

Structural Adaptive Models in Financial Econometrics

DISSERTATION

zur Erlangung des akademischen Grades

doctor rerum politicarum
(Doktor der Wirtschaftswissenschaft)
im Fach Statistik

eingereicht an der
Wirtschaftswissenschaftlichen Fakultät
Humboldt-Universität zu Berlin

von
M.Sc. Andrija Mihoci

Präsident der Humboldt-Universität zu Berlin:
Prof. Dr. Jan-Hendrik Olbertz

Dekan der Wirtschaftswissenschaftlichen Fakultät:
Prof. Dr. Ulrich Kamecke

Gutachter:
1. Prof. Dr. Wolfgang Karl Härdle
2. Prof. Dr. Nikolaus Hautsch

eingereicht am: 10. August 2012

Tag der mündlichen Prüfung: 24. September 2012

Abstract

Modern methods in statistics and econometrics successfully deal with stylized facts observed on financial markets. The presented techniques aim to understand the dynamics of financial market data more accurate than traditional approaches. Economic and financial benefits are achievable. The results are here evaluated in practical examples that mainly focus on forecasting of financial data. Our applications include: (i) modelling and forecasting of liquidity supply, (ii) localizing multiplicative error models and (iii) providing evidence for the empirical pricing kernel paradox across countries.

Zusammenfassung

Moderne statistische und ökonometrische Methoden behandeln erfolgreich stilierte Fakten auf den Finanzmärkten. Die vorgestellten Techniken erstreben die Dynamik von Finanzmarktdaten genauer als traditionelle Ansätze zu verstehen. Wirtschaftliche und finanzielle Vorteile sind erzielbar. Die Ergebnisse werden hier in praktischen Beispielen ausgewertet, die sich vor allem auf die Prognose von Finanzmarktdaten fokussieren. Unsere Anwendungen umfassen: (i) die Modellierung und die Vorhersage des Liquiditätsangebotes, (ii) die Lokalisierung des 'Multiplicative Error Model' und (iii) die Erbringung von Evidenz für den empirischen Zustandsfaktorparadox über Ländern.

Acknowledgments

An essential aspect of life is to appreciate those who sacrifice their time and energy in making contributions to the wellbeing of others. It is therefore no mistake that people usually say “give thanks to whom it is due”. Upon reflection of my entire experience at the Humboldt-Universität zu Berlin, I recount instances of challenges and in such situations there have always been people who made meaningful interventions that brought me thus far.

First, I am highly thankful to God, who gave me the strength to sail through the waters of the programme. On the human point of view, I am indebted to my supervisor Prof. Dr. Wolfgang Karl Härdle who has been, and continues to be, a source of inspiration to me in academic matters. There is no doubt whatsoever, that he is my academic mentor. His advice and recommendations have been essential to the start and completion of this work. For these reasons I appreciate him greatly.

I recognize and appreciate furthermore Prof. Dr. Nikolaus Hautsch for a constant source of motivation through my academic experience. My colleagues at the Chairs of Statistics and Econometrics are appreciated for their constant supportive help.

The relentless support of my darling wife Željka continues to bear significance to the success story of my academic endeavors. She and our sons Gabriel and Mihael continue to be a strong pillar within and without academia. Let me take this opportunity of thanking my parents Ivan and Anđela, and my wife’s parents Stjepan and Jelena, whose individual roles to my success cannot be downplayed. Sister Ivana, my wife’s sister Slađana, nephews David, Dino and Bruno and all our relatives are well appreciated for their respective roles.

In closing, I recognize and appreciate my numerous friends in Croatia and Germany, especially Dražen, Valentina, Marko, Josip, Marijan, Vladimir, Razvan and Augustine! To them I am thankful!

Contents

1	Introduction	1
1.1	Modelling and Forecasting Liquidity Supply	1
1.2	Localizing Multiplicative Error Models	4
1.3	Cross Country Evidence for the EPK Paradox	5
2	Theoretical Modelling	7
2.1	Dynamic Semiparametric Factor Model (DSFM)	7
2.1.1	Model Structure	8
2.1.2	Estimation	8
2.1.3	Application Details	9
2.2	Multiplicative Error Model (MEM)	10
2.2.1	Model Structure	10
2.2.2	Distributional Assumptions	11
2.2.3	Estimation Quality	11
2.3	Local Parametric Approach (LPA)	12
2.3.1	Statistical Framework	12
2.3.2	Local Change Point (LCP) Detection Test	13
2.3.3	Adaptive Estimation	14
2.3.4	Critical Values	15
2.4	Pricing Kernel under State-Dependent Utility	18
2.4.1	Economic Setup	18
2.4.2	Pricing Kernel	19
2.4.3	Moment Conditions	19
2.5	Generalized Method of Moments	19
2.5.1	Parameter Estimation	20
2.5.2	Hypothesis Testing	21
3	Data	23
3.1	Stock Markets around the Globe	23
3.1.1	Descriptive Statistics	23
3.1.2	Portfolio Selection	24
3.2	Trading at the Australian Stock Exchange (ASX)	26
3.3	Trading at the NASDAQ Stock Market	29

4	Applications	33
4.1	Modelling and Forecasting Liquidity Supply using Semiparametric Factor Dynamics	33
4.1.1	Limit Order Book Modelling using the DSFM	34
4.1.2	Modelling Limit Order Book Dynamics	39
4.1.3	Drivers of the Order Book Shape	42
4.1.4	Forecasting Liquidity Supply	48
4.1.5	Forecasting Results	49
4.1.6	Financial and Economic Applications	51
4.2	Local Adaptive Multiplicative Error Models for High-Frequency Forecasts	56
4.2.1	Parameter Dynamics	56
4.2.2	Adaptive Estimation	61
4.2.3	Empirical Findings	62
4.2.4	Forecasting Trading Volumes	64
4.2.5	Financial Applications	71
4.3	Cross Country Evidence for the EPK Paradox	74
4.3.1	Modelling Setup	74
4.3.2	Estimation Quality	74
4.3.3	Parameter Dynamics and Reference Point Analysis	76
4.3.4	Empirical Pricing Kernels across Stock Markets	78
5	Conclusions	79
5.1	Modelling and Forecasting Liquidity Supply	79
5.2	Localizing Multiplicative Error Models	80
5.3	Cross Country Evidence for the EPK Paradox	81

List of Figures

1.1	Limit order book for selected stocks traded at the ASX on July 8, 2002 at 10:15. Red: bid curve, blue: ask curve.	3
2.1	Graphical illustration of sequential testing for parameter homogeneity in interval I_k with length $n_k = I_k $ ending at fixed time point i_0 . Suppose we have not rejected homogeneity in interval I_{k-1} , we search within the interval $J_k = I_k \setminus I_{k-1}$ for a possible change point τ . The red interval marks $A_{k,\tau}$ and the blue interval marks $B_{k,\tau}$ (blue) splitting the interval I_{k+1} into two parts depending upon the position of the unknown change point τ	14
2.2	Simulated critical values of an EACD(1,1) model for the 'moderate risk case' ($r = 0.5$, upper panel) and the 'conservative risk case' ($r = 1$, lower panel), $\rho = 0.25$, $K = 8$ and chosen parameters constellations according to Table 4.11. The low (blue), middle (green) and upper (red) curves are associated with the corresponding ratio levels $\tilde{\beta}/(\tilde{\alpha} + \tilde{\beta})$	16
2.3	Simulated critical values of an EACD(1,1) model for the 'moderate risk case' ($r = 0.5$, upper panel) and the 'conservative risk case' ($r = 1$, lower panel), $\rho = 0.25$, $K = 13$ and chosen parameters constellations according to Table 4.11. The low (blue), middle (green) and upper (red) curves are associated with the corresponding ratio levels $\tilde{\beta}/(\tilde{\alpha} + \tilde{\beta})$	16
2.4	Simulated critical values of a WACD(1,1) model for the 'moderate risk case' ($r = 0.5$, upper panel) and the 'conservative risk case' ($r = 1$, lower panel), $\rho = 0.25$, $K = 8$ and chosen parameters constellations according to Table 4.11. The low (blue), middle (green) and upper (red) curves are associated with the corresponding ratio levels $\tilde{\beta}/(\tilde{\alpha} + \tilde{\beta})$	17
2.5	Simulated critical values of a WACD(1,1) model for the 'moderate risk case' ($r = 0.5$, upper panel) and the 'conservative risk case' ($r = 1$, lower panel), $\rho = 0.25$, $K = 13$ and chosen parameters constellations according to Table 4.11. The low (blue), middle (green) and upper (red) curves are associated with the corresponding ratio levels $\tilde{\beta}/(\tilde{\alpha} + \tilde{\beta})$	17
3.1	Market returns for selected stock market indices on major stock markets from 1 January 1990 until 31 May 2012 (5827 observations per series). . .	24

