \leq \frac{6 \sup_{\|a\|=1} a^\top (\hat{\Sigma}_m - \Sigma_m)^2 a}{\min_s |\lambda_s - \lambda_r|^2},$$

where $S_{r,m} = \sum_{s \neq r} \frac{1}{\lambda_s - \lambda_r} e_{s,m} e_{s,m}^\top$.

Assumption 1 implies $E(\bar{\beta}_r) = 0$, $\text{Var}(\bar{\beta}_r) = \frac{\lambda_r}{n}$, and with $\delta_{ii} = 1$, as well as $\delta_{ij} = 0$ for $i \neq j$, we obtain

$$\begin{aligned}
& E \left\{ \sup_{\|a\|=1} a^\top (\hat{\Sigma}_m - \Sigma_m)^2 a \right\} \\
& \leq E \{ \text{tr}[(\hat{\Sigma}_m - \Sigma_m)^2] \} \\
& = E \left\{ \sum_{j,k=1}^m \left[\frac{1}{n} \sum_{i=1}^n (\beta_{ji} - \bar{\beta}_j)(\beta_{ki} - \bar{\beta}_k) - \delta_{jk} \lambda_j \right]^2 \right\} \\
(29) \quad & \leq E \left\{ \sum_{j,k=1}^\infty \left[\frac{1}{n} \sum_{i=1}^n (\beta_{ji} - \bar{\beta}_j)(\beta_{ki} - \bar{\beta}_k) - \delta_{jk} \lambda_j \right]^2 \right\} \\
& = \frac{1}{n} \left(\sum_j \sum_k E \{ \beta_{ji}^2 \beta_{ki}^2 \} \right) + \mathcal{O}(n^{-1}) = \mathcal{O}(n^{-1}),
\end{aligned}$$

for all m . Since $\text{tr}\{e_{r,m} e_{r,m}^\top (\hat{\Sigma}_m - \Sigma_m)\} = \frac{1}{n} \sum_{i=1}^n (\beta_{ri} - \bar{\beta}_r)^2 - \lambda_r$, (25), (26), (27) and (29) together with standard central limit theorems imply that

$$\begin{aligned}
\sqrt{n}(\hat{\lambda}_r - \lambda_r) &= \frac{1}{\sqrt{n}} \sum_{i=1}^n (\beta_{ri} - \bar{\beta}_r)^2 - \lambda_r + \mathcal{O}_p(n^{-1/2}) \\
(30) \quad &= \frac{1}{\sqrt{n}} \sum_{i=1}^n [(\beta_{ri})^2 - E\{(\beta_{ri})^2\}] + \mathcal{O}_p(n^{-1/2}) \\
&\xrightarrow{\mathcal{L}} N(0, \Lambda_r).
\end{aligned}$$

It remains to prove assertion (iii). Relations (26) and (28) lead to

$$\begin{aligned}
\hat{\gamma}_{r,m}(t) - \gamma_r(t) &= g_m(t)^\top (\hat{\zeta}_{r,m} - e_{r,m}) \\
(31) \quad &= - \sum_{s \neq r}^m \left\{ \frac{1}{n(\lambda_s - \lambda_r)} \sum_{i=1}^n (\beta_{si} - \bar{\beta}_s)(\beta_{ri} - \bar{\beta}_r) \right\} \gamma_s(t) \\
&\quad + g_m(t)^\top R_{r,m}^*,
\end{aligned}$$

where due to (29) the function $g_m(t)^\top R_{r,m}^*$ satisfies

$$\begin{aligned}
E(\|g_m^\top R_{r,m}^*\|) &= E(\|R_{r,m}^*\|) \\
&\leq \frac{6}{n \min_s |\lambda_s - \lambda_r|^2} \left(\sum_j \sum_k E \{ \beta_{ji}^2 \beta_{ki}^2 \} \right) + \mathcal{O}(n^{-1}),
\end{aligned}$$

for all m . By Assumption 1, the series in (31) converge with probability 1 as $m \rightarrow \infty$.