List of Figures

3.2	Coefficients of determination (R^2) implied by linear regression of 1 min (red) and 5 min (blue) mid-quote returns on lagged order imbalances for selected stocks traded at the ASX from 8 July to 16 August 2002 (30 trading days). The horizontal axis depicts the number of included imbalance levels.	28
3.3	Estimated intra-day seasonality factors for quantities offered at best bid prices (red) and for quantities supplied at best ask prices (blue) across selected stocks traded at the ASX from 8 July to 16 August 2002 (30 trading days).	30
3.4	Estimated intra-day periodicity components for cumulative one-minute trading volumes (in units of 100,000 and plotted against the time of the day) of selected companies at NASDAQ on 2 September (blue, lowest 30-day trading volume) and 30 October 2008 (red, highest 30-day volume).	32
4.1	Root mean squared errors (RMSEs) for different absolute price levels, $\check{S}_{t,j}^b$ (red) and $\check{S}_{t,j}^a$ (blue), using the DSFM-Separated (solid) and the DSFM-Combined approach (dashed).	35
4.2	Estimated first and second factor of the limit order book depending on relative price levels using the DSFM-separated approach with two factors for selected stocks traded at the ASX from July 8 to August 16, 2002 (30 trading days). Red: bid curve, blue: ask curve.	36
4.3	Estimated first and second factor loadings of the limit order book depending on relative price levels using the DSFM-separated approach with two factors for selected stocks traded at the ASX from July 8 to August 16, 2002 (30 trading days). Red: bid curve, blue: ask curve.	37
4.4	True (solid) and estimated (dashed) limit order book using the DSFM-separated approach with two factors ($EV \approx 95\%$) on 8 July 2002 for NAB. Red: bid curve, blue: ask curve.	38
4.5	Orthogonalized impulse-response analysis: responses of the best bid quote return to a one standard deviation shock in the estimated first bid factor loadings (upper panel) and response of the best ask quote return to a one standard deviation shock in the estimated first ask factor loadings (lower panel). We employ the DSFM-separated approach with two factors and a VEC specification for selected stocks traded at the ASX from 8 July to 16 August 2002 (30 trading days). The response variable always enters the VEC specification in the first position. 95% confidence intervals are shown with dashed lines.	41
4.6	Estimated first factors of the bid side with respect to relative price levels and the past log traded sell volume using the DSFM-Separated approach with two factors for selected stocks traded at the ASX from 8 July to 16 August 2002 (30 trading days).	43

4.7	Estimated first factors of the ask side with respect to relative price levels and the past log traded buy volume using the DSFM-Separated approach with two factors for selected stocks traded at the ASX from 8 July to 16 August 2002 (30 trading days).	44
4.8	Estimated first factors of the bid side with respect to relative price levels and the best bid price using the DSFM-Separated approach with two factors for selected stocks traded at the ASX from 8 July to 16 August 2002 (30 trading days).	45
4.9	Estimated first factors of the ask side with respect to relative price levels and the best ask price using the DSFM-Separated approach with two factors for selected stocks traded at the ASX from 8 July to 16 August 2002 (30 trading days).	46
4.10	Predicted limit order book curves (dashed) and the true ones (solid) on July 22, 2002, at 11:00 (upper panels) and 15:00 (lower panels) at different absolute price levels in AUD. The predicted curve using the naive approach is shown with black solid line.	50
4.11	Root mean squared prediction errors (RMSPEs) implied by the DSFM-separated approach with two factors for the bid side (red) as well as the ask side (blue) and by the naive approach (black) for all intra-day forecasting horizons (in hours) for selected stocks traded at the ASX. Prediction period: July 22 to August 16, 2002 (20 trading days).	53
4.12	Average percentage gains by reduced transaction costs compared to an equal-splitting strategy when buying (blue) and selling (red) shares based on m DSFM-predicted time points per day. Upper panel: Daily volumes corresponding to 5 (2) times the average first level market depth for BHP, NAB, WOW (MIM). Lower panel: Daily volumes corresponding to 10 (5) times the average first level market depth for BHP, NAB, WOW (MIM). Prediction period: 22 July to 16 August 2002 (20 trading days).	53
4.13	Predicted demand and supply elasticities at best bid (red) and best ask prices (blue) using the DSFM-separated approach with two factors for selected stocks traded at the ASX from 22 July to 2 August 2002 (upper panels, 10 trading days) and from 5 August to 16 August 2002 (lower panels, 10 trading days).	55
4.14	Predicted (dashed) and realized (solid) limit order book curves for BHP on 22 July 2002, between 15:05-15:20 using the DSFM-Separated approach with two factors.	56
4.15	Time series of estimated 'weekly' (left panel, rolling windows covering 1800 observations) and 'daily' (right panel, rolling windows covering 360 observations) EACD(1, 1) parameters and functions thereof based on seasonally adjusted one-minute trading volumes for Intel Corporation (INTC) at each minute from 22 February to 31 December 2008 (215 trading days). First 35 days are used for initialization. Based on 154,800 individual estimations.	57

4.16	Kernel density plots (Gaussian kernel with optimal bandwidth) of estimated EACD(1,1) parameters for seasonally adjusted trading volumes over weekly (red) and daily window (blue).	58
4.17	Kernel density plots (Gaussian kernel with optimal bandwidth) of estimated WACD(1,1) parameters for seasonally adjusted trading volumes over weekly (red) and daily window (blue).	60
4.18	Estimated length of <i>intervals of homogeneity</i> $n_{\hat{k}}$ (in hours) for seasonally adjusted one-minute cumulative trading volumes of selected companies in case of a modest ($r = 0.5$, blue) and conservative ($r = 1$, red) modelling risk level. We use the interval scheme with $K = 13$ and $\rho = 0.25$. Underlying model: EACD(1,1). NASDAQ trading on 22 February 2008. .	62
4.19	Distribution of estimated interval length $n_{\hat{k}}$ (in hours) for seasonally adjusted trading volumes of selected companies in case of modest ($r = 0.5$, red) and conservative modelling risk ($r = 1$, blue), using an EACD (upper panel) and a WACD model (lower panel) from 22 February to 31 December 2008 (215 trading days). We select 13 estimation windows based on significance level $\rho = 0.25$	63
4.20	Average estimated interval length $n_{\hat{k}}$ (in hours) over the course of a trading day for seasonally adjusted trading volumes of selected companies in case of modest ($r = 0.5$, red) and conservative modelling risk ($r = 1$, blue), using an EACD (upper panel) and a WACD model (lower panel) from 22 February to 31 December 2008 (215 trading days). We select 13 based on significance level $\rho = 0.25$	64
4.21	Coefficient of determination R^2 of forecasting regressions (4.6) and (4.7) computed at fixed horizon $h = 1, \dots, 60$ for the EACD model from 22 February to 22 December 2008 (210 trading days). Results are shown for the LPA technique with 'significance' level $\rho = 0.25$ for two risk levels, i.e., $r = 0.5$ (solid red line) and $r = 1$ (dashed red line), as well as for two specifications of the standard method, i.e., ad hoc selected estimation window of 360 (solid blue line) and 1800 observations (dashed blue line). .	65
4.22	Test statistic $T_{DM,h}$ across all 60 forecasting horizons for five large companies traded at NASDAQ from 22 February to 22 December 2008 (210 trading days). The red curve depicts the statistic based on a test of the LPA against a fixed-window scheme using 360 observations (6 trading hours). The blue curve depicts the statistic based on a test of the LPA against a fixed-window scheme using 1800 observations (30 trading hours). The upper panel shows the results for the 'modest risk case' ($r = 0.5$) and the lower panel shows the results for the 'conservative risk case' ($r = 1$) given a significance level of $\rho = 0.25$	68
4.23	Ratio between the RMSPEs of the LPA and of a fixed-window approach (covering 6 trading hours) over the sample from 22 February to 22 December 2008 (210 trading days). Upper panel: EACD model, lower panel: WACD model.	70

4.24	Ratio between the RMSPEs of the LPA and of a fixed-window approach (covering 6 trading hours) over the sample from 22 February to 22 December 2008 (210 trading days). Upper panel: Results for underlying (local) EACD model. Lower panel: Results for underlying (local) WACD model.	70
4.25	Ratio between the RMSPEs of the LPA method and the RMSPE of the standard approach over the course a typical trading day using an EACD (upper panel) and WACD (lower panel) model from 22 February to 22 December 2008 (210 trading days).	71
4.26	Daily cash flow in USD (blue) and cumulated daily cash flow in USD (red) of the investment strategy. The investor uses an EACD (left panel) and a WACD model (right panel) from 22 February to 22 December 2008 (210 trading days).	73
4.27	Time series of the estimated parameters β_1 (red) and β_2 (blue) across six worldwide largest stock markets. We employ the iterated GMM estimation technique with $n = 500$ (2 years).	76
4.28	Time series of the estimated parameter β across six worldwide largest stock markets. We employ the GMM estimation technique with the HJ weighting matrix and with $n = 500$ (2 years).	77
4.29	Kernel density plots (Gaussian kernel with optimal bandwidth) of optimal reference point x . We employ the iterated GMM estimation technique with $n = 500$ (2 years).	77
4.30	Empirical pricing kernels across six worldwide largest stock markets for two scenarios, i.e., (a) $\beta_1, \beta_2 > 0$ (blue) and (c) $\beta_1 = \beta_2 = \beta$ (red). The empirical pricing kernels are fitted to average parameter estimates and the average value of the optimal reference point. We employ the iterated GMM estimation technique in (a) and the GMM estimation technique with the HJ weighting matrix in (c) with $n = 500$ (2 years).	78

List of Tables

3.1	Descriptive statistics of monthly returns of the selected market indices from 1 January 1990 until 31 May 2012 (5827 observations per series). . .	24
3.2	Average return of selected portfolios for companies with a high book-to-market ratio from 1 January 1990 until 31 May 2012.	25
3.3	Average return of selected portfolios for companies with a medium book-to-market ratio from 1 January 1990 until 31 May 2012.	25
3.4	Average return of selected portfolios for companies with a low book-to-market ratio from 1 January 1990 until 31 May 2012.	26
3.5	Total number of market and limit orders for selected stocks traded at the ASX from 8 July to 16 August 2002.	27
3.6	Descriptive statistics and Ljung-Box statistics (based on 10 lags) of daily and one-minute cumulated trading volumes of five large companies traded at NASDAQ between January 2 and December 31, 2008 (250 trading days, 90000 observations per stock).	31
4.1	Explained variance (EV) of estimated order book variations depending on relative prices based on different number of factors L using both DSFM approaches.	34
4.2	Root mean squared errors (RMSEs) implied by estimated order book variations depending on relative prices based on different number of factors L using both DSFM approaches.	35
4.3	Explained variance (EV) of estimated order books depending on relative prices based on different number of factors L and price grid points using the DSFM-Separated approach	37
4.4	Schmidt-Phillips test statistics for estimated factor loadings (H_0 : unit root, critical values are -15.0, -18.10 and -25.20 for significance levels 10%, 5% and 1%, respectively.)	38
4.5	KPSS test statistics for estimated factor loadings (H_0 : weak stationarity, critical values are 0.12, 0.15 and 0.22 for significance levels 10%, 5% and 1%, respectively.)	38
4.6	Root mean squared error (RMSE) of the estimated limit order book data for all selected stocks based on the traded buy quantities evaluated for different number of factors L using both DSFM approaches	47
4.7	Root mean squared error (RMSE) of the estimated limit order book data for all selected stocks based on the traded sell quantities evaluated for different number of factors L using both DSFM approaches	47

4.8	Root mean squared error (RMSE) of the estimated limit order book data for all selected stocks based on the intraday log-return evaluated for different number of factors L using both DSFM approaches	50
4.9	Average root mean squared prediction errors (RMPSEs) of both limit order book sides implied by the DSFM-separated approach with two factors and the naive model for selected stocks traded at the ASX in the period from 22 July to 6 August 2002 (20 forecasting days).	51
4.10	Quartiles of estimated persistence levels $(\tilde{\alpha} + \tilde{\beta})$ for all five stocks at each minute from 22 February to 31 December 2008 (215 trading days) and six lengths of local estimation windows based on EACD and WACD specifications. We label the first quartile as 'low', the second quartile as 'moderate' and the third quartile as 'high'.	59
4.11	Quartiles of 774,000 estimated ratios $\tilde{\beta}/(\tilde{\alpha} + \tilde{\beta})$ (based on estimation windows covering 1800 observations) for all five stocks at each minute from 22 February to 31 December 2008 (215 trading days) and both model specifications (EACD and WACD) conditional on the persistence level (low, moderate or high). We label the first quartile as 'low', the second quartile as 'mid' and the third quartile as 'high'. The shape parameter for the WACD model equals the median value in all cases ($\tilde{s} = 1.57$). . . .	61
4.12	Number of non-rejections of the unbiasedness or the efficiency null hypothesis given horizon $h = 1, \dots, 60$ of five large companies traded at the NASDAQ from 22 February to 22 December 2008 (210 trading days) using an EACD(1,1) and a WACD(1,1) model at a 5% significance level. We specify four tuning parameter constellations when using the LPA technique and two ad hoc selected window lengths when employing the standard method (1 week or 1 day). Maximum number of non-rejections is 60 in each case.	66
4.13	Number of non-rejections of the 'forecasting information coverage' null hypothesis ($H_0 : \eta_1 = 1$ and $H_0 : \eta_2 = 0$) given horizon $h = 1, \dots, 60$ of five large companies traded at the NASDAQ from 22 February to 22 December 2008 (210 trading days) using an EACD(1,1) and a WACD(1,1) model with a significance level of 5%. We specify four tuning parameter constellations when comparing the LPA technique to the standard method with an ad hoc selected window length of 1800 observations.	67
4.14	Largest (in absolute terms) test statistic $T_{ST,h}$ across all 60 forecasting horizons as well as EACD and WACD specifications for five large companies traded at NASDAQ from 22 February to 22 December 2008 (210 trading days). We compare LPA-implied forecasts with those based on rolling windows using a priori fixed lengths of one week and one day, respectively. Negative values indicate lower squared prediction errors resulting from the LPA. According to the Diebold-Mariano test (4.10), the average loss differential is significantly negative in all cases (significance level 5%).	69

List of Tables

4.15	Average optimal objective function value across six worldwide largest markets for two competing estimation techniques and three scenarios: (a) $\beta_1, \beta_2 > 0$, (b) $\beta_1 > \beta_2 > 0$ and (c) $\beta_1 = \beta_2 = \beta = 0$. The estimation window covers $n = 250$ observations (1 year).	75
4.16	Average optimal objective function value across six worldwide largest markets for two competing estimation techniques and three scenarios: (a) $\beta_1, \beta_2 > 0$, (b) $\beta_1 > \beta_2 > 0$ and (c) $\beta_1 = \beta_2 = \beta = 0$. The estimation window covers $n = 500$ observations (2 years).	75
4.17	Average optimal objective function value across six worldwide largest markets for two competing estimation techniques and three scenarios: (a) $\beta_1, \beta_2 > 0$, (b) $\beta_1 > \beta_2 > 0$ and (c) $\beta_1 = \beta_2 = \beta$. The estimation window covers $n = 1250$ observations (5 years).	75
4.18	Percentage of rejections of the null hypothesis of the D -test ($H_0 : \beta_1 = \beta_2 = \beta$) as indicator for the existence of the EPK paradox across the worldwide largest six stock markets. We employed two GMM estimation techniques.	78

1 Introduction

The key to growth is the introduction of higher dimensions of consciousness into our awareness.