Obviously, the event $\hat{\lambda}_{r-1} > \hat{\lambda}_r > \hat{\lambda}_{r+1}$ occurs with probability 1. Since m is arbitrary, we can therefore conclude from (25) and (31) that

$$\begin{aligned}
& \hat{\gamma}_r(t) - \gamma_r(t) \\
(32) \quad &= - \sum_{s \neq r} \left\{ \frac{1}{n(\lambda_s - \lambda_r)} \sum_{i=1}^n (\beta_{si} - \bar{\beta}_s)(\beta_{ri} - \bar{\beta}_r) \right\} \gamma_s(t) + R_r^*(t) \\
&= - \sum_{s \neq r} \left\{ \frac{1}{n(\lambda_s - \lambda_r)} \sum_{i=1}^n \beta_{si} \beta_{ri} \right\} \gamma_s(t) + R_r(t),
\end{aligned}$$

where $\|R_r^*\| = \mathcal{O}_p(n^{-1})$, as well as $\|R_r\| = \mathcal{O}_p(n^{-1})$. Moreover, $\sqrt{n} \times \sum_{s \neq r} \left\{ \frac{1}{n(\lambda_s - \lambda_r)} \sum_{i=1}^n \beta_{si} \beta_{ri} \right\} \gamma_s(t)$ is a zero mean random variable with variance $\sum_{q \neq r} \sum_{s \neq r} \frac{E[\beta_{ri}^2 \beta_{qi} \beta_{si}]}{(\lambda_q - \lambda_r)(\lambda_s - \lambda_r)} \gamma_q(t) \gamma_s(t) < \infty$. By Assumption 1, it follows from standard central limit arguments that for any $q \in \mathbb{N}$ the truncated series $\sqrt{n} W(q) \stackrel{\text{def}}{=} \sqrt{n} \sum_{s=1, s \neq r}^q \left[\frac{1}{n(\lambda_s - \lambda_r)} \sum_{i=1}^n \beta_{si} \beta_{ri} \right] \gamma_s(t)$ is asymptotically normal distributed. The asserted asymptotic normality of the complete series then follows from an argument similar to the one used in the proof of assertion (i). \square

PROOF OF THEOREM 3. The results of Theorem 2 imply that

$$\begin{aligned}
(33) \quad n\Delta_1 &= \int \left(\sum_r \left(\frac{1}{\sqrt{q_1 n_1}} \sum_{i=1}^{n_1} \beta_{ri}^{(1)} \gamma_r^{(1)}(t) \right. \right. \\
&\quad \left. \left. - \sum_r \frac{1}{\sqrt{q_2 n_2}} \sum_{i=1}^{n_2} \beta_{ri}^{(2)} \gamma_r^{(2)}(t) \right) \right)^2 dt.
\end{aligned}$$

Furthermore, independence of $X_i^{(1)}$ and $X_i^{(2)}$ together with (30) imply that

$$\begin{aligned}
(34) \quad \sqrt{n}[\hat{\lambda}_r^{(1)} - \lambda_r^{(1)} - \{\hat{\lambda}_r^{(2)} - \lambda_r^{(2)}\}] &\xrightarrow{\mathcal{L}} N\left(0, \frac{\Lambda_r^{(1)}}{q_1} + \frac{\Lambda_r^{(2)}}{q_2}\right) \quad \text{and} \\
\frac{n}{\Lambda_r^{(1)}/q_1 + \Lambda_r^{(2)}/q_2} \Delta_{3,r} &\xrightarrow{\mathcal{L}} \chi_1^2.
\end{aligned}$$

Furthermore, (32) leads to

$$\begin{aligned}
(35) \quad n\Delta_{2,r} &= \left\| \sum_{s \neq r} \left\{ \frac{1}{\sqrt{q_1 n_1} (\lambda_s^{(1)} - \lambda_r^{(1)})} \sum_{i=1}^{n_1} \beta_{si}^{(1)} \beta_{ri}^{(1)} \right\} \gamma_s^{(1)} \right. \\
&\quad \left. - \sum_{s \neq r} \left\{ \frac{1}{\sqrt{q_2 n_2} (\lambda_s^{(2)} - \lambda_r^{(2)})} \sum_{i=1}^{n_2} \beta_{si}^{(2)} \beta_{ri}^{(2)} \right\} \gamma_s^{(2)} \right\|^2 + \mathcal{O}_p(n^{-1/2})
\end{aligned}$$