Lao Tzu

It is a challenging task to understand the dynamics of financial market processes. The goal of this work is to mitigate this challenge from an academic and a practical perspective. Our theoretical framework discusses recently developed statistical and econometric (structural and adaptive) techniques adequate for modelling and forecasting of financial data. Applying the framework to transaction and market data, we achieve (and evaluate) economic and financial benefits. Special care is given to the treatment of high-frequency data since the ultimate research goal in academia and in practice is to understand the time evolution of transaction level events. At glance, our work deals with high-dimensional and time-varying data structures and demonstrates the full power of non- and semiparametric techniques.

The theoretical background is provided in Chapter 2. When dealing with high-dimensional data structures evolving over time on a high frequency, we recommend to employ a successful dimension-reduction technique, the, so-called, dynamic semiparametric factor model (DSFM). A powerful approach applied in modelling and forecasting univariate time series is the local adaptive multiplicative error model (MEM). The local MEM is based on the local parametric approach (LPA) which was gradually introduced into the econometric literature. While studying the market microstructure, in particular asset pricing, a state-dependent pricing kernel specification allows to investigate preferences over time in a realistic setup. The methodology enables us to provide evidence on the empirical pricing kernel (EPK) paradox across stock markets.

The financial data collected at the world's leading stock markets is described in Chapter 3. The theoretical methods are illustrated in the following key applications in Chapter 4: (i) Modelling and forecasting of liquidity supply, (ii) Local adaptive multiplicative error models and (iii) Cross country evidence for the empirical pricing kernel paradox. Our applications cover structural and adaptive modelling of data and achieve financial and economic benefits, as summarized in Chapter 5. In this introductory chapter we correspondingly motivate all key applications.

1.1 Modelling and Forecasting Liquidity Supply

Electronic limit order book (LOB) trading has become the dominant trading form for equities. The limit order book provides important information about the current liquidity supply, i.e., excess supply for shares on the market. The volume quoted above

the realized market quantity summarizes traders' price expectations, the current implied trading costs as well as the marginal costs (i.e., the demand and supply elasticities). As a high-dimensional object, the limit order book changes on a transaction level basis.

The dynamic behaviour of liquidity supply has been firstly studied by Härdle et al. [2012a]. Due to the flexibility of the employed dynamic semiparametric factor model it is now possible to model the high-dimensional bid and ask curves in a dynamic setting. Our work rewrites the cited study and provides more implementation details. In financial and economic applications we focus on liquidity supply prediction and the improvement of trading strategies.

The underlying modelling idea is to capture the shape of high-dimensional ask and bid curves by a low dimensional factor structure. The factors are estimated non-parametrically ('smooth in space'). A dynamic semiparametric factor model (DSFM) enables us to capture the shape of order schedules by a non-parametric factor structure while the curves' dynamic behavior is driven by time-varying factor loadings. The latter are modelled parametrically employing a vector error correction (VEC) specification ('parametric in time'). The governing modelling philosophy of the DSFM is thus 'smooth in space and parametric in time'.

Recent empirical literature is enriched by our evidence on the dynamics and predictability of order book schedules and we complement recent (theoretical) work on order splitting and dynamic order execution strategies. For example, splitting a large order is highly relevant in financial practice. Optimal splitting strategies require predictions of liquidity demand and liquidity supply, see, e.g., Obizhaeva and Wang [2005] and Engle and Ferstenberg [2007]. Execution strategies are derived by minimizing expected costs of order execution, see, e.g., Bertsimas and Lo [1998] and Almgren and Chriss [2000] and they are analyzed in a limit order book market by Alfonsi et al. [2010]. The studies require knowledge about the (future) order book shape which we can provide.

The study by Härdle et al. [2012a] is the first to model jointly the shape and the dynamics of liquidity supply. Their methodology is used to improve order execution strategies. Note that not only the future shape of the high-dimensional limit order book is predicted, but also the curves' position. Quotes are shown to be in this context short term predictable based on the shape knowledge. Other studies discuss only partially the shape, dynamics or order splitting strategies. We refer here to studies by, e.g., Biais et al. [1995], Griffiths et al. [2000], Ahn et al. [2001], Ranaldo [2004], Hollifield et al. [2004], Bloomfield et al. [2005], Degryse et al. [2005], Hall and Hautsch [2006], Hall and Hautsch [2007], Large [2007], Hasbrouck and Saar [2009] and Cao et al. [2009]. In the central focus of recent literature are furthermore the analysis of liquidity risks, see, e.g., Johnson [2008], Liu [2009], Garvey and Wu [2009] or Goyenko et al. [2009], and the treatment of liquidity costs, see, e.g., Chacko et al. [2008] and Hasbrouck [2009].

The bid and ask curves are in our study high-dimensional objects. Our objective is to capture quantities close to the best quotes as well as volumes more deeply in the book. A snapshot of a typical limit order book for four stocks traded at the Australian Securities Exchange (ASX) is provided in Figure 1.1. Looking at the order book dynamics, we note that volume can be substantially dispersed over a wider range of price levels. Dynamic and individual modelling of all volume levels becomes complicated and intractable.




Figure 1.1: Limit order book for selected stocks traded at the ASX on July 8, 2002 at 10:15. Red: bid curve, blue: ask curve.

The high dimensionality of the order book is reduced by a the so-called dynamic semiparametric factor model (DSFM), proposed by Fengler et al. [2007], Brüggemann et al. [2008], Park et al. [2009] and Cao et al. [2009]. The shape of the book is captured by latent factors which are defined on a grid space around the best ask or bid quotes and can depend on explanatory variables, i.e., the state of the market. The factors are in the first step estimated nonparametrically. In the second step the corresponding factor loadings are modelled jointly with the best bid and the best ask quotes using a VEC specification.

The research questions are: (i) How many factors are sufficient to model liquidity supply reasonably well? (ii) What does the shape of the factors look like? (iii) What do the dynamics of the estimated factor loadings look like? (iv) Does there exist evidence for a strong cross-dependence between both sides of the order book? (v) Can quotes be predictable in the short run? (vi) Does the shape of the order book curves depend on past price movements, past trading volume as well as past volatility? (vii) How successful is the model in predicting future liquidity supply and can it be used to improve order execution strategies?

Applying the methodology to limit order book data of four stocks traded at the ASX one observes that approximately 95% of the order book variations can be explained by two factors. The first factor captures the overall order book slope, whereas the second factor is associated with its curvature. The estimated factor loadings follow highly persistent though stationary dynamics. There are less significant spill-over effects between the bid and the ask side of the market. Quotes are predictable in the short run given a knowledge upon the order book shape. Recent liquidity demand, past returns and corresponding (realized) volatility have an effect on the shape of the order book. The DSFM approach outperforms a naive prediction, where the current book is used as a predictor, in a realistic forecasting exercise. The intra-day order execution strategies are improved by our approach, i.e., our model reduces the implied transaction costs.

1.2 Localizing Multiplicative Error Models

Researchers in academia and in practice aim to understand the dynamics of processes when all single events are recorded. The ultimate goal is to account for external shocks and structural shifts on financial markets. The present study expands the local adaptive multiplicative error model (MEM) study by Härdle et al. [2012b] which can be used to accomplish the above mentioned goals. The methodology is suitable to be applied to (positive-valued) univariate time series, such as, for example, durations, bid-ask spreads, trading volumes, transactions costs or volatilities.

A multiplicative error model (MEM) by Engle [2002], serves as a workhorse for the modelling of positive valued, serially dependent high-frequency data. It has been successfully applied to, for example, financial duration data by Engle and Russell [1998] in the context of an autoregressive conditional duration (ACD) model, or to intraday trading volumes, see, e.g., Manganelli [2005], Brownlees et al. [2011] and Hautsch and Huang [2011]. The model parameters are typically estimated over long estimation windows in order to increase estimation efficiency. Empirical evidence makes parameter constancy in high-frequency models over long time intervals questionable.