and

$$\begin{aligned}
n\Delta_{4,L} &= n \iint \left[\sum_{r=1}^L \gamma_r^{(1)}(t) \{ \hat{\gamma}_r^{(1)}(u) - \gamma_r^{(1)}(u) \} \right. \\
&\quad + \gamma_r^{(1)}(u) \{ \hat{\gamma}_r^{(1)}(t) - \gamma_r^{(1)}(t) \} \\
&\quad - \sum_{r=1}^L \gamma_r^{(2)}(t) \{ \hat{\gamma}_r^{(2)}(u) - \gamma_r^{(2)}(u) \} \\
&\quad \left. + \gamma_r^{(2)}(u) \{ \hat{\gamma}_r^{(2)}(t) - \gamma_r^{(2)}(t) \} \right]^2 dt du + \mathcal{O}_p(n^{-1/2}) \\
(36) \quad &= \iint \left[\sum_{r=1}^L \sum_{s>L} \left\{ \frac{1}{\sqrt{q_1 n_1} (\lambda_s^{(1)} - \lambda_r^{(1)})} \sum_{i=1}^{n_1} \beta_{si}^{(1)} \beta_{ri}^{(1)} \right\} \right. \\
&\quad \times \{ \gamma_r^{(1)}(t) \gamma_s^{(1)}(u) + \gamma_r^{(1)}(u) \gamma_s^{(1)}(t) \} \\
&\quad - \sum_{r=1}^L \sum_{s>L} \left\{ \frac{1}{\sqrt{q_2 n_2} (\lambda_s^{(2)} - \lambda_r^{(2)})} \sum_{i=1}^{n_2} \beta_{si}^{(2)} \beta_{ri}^{(2)} \right\} \\
&\quad \left. \times \{ \gamma_r^{(2)}(t) \gamma_s^{(2)}(u) + \gamma_r^{(2)}(u) \gamma_s^{(2)}(t) \} \right]^2 dt du \\
&\quad + \mathcal{O}_p(n^{-1/2}).
\end{aligned}$$

In order to verify (36), note that $\sum_{r=1}^L \sum_{s=1, s \neq r}^L \frac{1}{(\lambda_s^{(p)} - \lambda_r^{(p)})} a_r a_s = 0$ for $p = 1, 2$ and all possible sequences a_1, \dots, a_L . It is clear from our assumptions that all sums involved converge with probability 1. Recall that $E(\beta_{ri}^{(p)} \beta_{si}^{(p)}) = 0$, $p = 1, 2$ for $r \neq s$.

It follows that $\tilde{X}_r^{(p)} := \frac{1}{\sqrt{q_p n_p}} \sum_{s \neq r} \sum_{i=1}^{n_p} \frac{\beta_{si}^{(p)} \beta_{ri}^{(p)}}{\lambda_s^{(p)} - \lambda_r^{(p)}} \gamma_s^{(p)}$, $p = 1, 2$, is a continuous, zero mean random function on $L^2[0, 1]$, and, by assumption, $E(\|\tilde{X}_r^{(p)}\|^2) < \infty$. By Hilbert space central limit theorems [see, e.g., [Araujo and Giné \(1980\)](#)], $\tilde{X}_r^{(p)}$ thus converges in distribution to a Gaussian random function $\xi_r^{(p)}$ as $n \rightarrow \infty$. Obviously, $\xi_r^{(1)}$ is independent of $\xi_r^{(2)}$. We can conclude that $n\Delta_{4,L}$ possesses a continuous limit distribution $F_{4,L}$ defined by the distribution of $\iint [\sum_{r=1}^L \{ \xi_r^{(1)}(t) \gamma_r^{(1)}(u) + \xi_r^{(1)}(u) \gamma_r^{(1)}(t) \} - \sum_{r=1}^L \{ \xi_r^{(2)}(t) \gamma_r^{(2)}(u) + \xi_r^{(2)}(u) \gamma_r^{(2)}(t) \}]^2 dt du$. Similar arguments show the existence of continuous limit distributions F_1 and $F_{2,r}$ of $n\Delta_1$ and $n\Delta_{2,r}$.

For given $q \in \mathbb{N}$, define vectors $b_{i1}^{(p)} = (\beta_{1i}^{(p)}, \dots, \beta_{qi}^{(p)})^\top \in \mathbb{R}^q$, $b_{i2}^{(p)} = (\beta_{1i}^{(p)} \beta_{ri}^{(p)}, \dots, \beta_{r-1,i}^{(p)} \beta_{ri}^{(p)}, \beta_{r+1,i}^{(p)} \beta_{ri}^{(p)}, \dots, \beta_{qi}^{(p)} \beta_{ri}^{(p)})^\top \in \mathbb{R}^{q-1}$ and $b_{i3} = (\beta_{1i}^{(p)} \beta_{2i}^{(p)},$