Structural breaks in MEM parameters have been reported in current literature, see, e.g., Zhang et al. [2001] who identify regime shifts in trade durations and suggest a threshold ACD (TACD) specification in the spirit of threshold ARMA models, see, e.g., Tong [1990]. A smooth transition ACD (STACD) model may capture the transition of parameters between different states. Regime-switching MEM approaches allow for changing parameters on possibly high frequencies (in the extreme case from observation to observation). They, however, require to impose a priori structures on the transition form, the number of underlying regimes or on the type of the transition variable.

The local MEM adapts a local parametric approach and only assumes that financial data locally (i.e., over short data intervals) follow an underlying MEM model. The *local parametric approach* (LPA), originally proposed by Spokoiny [1998], has been gradually introduced into the time series literature. For applications to daily exchange rates, see, e.g., Mercurio and Spokoiny [2004] and Čížek et al. [2009] for an adaptation of the approach to GARCH models in modelling of daily market index volatility. In realized volatility analysis, the LPA has been applied by Chen et al. [2010] to daily stock index returns.

Implementing a local parametric framework for multiplicative error processes, Härdle et al. [2012b] illustrate its usefulness when it comes to out-of-sample forecasting under possibly non-stable market conditions. A flexible statistical approach presented here allows to statistically select a data window over which it is appropriate to fit a local constant-parameter model. Insights into the time evolution of high-frequency data are provided. The length of (local) estimation windows is data-driven, i.e., a sequential testing procedure is used to determine the so-called *interval of homogeneity*. Within this interval one can safely fit a MEM with constant (homogeneous) parameters.

The *interval of homogeneity* leads to the adaptive estimate which in turn is used to produce (multi-step ahead) forecasts of financial data. These steps are repeated in every period. Period-to-period variations in parameters are therefore captured and

rolling-window out-of-sample forecasts account only for information which is statistically identified as being 'relevant'.

In order to control liquidity risk, i.e., the exposure of a trader's position to volume dynamics, we focus on one-minute cumulative trading volumes of five highly liquid stocks traded at NASDAQ. Note that our findings may be carried over to other high-frequency (univariate) series, as the stochastic properties of high-frequency volumes are quite similar to those of other processes, such as, trade counts, squared midquote returns, market depth or bid-ask spreads.

Our research questions include: (i) How strong is the time variation of MEM parameters? (ii) What are typical interval lengths of parameter homogeneity implied by the local MEM? (iii) How good are out-of-sample short-term forecasts compared to adaptive procedures where the length of the estimation windows is fixed on an ad hoc basis? (iv) Is it possible to achieve financial gains using the LPA?

Based on trading volumes at the NASDAQ stock market, one observes that MEM parameters and the estimation quality considerably change over time. On average a more precise adaptive estimates require local estimation windows of approximately 3 to 4 hours. A less conservative approach would select 2-3 hours of data. The local MEM yields statistically better short-term forecasts than competing approaches using fixed-length rolling windows of comparable sizes. It also results in financial benefits when considering a (hypothetical) trading strategy. Implementing the proposed framework requires re-estimating and re-evaluating the model based on rolling windows of different lengths from minute to minute, yielding extensive insights into the time-varying nature of high-frequency trading processes.

1.3 Cross Country Evidence for the EPK Paradox

In market microstructure literature, the empirical pricing kernel (EPK) paradox has been reported in the options markets with payoffs dependent on the stock index holdings. The 'anomaly' regarding the monotonicity of the pricing kernel has been shown to be time persistent and significant.

On stock markets there does not exist much evidence about the (statistical) existence or the persistence of the EPK paradox. It appears as a U-shaped functional estimate, see, e.g., Dittmar [2002] who investigated equity returns of US data. Schweri [2010] often rejects the non-monotonicity (more precisely the U-shaped EPK form) for high values of wealth due to sparsity of data.

Following a recent and promising approach, we consider the pricing kernel derived under state-dependent utility, see, e.g., Grith et al. [2011]. The methodology is here applied to portfolios at six worldwide largest stock markets.

We aim answering the following research questions: (i) Does there exist an EPK paradox when considering cross-sections of equity returns? (ii) How well does the proposed methodology (i.e. state dependent preference specification) explains the EPK paradox? (iii) How strong is the variation of the EPK and its parameters over time? (iv) Is the locally increasing EPK part statistically significant? (v) How strong is the cross country

1 Introduction

variation of the EPK in equity returns?

Based on stock market data at the worldwide leading markets we show that statistically there exists a EPK paradox. The results are quite robust across countries and for the underlying framework specifications. Estimated EPK preference parameters exhibit a time-varying pattern.

2 Theoretical Modelling

If the facts don't fit the theory, change the facts.
Albert Einstein

In this chapter we introduce the theoretical background relevant for the econometric modelling of financial market data. We firstly introduce the dynamic semiparametric factor model (DSFM) that deals with high-dimensional data structures evolving over time. The model's power lies in the joint modelling of the object's spatial structure (e.g. the LOB's shape) and its temporal structure related to the most influential economic variables, such as, e.g. the observed prices and liquidity demand, the bid-ask spread, realized volatility or stock returns. We will therefore use the model for prediction of liquidity supply, i.e. the order book's shape and its (relative or absolute) position depending upon the current market conditions, see Chapter 4.

Secondly, in the context of univariate financial time series modelling, the localized Multiplicative Error Model (MEM) addresses the question of time-varying MEM parameters. A statistical technique, the so-called local parametric approach (LPA), is used to locally (i.e. over short time intervals) approximate a financial time series' dynamics by a MEM structure. Here we strike a balance between parameter variability and the modelling bias. The presented methodology allows to model any (positive) financial process, such as, e.g., durations, bid-ask spreads, trading volumes, stock volatilities or trading costs. We will illustrate its usefulness in a case study of predicting trading volumes for stocks over the course of a trading day, see Chapter 4.

Finally, focusing on the market microstructure modelling, we investigate how agents price assets on stock markets. The theory deals with the derivation of a pricing kernel under state-dependent utility, as well as with a testing procedure on parameter significance. By applying the introduced methodology in Chapter 4, we provide worldwide evidence for the empirical pricing kernel paradox (i.e. non-monotonicity of the pricing kernel) on stock markets.

2.1 Dynamic Semiparametric Factor Model (DSFM)

In modelling liquidity supply we observe a high-dimensional object of order volume inventories moving on a high frequency. We therefore apply a flexible statistical framework that allows us to reduce the object's dimensionality while preserving the spatial structure of the data, to parametrize the temporal dependence of the order book, as well as to relate the book's shape and dynamics to various economic variables.

Denote the (periodicity adjusted) volumes pending on the bid side at time t by $Y_{t,j}^b \in \mathbb{R}^J$ and that on the ask side by $Y_{t,j}^a \in \mathbb{R}^J$. The seasonal adjustment of data is explained

2 Theoretical Modelling

in Section 3.2. Our analysis considers four different depth levels of the book, namely $J \in \{25, 50, 75, 101\}$. Due to the present modelling challenges we have to strike a balance between information loss on the spatial and temporal structure of the process, involved computational costs and the stability of employed numerical techniques.

In order to reduce the dimensionality a factor decomposition is commonly applied. A high-dimensional object's shape is then explained by few common (parametric) factors. For example, in modelling yield curves, Nelson and Siegel [1987] propose a parametric factor model based on Laguerre polynomials.

Since there is no (theoretical) form for the shape or dynamics of limit order book curves, we capture the book's spatial structure non-parametrically employing a Dynamic Semiparametric Factor Model (DSFM), proposed by Fengler et al. [2007], Brüggemann et al. [2008], Park et al. [2009] and Cao et al. [2009]. As the model is stipulated under the modelling philosophy *smooth in space and parametric in time* it combines the advantages of a nonparametric approach (spatial dependence) and parametric modelling (multivariate time series).

2.1.1 Model Structure

Let a random vector $Y_{t,j} \in \mathbb{R}^J$ be decomposed based on the orthogonal L -factor model

$$Y_{t,j} = m_{0,j} + Z_{t,1}m_{1,j} + \dots + Z_{t,L}m_{L,j} + \varepsilon_{t,j}, \quad (2.1)$$

with *time-invariant factors* $m(\cdot) = (m_0, m_1, \dots, m_L)^\top$, $m_l : \mathbb{R}^d \rightarrow \mathbb{R}$, $l = 0, \dots, L$ and *factor loadings* $Z_t = (1_T, Z_{t,1}, \dots, Z_{t,L})^\top$. Here $\varepsilon_{t,j}$ represents a white noise error term. The time index is denoted by $t = 1, \dots, T$, whereas the cross-sectional (i.e. order book level) index is $j = 1, \dots, J$.

The DSFM allows the factors m_l to depend upon explanatory variables and thus can be seen as a generalization of the factor model (2.1)

$$Y_{t,j} = \sum_{l=0}^L Z_{t,l}m_l(X_{t,j}) + \varepsilon_{t,j} = Z_t^\top m(X_{t,j}) + \varepsilon_{t,j}, \quad (2.2)$$

assuming that the processes $X_{t,j}$, $\varepsilon_{t,j}$ and Z_t are independent. Moreover, the number of factors L should not exceed the object's dimension J .

As explanatory variable $X_{t,j}$, we consider either the 'relative price levels' on the bid side $S_{t,j}^b$ or those on the ask side $S_{t,j}^a$. In studying the order book shape predictability we consider additionally one of the three key (weakly exogenous) market variables: past 5-min aggregated trading volume on both sides of the market representing the recent liquidity demand, the past 5-min log mid-quote return as well as the past 5-min volatility, see Section 4.1.3.

2.1.2 Estimation

Factors m_l are estimated using a series estimator, see, e.g. Park et al. [2009]. For $K \geq 1$, one selects functions $\psi_k : [0, 1]^d \rightarrow \mathbb{R}$, $k = 1, \dots, K$, such that $\int \psi_k^2(x) dx = 1$

2.1 Dynamic Semiparametric Factor Model (DSFM)

holds. Park et al. [2009] select tensor B-spline basis functions, whereas Fengler et al. [2007] use a kernel smoothing approach. We follow the former strategy. The factors $m(\cdot) = (m_0, m_1, \dots, m_L)^\top$ are approximated by $A\psi$, with a coefficient matrix $A = (a_{l,k}) \in \mathbb{R}^{(L+1)K}$ and a vector of selected functions $\psi(\cdot) = (\psi_1, \dots, \psi_K)^\top$. K denotes the number of knots and it can be seen as a bandwidth parameter. This allows us to rewrite a part of model (2.2)

$$Z_t^\top m(X_{t,j}) = \sum_{l=0}^L Z_{t,l} m_l(X_{t,j}) = \sum_{l=0}^L Z_{t,l} \sum_{k=1}^K a_{l,k} \psi_k(X_{t,j}) = Z_t^\top A \psi(X_{t,j}). \quad (2.3)$$

The coefficient matrix A and time series of factor loadings Z_t is estimated using least squares. The estimated matrix \hat{A} and factor loadings $\hat{Z}_t = (1_T, \hat{Z}_{t,1}, \dots, \hat{Z}_{t,L})^\top$ minimize the sum of squared residuals, $S(A, Z_t)$

$$(\hat{Z}_t, \hat{A}) = \arg \min_{Z_t, A} S(A, Z_t) \quad (2.4)$$

$$= \arg \min_{Z_t, A} \sum_{t=1}^T \sum_{j=1}^J \left\{ Y_{t,j} - Z_t^\top A \psi(X_{t,j}) \right\}^2. \quad (2.5)$$

Newton-Raphson algorithm is used to find a solution of the minimization in (2.5). This algorithm converges to a solution at a geometric rate under some weak conditions on the initial choice $\{vec(A)^{(0)}, Z_t^{(0)}\}$, see, e.g. Park et al. [2009]. Furthermore, they prove that the differences between the estimated loadings \hat{Z}_t and the true loadings Z_t are asymptotically negligible. It is therefore justified to model the estimated factor loadings and consequently the object's dynamics by a parametric multivariate time series. In our work we consider a vector error correction (VEC) specification.

2.1.3 Application Details

The number of time-invariant factors L and the number of knots K is performed by evaluating the proportion of explained variance (EV)

$$EV(L) = 1 - RV(L) = 1 - \frac{\sum_{t=1}^T \sum_{j=1}^J \left\{ Y_{t,j} - \sum_{l=0}^L \hat{Z}_{t,l} \hat{m}_l(X_{t,j}) \right\}^2}{\sum_{t=1}^T \sum_{j=1}^J \{Y_{t,j} - \bar{Y}\}^2}. \quad (2.6)$$

We choose linearly spaced knots. A starting point is determined by the minimal value of the explanatory variable (corrected by -5%), whereas the end point corresponds to the maximal value (corrected by 5%). Results of a sensitivity results are quite stable regarding the choice of grid points. Due to the use of tensor B-spline functions for bid and ask curves (which are monotone in the price), our estimated first factor \hat{m}_1 and the estimated quantities $\hat{Y}_{t,j}$ are adjusted for extreme price levels. For the bid side we keep

2 Theoretical Modelling

constant the first (lowest) ten level values, and analogously, for the ask side we fix the last (highest) ten level values. Finally, the model's goodness-of-fit is evaluated using the root mean squared error (RMSE) criterion,

$$RMSE = \sqrt{\frac{1}{TJ} \sum_{t=1}^T \sum_{j=1}^J \{Y_{t,j} - \sum_{l=0}^L \hat{Z}_{t,l} \hat{m}_l(X_{t,j})\}^2}. \quad (2.7)$$

In summary, the DSFM framework allows us to simultaneously model the order book's shape and its dynamics. It will help us to discuss the spatial and temporal dependencies (structures) between liquidity supply and various economic variables, such as, e.g., the 'relative price levels', trading volume (liquidity demand), mid-quote returns, realized volatility and the best 'bid' and 'ask' prices (with corresponding returns and the bid-ask spread).

2.2 Multiplicative Error Model (MEM)

While the DSFM framework successfully deals with high-dimensional data structures, a Multiplicative Error Model (MEM), as discussed by Engle [2002], is widely used in univariate financial time series analysis like, e.g., in modelling trading volumes, durations, bid-ask spreads, price volatilities, market depth or trading costs. A MEM specification assumes, contrary to empirical evidence, time-invariant parameter structures over long estimation windows. In *localizing MEM* we therefore aim to find the longest estimation window over which one can safely apply a parametric MEM structure.

2.2.1 Model Structure

The model structure is based on the idea of an autoregressive conditional heteroscedasticity (ARCH) specification introduced by Engle [1982]. The framework is essentially used for modelling the temporal dependence and clustering effects in financial data. In high-frequency finance, Engle and Russell [1998] were first to apply the MEM to trade durations. Since then, MEM literature grew, see, e.g., Hautsch [2012] for a comprehensive overview.

Denote by $y = \{y_i\}_{i=1}^n$ a positive valued (financial) process. The data is modelled as a product of its conditional mean μ_i and a positive valued (unit mean) error term ε_i

$$y_i = \mu_i \varepsilon_i, \quad \mathbb{E}[\varepsilon_i | \mathcal{F}_{i-1}] = 1, \quad (2.8)$$

given the information set \mathcal{F}_i up to observation i . The conditional mean process follows an ARMA-type specification

$$\mu_i = \mu_i(\theta) = \omega + \sum_{j=1}^p \alpha_j y_{i-j} + \sum_{j=1}^q \beta_j \mu_{i-j}, \quad (2.9)$$

with parameters ω , $\alpha = (\alpha_1, \dots, \alpha_p)^\top$ and $\beta = (\beta_1, \dots, \beta_q)^\top$.

In modelling squared (de-measured) log returns the model structure resembles the conditional variance equation of a GARCH(p, q) model. In the context of modelling duration data, econometric literature refers to it as an autoregressive conditional duration (ACD) model. Here both labels are used as synonyms.