$\dots, \beta_{qi}^{(p)} \beta_{Li}^{(p)})^\top \in \mathbb{R}^{(q-1)L}$. When the infinite sums over r in (33), respectively $s \neq r$ in (35) and (36), are restricted to $q \in \mathbb{N}$ components (i.e., \sum_r and $\sum_{s>L}$ are replaced by $\sum_{r \leq q}$ and $\sum_{L < s \leq q}$), then the above relations can generally be presented as limits $n\Delta = \lim_{q \rightarrow \infty} n\Delta(q)$ of quadratic forms

$$(37) \quad \begin{aligned} n\Delta_1(q) &= \begin{pmatrix} \frac{1}{\sqrt{n_1}} \sum_{i=1}^{n_1} b_{i1}^{(1)} \\ \frac{1}{\sqrt{n_2}} \sum_{i=1}^{n_2} b_{i1}^{(2)} \end{pmatrix}^\top Q_1^q \begin{pmatrix} \frac{1}{\sqrt{n_1}} \sum_{i=1}^{n_1} b_{i1}^{(1)} \\ \frac{1}{\sqrt{n_2}} \sum_{i=1}^{n_2} b_{i1}^{(2)} \end{pmatrix}, \\ n\Delta_{2,r}(q) &= \begin{pmatrix} \frac{1}{\sqrt{n_1}} \sum_{i=1}^{n_1} b_{i2}^{(1)} \\ \frac{1}{\sqrt{n_2}} \sum_{i=1}^{n_2} b_{i2}^{(2)} \end{pmatrix}^\top Q_2^q \begin{pmatrix} \frac{1}{\sqrt{n_1}} \sum_{i=1}^{n_1} b_{i2}^{(1)} \\ \frac{1}{\sqrt{n_2}} \sum_{i=1}^{n_2} b_{i2}^{(2)} \end{pmatrix}, \\ n\Delta_{4,L}(q) &= \begin{pmatrix} \frac{1}{\sqrt{n_1}} \sum_{i=1}^{n_1} b_{i3}^{(1)} \\ \frac{1}{\sqrt{n_2}} \sum_{i=1}^{n_2} b_{i3}^{(2)} \end{pmatrix}^\top Q_3^q \begin{pmatrix} \frac{1}{\sqrt{n_1}} \sum_{i=1}^{n_1} b_{i3}^{(1)} \\ \frac{1}{\sqrt{n_2}} \sum_{i=1}^{n_2} b_{i3}^{(2)} \end{pmatrix}, \end{aligned}$$

where the elements of the $2q \times 2q$, $2(q-1) \times 2(q-1)$ and $2L(q-1) \times 2L(q-1)$ matrices Q_1^q , Q_2^q and Q_3^q can be computed from the respective (q -element) version of (33)–(36). Assumption 1 implies that all series converge with probability 1 as $q \rightarrow \infty$, and by (33)–(36), it is easily seen that for all $\epsilon, \delta > 0$ there exist some $q(\epsilon, \delta), n(\epsilon, \delta) \in \mathbb{N}$ such that

$$(38) \quad \begin{aligned} P(|n\Delta_1 - n\Delta_1(q)| > \epsilon) &< \delta, & P(|n\Delta_{2,r} - n\Delta_{2,r}(q)| > \epsilon) &< \delta, \\ P(|n\Delta_{4,L} - n\Delta_{4,L}(q)| > \epsilon) &< \delta \end{aligned}$$

hold for all $q \geq q(\epsilon, \delta)$ and all $n \geq n(\epsilon, \delta)$. For any given q , we have $E(b_{i1}) = E(b_{i2}) = E(b_{i3}) = 0$, and it follows from Assumption 1 that the respective covariance structures can be represented by finite covariance matrices $\Omega_{1,q}$, $\Omega_{2,q}$ and $\Omega_{3,q}$. It therefore follows from our assumptions together with standard multivariate central limit theorems that the vectors $\{\frac{1}{\sqrt{n_1}} \sum_{i=1}^{n_1} (b_{ik}^{(1)})^\top, \frac{1}{\sqrt{n_2}} \sum_{i=1}^{n_2} (b_{ik}^{(2)})^\top\}^\top$, $k = 1, 2, 3$, are asymptotically normal with zero means and covariance matrices $\Omega_{1,q}$, $\Omega_{2,q}$ and $\Omega_{3,q}$. One can thus conclude that, as $n \rightarrow \infty$,

$$(39) \quad n\Delta_1(q) \xrightarrow{\mathcal{L}} F_{1,q}, \quad n\Delta_{2,r}(q) \xrightarrow{\mathcal{L}} F_{2,r,q}, \quad n\Delta_{4,L}(q) \xrightarrow{\mathcal{L}} F_{4,L,q},$$

where $F_{1,q}, F_{2,r,q}, F_{4,L,q}$ denote the continuous distributions of the quadratic forms $z_1^\top Q_1^q z_1, z_2^\top Q_2^q z_2, z_3^\top Q_3^q z_3$ with $z_1 \sim N(0, \Omega_{1,q}), z_2 \sim N(0, \Omega_{2,q}), z_3 \sim$

$N(0, \Omega_{3,q})$. Since ϵ, δ are arbitrary, (38) implies

$$(40) \quad \lim_{q \rightarrow \infty} F_{1,q} = F_1, \quad \lim_{q \rightarrow \infty} F_{2,r,q} = F_{2,r}, \quad \lim_{q \rightarrow \infty} F_{4,L,q} = F_{4,L}.$$

We now have to consider the asymptotic properties of bootstrapped eigenvalues and eigenfunctions. Let $\bar{X}^{(p)*} = \frac{1}{n_p} \sum_{i=1}^{n_p} X_i^{(p)*}$, $\beta_{ri}^{(p)*} = \int \gamma_r^{(p)}(t) \{X_i^{(p)*}(t) - \mu(t)\}$, $\bar{\beta}_r^{(p)*} = \frac{1}{n_p} \sum_{i=1}^{n_p} \beta_{ri}^{(p)*}$, and note that $\int \gamma_r^{(p)}(t) \{X_i^{(p)*}(t) - \bar{X}^{(p)*}(t)\} = \beta_{ri}^{(p)*} - \bar{\beta}_r^{(p)*}$. When considering unconditional expectations, our assumptions imply that for $p = 1, 2$

$$(41) \quad \begin{aligned} \mathbb{E}[\beta_{ri}^{(p)*}] &= 0, & \mathbb{E}[(\beta_{ri}^{(p)*})^2] &= \lambda_r^{(p)}, \\ \mathbb{E}[(\bar{\beta}_r^{(p)*})^2] &= \frac{\lambda_r^{(p)}}{n_p}, & \mathbb{E}\{[(\beta_{ri}^{(p)*})^2 - \lambda_r^{(p)}]^2\} &= \Lambda_r^{(p)}, \\ \mathbb{E}\left\{ \sum_{l,k=1}^{\infty} \left[\frac{1}{n_p} \sum_{i=1}^{n_p} (\beta_{li}^{(p)*} - \bar{\beta}_l^{(p)*})(\beta_{ki}^{(p)*} - \bar{\beta}_k^{(p)*}) - \delta_{lk} \lambda_l^{(p)} \right]^2 \right\} \\ &= \frac{1}{n_p} \left(\sum_l \Lambda_l^{(p)} + \sum_{l \neq k} \lambda_l^{(p)} \lambda_k^{(p)} \right) + \mathcal{O}(n_p^{-1}). \end{aligned}$$

One can infer from (41) that the arguments used to prove Theorem 1 can be generalized to approximate the difference between the bootstrap eigenvalues and eigenfunctions $\hat{\lambda}_r^{(p)*}$, $\hat{\gamma}_r^{(p)*}$ and the true eigenvalues $\lambda_r^{(p)}$, $\gamma_r^{(p)}$. All infinite sums involved converge with probability 1. Relation (30) then generalizes to

$$(42) \quad \begin{aligned} &\sqrt{n_p}(\hat{\lambda}_r^{(p)*} - \hat{\lambda}_r^{(p)}) \\ &= \sqrt{n_p}(\hat{\lambda}_r^{(p)*} - \lambda_r^{(p)}) - \sqrt{n_p}(\hat{\lambda}_r^{(p)} - \lambda_r^{(p)}) \\ &= \frac{1}{\sqrt{n_p}} \sum_{i=1}^{n_p} (\beta_{ri}^{(p)*} - \bar{\beta}_r^{(p)*})^2 \\ &\quad - \frac{1}{\sqrt{n_p}} \sum_{i=1}^{n_p} (\beta_{ri}^{(p)} - \bar{\beta}_r^{(p)})^2 + \mathcal{O}_p(n_p^{-1/2}) \\ &= \frac{1}{\sqrt{n_p}} \sum_{i=1}^{n_p} \left\{ (\beta_{ri}^{(p)*})^2 - \frac{1}{n_p} \sum_{k=1}^{n_p} (\beta_{rk}^{(p)})^2 \right\} + \mathcal{O}_p(n_p^{-1/2}). \end{aligned}$$