2.2.2 Distributional Assumptions

We choose the (standard) exponential and the Weibull distribution for modelling of the error term ε_i , see, e.g. Engle and Russell [1998]. A quasi maximum likelihood approach results in consistent MEM parameter estimation even in the case of distributional misspecification. Denote by $I = [i_0 - n, i_0]$ a (right-end) fixed interval of $(n + 1)$ observations at observation i_0 . We consider the following ACD models:

- (i) Exponential-ACD model (EACD) - $\varepsilon_i \sim \text{Exp}(1)$, $\theta_E = (\omega, \alpha^\top, \beta^\top)^\top$, with (quasi) log likelihood function over $I = [i_0 - n, i_0]$ given i_0 ,

$$L_I(y; \theta_E) = \sum_{i=\max(p,q)+1}^n \left(-\log \mu_i - \frac{y_i}{\mu_i} \right) \mathbf{I}(i \in I); \quad (2.10)$$

- (ii) Weibull-ACD model (WACD) - $\varepsilon_i \sim \mathcal{G}(s, 1)$, $\theta_W = (\omega, \alpha^\top, \beta^\top, s)^\top$, with (quasi) log likelihood function over $I = [i_0 - n, i_0]$ given i_0 ,

$$L_I(y; \theta_W) = \sum_{i=\max(p,q)+1}^n \left[\log \frac{s}{y_i} + s \log \frac{\Gamma(1 + 1/s) y_i}{\mu_i} - \left\{ \frac{\Gamma(1 + 1/s) y_i}{\mu_i} \right\}^s \right] \cdot \mathbf{I}(i \in I). \quad (2.11)$$

The quasi-maximum likelihood estimates (QMLEs) of θ_E and θ_W over the data interval I are given by

$$\tilde{\theta}_I = \arg \max_{\theta \in \Theta} L_I(y; \theta). \quad (2.12)$$

2.2.3 Estimation Quality

We assess the quality of the QMLE $\tilde{\theta}_I$ of the *true* parameter vector θ^* by the Kullback-Leibler divergence. For a fixed estimation interval I consider the (positive) difference $L_I(\tilde{\theta}_I) - L_I(\theta^*)$. The expressions for the log likelihood expressions for the EACD and WACD models given by (2.10) and (2.11), respectively. We define the loss function as the r -th power of that difference, i.e. $L_I(\tilde{\theta}_I, \theta^*) \stackrel{\text{def}}{=} |L_I(\tilde{\theta}_I) - L_I(\theta^*)|^r$.

For any risk power $r > 0$, there exists a constant, the so-called (parametric) risk bound $\mathcal{R}_r(\theta^*)$, such that it bounds the expected loss function

$$\mathbb{E}_{\theta^*} |L_I(\tilde{\theta}_I, \theta^*)|^r \leq \mathcal{R}_r(\theta^*), \quad (2.13)$$

see, e.g., Spokoiny [2009] and Čížek et al. [2009].

Recall that the idea behind *localizing MEM* is to find a data interval, i.e. the so-called *interval of homogeneity*, over which one can safely use the parametric MEM structure with constant parameters. The power r can be seen in this context as a steering parameter. Higher (lower) values of r lead to, ceteris paribus, longer (shorter) *intervals of homogeneity*. In this work we consider two scenarios, namely the 'modest risk case' ($r = 0.5$) and a 'conservative risk case' ($r = 1$). Practically, the 'modest risk case' would result, contrary to the 'conservative risk case', with shorter *intervals of homogeneity*, more variable parameters and essentially a lower modelling bias. We therefore suspect that the 'modest risk case' may outperform the 'conservative risk case' in short-term forecasting.

2.3 Local Parametric Approach (LPA)

A *local parametric approach* (LPA) requires that a time series can be locally, i.e., over short time periods, approximated by a parametric model. The idea behind the LPA is to statistically find the longest interval over which parameter homogeneity cannot be rejected. We consequently assume that high-frequency data locally follow the MEM structure presented in Section 2.2. This assumption is quite realistic in practical applications.

The LPA was proposed by Spokoyny [1998] and it has been gradually introduced into the econometrics literature. It was successfully used for daily volatility modelling, see, e.g., Mercurio and Spokoyny [2004] (local constant volatility), Čížek et al. [2009] (GARCH models), Chen et al. [2010] (realized volatility). We apply the LPA to high-frequency (transaction) data with the goal to understand the dynamics of financial processes when all single events are recorded.

2.3.1 Statistical Framework

Econometric literature suggests that local modelling outperforms global (parametric) modelling. Although long estimation windows reduce the parameter variability, they considerably enlarge the modelling bias. This effect is even more pronounced in high-frequency data analysis. By striking a balance between parameter variability and the modelling bias, the methodology presented below yields an *interval of homogeneity*. One can safely use this interval in modelling of transaction data.

We measure the theoretical differences between the underlying 'true' process μ_i and the parametric model $\mu_i(\theta)$ in (2.9) by the entity $\Delta_{I_k}(\theta) = \sum_{i \in I_k} \mathcal{K}\{\mu_i, \mu_i(\theta)\}$ over a given data interval I_k , where $\mathcal{K}(\cdot)$ denotes the Kullback-Leibler divergence. Denote by $\Delta \geq 0$ the *small modelling bias* (SMB), i.e. a constant that bounds the expected value of the modelling differences

$$\mathbb{E}[\Delta_{I_k}(\theta)] \leq \Delta, \quad (2.14)$$

for a fixed interval I_k and for some $\theta \in \Theta$.

2.3 Local Parametric Approach (LPA)

The SMB condition implies that in our QML estimation framework with loss function $L_I(\tilde{\theta}_I, \theta^*)$ over a fixed data interval I

$$\mathbb{E} \left[\log \left\{ 1 + \left| L_I(\tilde{\theta}_I, \theta^*) \right|^r / \mathcal{R}_r(\theta^*) \right\} \right] \leq 1 + \Delta, \quad (2.15)$$

where $\mathcal{R}_r(\theta^*)$ denotes the (parametric) risk bound in (2.13).

Consider $(K + 1)$ nested intervals (with right-end point i_0 fixed) $I_k = [i_0 - n_k, i_0]$ of length n_k and $I_0 \subset I_1 \subset \dots \subset I_K$. The 'oracle' (i.e., theoretically optimal) selected I_{k^*} is the largest interval for which the SMB condition holds $\mathbb{E} [\Delta_{I_{k^*}}(\theta)] \leq \Delta$.

In practice the entity Δ_{I_k} is unknown and therefore we aim to mimic the oracle choice using a sequential testing procedure, see Section 2.3.2. The resulting *interval of homogeneity* I_k is then used for defining the adaptive estimate. An important property of the adaptive estimation is that the involved estimation errors during steps $k \leq k^*$ are not larger than those induced by QML estimation based on k^* (stability condition), see, e.g., Čížek et al. [2009] and Spokoiny [2009]. The adaptive estimation therefore does not incur a larger estimation error compared to the situation where k^* is known, see (2.15).

The lengths of the underlying intervals are chosen to evolve on a geometric grid with initial length n_0 and a multiplier $c > 1$, $n_k = \lceil n_0 c^k \rceil$. We select $n_0 = 60$ observations (i.e., minutes) and consider two schemes with $c = 1.50$ and $c = 1.25$ and $K = 8$ and $K = 13$, respectively:

- (i) 9 estimation windows: $n_0 = 60$ min, $n_1 = 90$ min, \dots , $n_8 = 1800$ min (1 week), and
- (ii) 14 estimation windows: $n_0 = 60$ min, $n_1 = 75$ min, \dots , $n_{13} = 1800$ min (1 week).

The later scheme bears a slightly finer granulation than the first one.

2.3.2 Local Change Point (LCP) Detection Test

The selection of the *interval of homogeneity* is performed based on a sequential testing procedure. At each interval $I_k, k = 1, \dots, K$ one tests the null on parameter homogeneity against a change point alternative. The alternative hypothesis at step k states that there exists a change point at unknown location τ within I_k .

The test statistic of the test is given by

$$T_{I_k, J_k} = \sup_{\tau \in J_k} \left\{ L_{A_k, \tau}(\tilde{\theta}_{A_k, \tau}) + L_{B_k, \tau}(\tilde{\theta}_{B_k, \tau}) - L_{I_{k+1}}(\tilde{\theta}_{I_{k+1}}) \right\}, \quad (2.16)$$

where J_k and B_k denote intervals $J_k = I_k \setminus I_{k-1}$, $A_{k, \tau} = [i_0 - n_{k+1}, \tau]$ and $B_{k, \tau} = (\tau, i_0]$. This intervals use a part of the observations within I_{k+1} . The test statistic considers the maximum (supremum) of the corresponding likelihood ratio statistics over all (unknown) change points $\tau \in I_k$.