Similarly, (32) becomes

$$(43) \quad \begin{aligned} &\hat{\gamma}_r^{(p)*} - \hat{\gamma}_r^{(p)} \\ &= \hat{\gamma}_r^{(p)*} - \gamma_r^{(p)} - (\hat{\gamma}_r^{(p)} - \gamma_r^{(p)}) \end{aligned}$$

$$\begin{aligned}
&= - \sum_{s \neq r} \left\{ \frac{1}{\lambda_s^{(p)} - \lambda_r^{(p)}} \frac{1}{n_p} \sum_{i=1}^{n_p} (\beta_{si}^{(p)*} - \bar{\beta}_s^{(p)*})(\beta_{ri}^{(p)*} - \bar{\beta}_r^{(p)*}) \right. \\
&\quad \left. - \frac{1}{\lambda_s^{(p)} - \lambda_r^{(p)}} \frac{1}{n_p} \sum_{i=1}^{n_p} (\beta_{si}^{(p)} - \bar{\beta}_s^{(p)})(\beta_{ri}^{(p)} - \bar{\beta}_r^{(p)}) \right\} \gamma_s^{(p)}(t) \\
&\quad + R_r^{(p)*}(t) \\
&= - \sum_{s \neq r} \left\{ \frac{1}{\lambda_s^{(p)} - \lambda_r^{(p)}} \frac{1}{n_p} \sum_{i=1}^{n_p} \left(\beta_{si}^{(p)*} \beta_{ri}^{(p)*} - \frac{1}{n_p} \sum_{k=1}^{n_p} \beta_{sk}^{(p)} \beta_{rk}^{(p)} \right) \right\} \gamma_s^{(p)}(t) \\
&\quad + \tilde{R}_r^{(p)*}(t),
\end{aligned}$$

where due to (28), (29) and (41), the remainder term satisfies $\|R_r^{(p)*}\| = \mathcal{O}_p(n_p^{-1})$.

We are now ready to analyze the bootstrap versions Δ^* of the different Δ . First consider $\Delta_{3,r}^*$ and note that $\{(\beta_{ri}^{(p)*})^2\}$ are i.i.d. bootstrap resamples from $\{(\beta_{ri}^{(p)})^2\}$. It therefore follows from basic bootstrap results that the conditional distribution of $\frac{1}{\sqrt{n_p}} \sum_{i=1}^{n_p} [(\beta_{ri}^{(p)*})^2 - \frac{1}{n_p} \sum_{k=1}^{n_p} (\beta_{rk}^{(p)})^2]$ given \mathcal{X}_p converges to the same $N(0, \Lambda_r^{(p)})$ limit distribution as $\frac{1}{\sqrt{n_p}} \sum_{i=1}^{n_p} [(\beta_{ri}^{(p)})^2 - E\{(\beta_{ri}^{(p)})^2\}]$. Together with the independence of $(\beta_{ri}^{(1)*})^2$ and $(\beta_{ri}^{(2)*})^2$, the assertion of the theorem is an immediate consequence.

Let us turn to Δ_1^* , $\Delta_{2,r}^*$ and $\Delta_{4,L}^*$. Using (41)–(43), it is then easily seen that $n\Delta_1^*$, $n\Delta_{2,r}^*$ and $n\Delta_{4,L}^*$ admit expansions similar to (33), (35) and (36), when replacing there $\frac{1}{\sqrt{n_p}} \sum_{i=1}^{n_p} \beta_{ri}^{(p)}$ by $\frac{1}{\sqrt{n_p}} \sum_{i=1}^{n_p} (\beta_{ri}^{(p)*} - \frac{1}{n_p} \sum_{k=1}^{n_p} \beta_{rk}^{(p)})$, as well as $\frac{1}{\sqrt{n_p}} \sum_{i=1}^{n_p} \beta_{si}^{(p)} \beta_{ri}^{(p)}$ by $\frac{1}{\sqrt{n_p}} \sum_{i=1}^{n_p} (\beta_{si}^{(p)*} \beta_{ri}^{(p)*} - \frac{1}{n_p} \sum_{k=1}^{n_p} \beta_{sk}^{(p)} \beta_{rk}^{(p)})$.

Replacing $\beta_{ri}^{(p)}$, $\beta_{si}^{(p)}$ by $\beta_{ri}^{(p)*}$, $\beta_{si}^{(p)*}$ leads to bootstrap analogs $b_{ik}^{(p)*}$ of the vectors $b_{ik}^{(p)}$, $k = 1, 2, 3$. For any $q \in \mathbb{N}$, define bootstrap versions $n\Delta_1^*(q)$, $n\Delta_{2,r}^*(q)$ and $n\Delta_{4,L}^*(q)$ of $n\Delta_1(q)$, $n\Delta_{2,r}(q)$ and $n\Delta_{4,L}(q)$ by using $(\frac{1}{\sqrt{n_1}} \sum_{i=1}^{n_1} (b_{ik}^{(1)*} - \frac{1}{n_1} \sum_{k=1}^{n_1} b_{ik}^{(1)})^\top, \frac{1}{\sqrt{n_2}} \sum_{i=1}^{n_2} (b_{ik}^{(2)*} - \frac{1}{n_2} \sum_{k=1}^{n_2} b_{ik}^{(2)})^\top)$ instead of $(\frac{1}{\sqrt{n_1}} \sum_{i=1}^{n_1} (b_{ik}^{(1)})^\top, \frac{1}{\sqrt{n_2}} \sum_{i=1}^{n_2} (b_{ik}^{(2)})^\top)$, $k = 1, 2, 3$, in (37). Applying again (41)–(43), one can conclude that for any $\epsilon > 0$ there exists some $q(\epsilon)$ such that, as $n \rightarrow \infty$,

$$\begin{aligned}
(44) \quad &P(|n\Delta_1^* - n\Delta_1^*(q)| < \epsilon) \rightarrow 1, \\
&P(|n\Delta_{2,r}^* - n\Delta_{2,r}^*(q)| < \epsilon) \rightarrow 1, \\
&P(|n\Delta_{4,L}^* - n\Delta_{4,L}^*(q)| < \epsilon) \rightarrow 1
\end{aligned}$$

hold for all $q \geq q(\epsilon)$. Of course, (44) generalizes to the conditional probabilities given $\mathcal{X}_1, \mathcal{X}_2$.

In order to prove the theorem, it thus only remains to show that for *any* given q and all δ

$$(45) \quad |\mathbb{P}(n\Delta(q) \geq \delta) - \mathbb{P}(n\Delta^*(q) \geq \delta) | \mathcal{X}_1, \mathcal{X}_2| = o_p(1)$$

hold for either $\Delta(q) = \Delta_1(q)$ and $\Delta^*(q) = \Delta_1^*(q)$, $\Delta(q) = \Delta_{2,r}(q)$ and $\Delta^*(q) = \Delta_{2,r}^*(q)$, or $\Delta(q) = \Delta_{4,L}(q)$ and $\Delta^*(q) = \Delta_{4,L}^*(q)$. But note that for $k = 1, 2, 3$, $\mathbb{E}(b_{ik}) = 0$, $\{b_{ik}^{(j)*}\}$ are i.i.d. bootstrap resamples from $\{b_{ik}^{(p)}\}$, and $\mathbb{E}(b_{ik}^{(p)*} | \mathcal{X}_1, \mathcal{X}_2) = \frac{1}{n_p} \sum_{k=1}^{n_p} b_{ik}^{(p)}$ are the corresponding conditional means. It therefore follows from basic bootstrap results that as $n \rightarrow \infty$ the conditional distribution of $(\frac{1}{\sqrt{n_1}} \sum_{i=1}^{n_1} (b_{ik}^{(1)*} - \frac{1}{n_1} \sum_{k=1}^{n_1} b_{ik}^{(1)})^\top, \frac{1}{\sqrt{n_2}} \sum_{i=1}^{n_2} (b_{ik}^{(2)*} - \frac{1}{n_2} \sum_{k=1}^{n_2} b_{ik}^{(2)})^\top)$ given $\mathcal{X}_1, \mathcal{X}_2$ converges to the same $N(0, \Omega_{k,q})$ limit distribution as $(\frac{1}{\sqrt{n_1}} \sum_{i=1}^{n_1} (b_{ik}^{(1)})^\top, \frac{1}{\sqrt{n_2}} \sum_{i=1}^{n_2} (b_{ik}^{(2)})^\top)$. This obviously holds for all $q \in \mathbb{N}$, and (45) is an immediate consequence. The theorem then follows from (38), (39), (40), (44) and (45). \square

REFERENCES

- ARAÚJO, A. and GINÉ, E. (1980). *The Central Limit Theorem for Real and Banach Valued Random Variables*. Wiley, New York. [MR0576407](#)
- BESSE, P. and RAMSAY, J. (1986). Principal components of sampled functions. *Psychometrika* **51** 285–311. [MR0848110](#)
- BLACK, F. and SCHOLES, M. (1973). The pricing of options and corporate liabilities. *J. Political Economy* **81** 637–654.
- DAUXOIS, J., POUSSE, A. and ROMAIN, Y. (1982). Asymptotic theory for the principal component analysis of a vector random function: Some applications to statistical inference. *J. Multivariate Anal.* **12** 136–154. [MR0650934](#)
- FENGLER, M. (2005a). Arbitrage-free smoothing of the implied volatility surface. SFB 649 Discussion Paper No. 2005–019, SFB 649, Humboldt-Universität zu Berlin.
- FENGLER, M. (2005b). *Semiparametric Modeling of Implied Volatility*. Springer, Berlin. [MR2183565](#)
- FENGLER, M., HÄRDLE, W. and VILLA, P. (2003). The dynamics of implied volatilities: A common principle components approach. *Rev. Derivative Research* **6** 179–202.
- FENGLER, M., HÄRDLE, W. and MAMMEN, E. (2007). A dynamic semiparametric factor model for implied volatility string dynamics. *Financial Econometrics* **5** 189–218.
- FERRATY, F. and VIEU, P. (2006). *Nonparametric Functional Data Analysis*. Springer, New York. [MR2229687](#)
- FLURY, B. (1988). *Common Principal Components and Related Models*. Wiley, New York. [MR0986245](#)
- GIHMAN, I. I. and SKOROHOD, A. V. (1973). *The Theory of Stochastic Processes. II*. Springer, New York. [MR0375463](#)
- HALL, P. and HOSSEINI-NASAB, M. (2006). On properties of functional principal components analysis. *J. Roy. Statist. Soc. Ser. B* **68** 109–126. [MR2212577](#)
- HALL, P., MÜLLER, H. G. and WANG, J. L. (2006). Properties of principal components methods for functional and longitudinal data analysis. *Ann. Statist.* **34** 1493–1517. [MR2278365](#)

- HALL, P., KAY, J. W. and TITTERINGTON, D. M. (1990). Asymptotically optimal difference-based estimation of variance in nonparametric regression. *Biometrika* **77** 520–528. [MR1087842](#)
- HAFNER, R. (2004). *Stochastic Implied Volatility*. Springer, Berlin. [MR2090447](#)
- HÄRDLE, W. and SIMAR, L. (2003). *Applied Multivariate Statistical Analysis*. Springer, Berlin. [MR2061627](#)
- KAHALÉ, N. (2004). An arbitrage-free interpolation of volatilities. *Risk* **17** 102–106.
- KNEIP, A. and UTIKAL, K. (2001). Inference for density families using functional principal components analysis. *J. Amer. Statist. Assoc.* **96** 519–531. [MR1946423](#)
- LACANTORE, N., MARRON, J. S., SIMPSON, D. G., TRIPOLI, N., ZHANG, J. T. and COHEN, K. L. (1999). Robust principal component analysis for functional data. *Test* **8** 1–73. [MR1707596](#)
- PEZZULLI, S. D. and SILVERMAN, B. (1993). Some properties of smoothed principal components analysis for functional data. *Comput. Statist.* **8** 1–16. [MR1220336](#)
- RAMSAY, J. O. and DALZELL, C. J. (1991). Some tools for functional data analysis (with discussion). *J. Roy. Statist. Soc. Ser. B* **53** 539–572. [MR1125714](#)
- RAMSAY, J. and SILVERMAN, B. (2002). *Applied Functional Data Analysis*. Springer, New York. [MR1910407](#)
- RAMSAY, J. and SILVERMAN, B. (2005). *Functional Data Analysis*. Springer, New York. [MR2168993](#)
- RAO, C. (1958). Some statistical methods for comparison of growth curves. *Biometrics* **14** 1–17.
- RICE, J. and SILVERMAN, B. (1991). Estimating the mean and covariance structure non-parametrically when the data are curves. *J. Roy. Statist. Soc. Ser. B* **53** 233–243. [MR1094283](#)
- SILVERMAN, B. (1996). Smoothed functional principal components analysis by choice of norm. *Ann. Statist.* **24** 1–24. [MR1389877](#)
- TYLER, D. E. (1981). Asymptotic inference for eigenvectors. *Ann. Statist.* **9** 725–736. [MR0619278](#)
- YAO, F., MÜLLER, H. G. and WANG, J. L. (2005). Functional data analysis for sparse longitudinal data. *J. Amer. Statist. Assoc.* **100** 577–590. [MR2160561](#)

M. BENKO
 W. HÄRDLE
 CASE—CENTER FOR APPLIED STATISTICS AND ECONOMICS
 HUMBOLDT-UNIVERSITÄT ZU BERLIN
 SPANDAUERSTR 1
 D-10178 BERLIN
 GERMANY
 E-MAIL: benko@wiwi.hu-berlin.de
haerdle@wiwi.hu-berlin.de
 URL: <http://www.case.hu-berlin.de/>

A. KNEIP
 STATISTISCHE ABTEILUNG
 DEPARTMENT OF ECONOMICS
 UNIVERSITÄT BONN
 ADENAUERALLEE 24-26
 D-53113 BONN
 GERMANY
 E-MAIL: akneip@uni-bonn.de