The testing procedure is graphically illustrated in Figure 2.1. Assume that for fixed i_0 parameter homogeneity in the interval I_{k-1} has been established. Homogeneity in interval I_k would mean that there is no break point τ in the interval $J_k = I_k \setminus I_{k-1}$. For

2 Theoretical Modelling

every τ , we compute the log likelihoods over the intervals $A_{k,\tau} = [i_0 - n_{k+1}, \tau]$ (colored in red) and $B_{k,\tau} = (\tau, i_0]$ (colored in blue). The test statistics (2.16) is then computed as the supremum of these likelihood values for each $\tau \in J_k$ relative to the log likelihood associated with the interval I_{k+1} .

For example, let us test for parameter homogeneity at step $k = 1$, i.e. consider the interval $I_1 = 75$ min. We search for possible change point(s) within $J_1 = I_1 \setminus I_0$, containing observations from y_{i_0-75} up to y_{i_0-60} . Sum the log likelihood values fitted over $A_{1,\tau}$ and $B_{1,\tau}$ and subtract the likelihood over I_2 . The test statistic (2.16) corresponds to the largest obtained likelihood ratio.




Figure 2.1: Graphical illustration of sequential testing for parameter homogeneity in interval I_k with length $n_k = |I_k|$ ending at fixed time point i_0 . Suppose we have not rejected homogeneity in interval I_{k-1} , we search within the interval $J_k = I_k \setminus I_{k-1}$ for a possible change point τ . The red interval marks $A_{k,\tau}$ and the blue interval marks $B_{k,\tau}$ (blue) splitting the interval I_{k+1} into two parts depending upon the position of the unknown change point τ .

2.3.3 Adaptive Estimation

At every step k we search for the longest *interval of homogeneity* $I_{\hat{k}}$ for which the null hypothesis (parameter homogeneity) is still not rejected. This is accomplished by comparison of the test statistic (2.16) and the corresponding (simulated) critical value. The construction of critical values is explained in the following Section 2.3.4. Then the *adaptive estimate* $\hat{\theta}$ is defined as the QMLE at the *interval of homogeneity*

$$\hat{\theta} = \tilde{\theta}_{\hat{k}}. \quad (2.17)$$

Note that if the null hypothesis is already rejected at the first step of the LCP detection test from Section 2.3.2, then $\hat{\theta}$ equals to the QMLE at the shortest interval. In our case, the shortest interval (by assumption homogeneous) is $I_0 = 60$ min. If no break point can (still) be detected within I_K , then $\hat{\theta}$ equals the QMLE over the longest window. The longest interval in our case contains 1800 observations: $I_K = 1800$ min = 1 week.

2.3.4 Critical Values

The test statistic for parameter homogeneity is a nonlinear function (supremum) of the likelihood and therefore it makes the derivation of distributional properties difficult. We therefore simulate the critical values under the null of parameter homogeneity over the interval sequence $\{I_k\}_{k=1}^K$. The idea is to control the loss due to adaptive estimation since step $\hat{k} < K$ is potentially selected. Under the null hypothesis of parameter homogeneity, the correct choice is the largest considered interval I_K . At each step $k = 1, \dots, K$ we impose the following ('propagation') condition on the corresponding k -th step critical value

$$\mathbb{E}_{\theta^*} \left| L_{I_k}(\tilde{\theta}_{I_k}) - L_{I_k}(\hat{\theta}_{I_k}) \right|^r \leq \rho_k \mathcal{R}_r(\theta^*), \quad (2.18)$$

with $\rho_k = \rho k/K$ for given significance level ρ and the risk bound $\mathcal{R}_r(\theta^*)$, recall (2.13).

So selected critical values ensure that the loss associated with 'false alarm' (i.e., selecting $k < K$) is at most a ρ -fraction of the parametric risk bound of the 'oracle' estimate $\tilde{\theta}_K$. For $r \rightarrow 0$, ρ can be interpreted as the false alarm probability.

Since the test statistic (2.16) is not studentized, one needs to check the critical values for different null parameters θ^* . The simulation of critical values depends furthermore upon the involved parameters (r, ρ, K), as well as the modelling framework employed (EACD or WACD). The parametric risk bound $\mathcal{R}_r(\theta^*)$ given in (2.13) has to be simulated as well.

We select nine parameter constellations based on a local parameter dynamics study, see Section 4.2.1 and Table 4.11. Due to the (high-frequency) nature of financial data, we rather follow a data-driven selection technique than a grid-search based approach. Here we consider two risk levels ($r = 0.5$ and $r = 1$), two interval granulation schemes ($K = 8$ and $K = 13$) and two significance levels ($\rho = 0.25$ and $\rho = 0.50$).

The resulting critical values satisfying (2.18) for the nine possibilities of 'true' parameter constellations from Table 4.11 of the EACD(1,1) model for interval schemes with $K = 8$ and $K = 13$ are displayed in Figures 2.2 and 2.3, respectively. Similarly, the critical values of the WACD(1,1) model for the two considered interval schemes ($K = 8$ and $K = 13$) are shown in Figures 2.4 and 2.5, respectively. We distinguish at each scenario between a 'moderate risk case' ($r = 0.5$) and a 'conservative risk case' ($r = 1$). We set $\rho = 0.25$ since the results with $\rho = 0.50$ are quite similar.

The critical values are almost invariable with respect to θ^* across the nine scenarios, recall the parameter constellations from Table 4.11. The largest difference between all cases appears for interval lengths up to 90 minutes, i.e. until step $k = 3$. We observe that the critical values are robust across the range of parameters, particularly for the underlying risk cases ($r = 0.5$ and $r = 1$), interval selection schemes and models under consideration (EACD and WACD).




Figure 2.2: Simulated critical values of an EACD(1,1) model for the 'moderate risk case' ($r = 0.5$, upper panel) and the 'conservative risk case' ($r = 1$, lower panel), $\rho = 0.25$, $K = 8$ and chosen parameters constellations according to Table 4.11. The low (blue), middle (green) and upper (red) curves are associated with the corresponding ratio levels $\tilde{\beta}/(\tilde{\alpha} + \tilde{\beta})$.




Figure 2.3: Simulated critical values of an EACD(1,1) model for the 'moderate risk case' ($r = 0.5$, upper panel) and the 'conservative risk case' ($r = 1$, lower panel), $\rho = 0.25$, $K = 13$ and chosen parameters constellations according to Table 4.11. The low (blue), middle (green) and upper (red) curves are associated with the corresponding ratio levels $\tilde{\beta}/(\tilde{\alpha} + \tilde{\beta})$.




Figure 2.4: Simulated critical values of a WACD(1, 1) model for the 'moderate risk case' ($r = 0.5$, upper panel) and the 'conservative risk case' ($r = 1$, lower panel), $\rho = 0.25$, $K = 8$ and chosen parameters constellations according to Table 4.11. The low (blue), middle (green) and upper (red) curves are associated with the corresponding ratio levels $\tilde{\beta}/(\tilde{\alpha} + \tilde{\beta})$.




Figure 2.5: Simulated critical values of a WACD(1, 1) model for the 'moderate risk case' ($r = 0.5$, upper panel) and the 'conservative risk case' ($r = 1$, lower panel), $\rho = 0.25$, $K = 13$ and chosen parameters constellations according to Table 4.11. The low (blue), middle (green) and upper (red) curves are associated with the corresponding ratio levels $\tilde{\beta}/(\tilde{\alpha} + \tilde{\beta})$.

2.4 Pricing Kernel under State-Dependent Utility

In market microstructure literature (particularly in asset pricing), a promising technique for preference modelling works with the pricing kernel under state-dependent utility, see, e.g., Grith et al. [2011]. The presented methodology considers therefore a more realistic economic framework.

2.4.1 Economic Setup

We consider a representative agent with an exogenous income ω_t at time t . The entire income is used to finance current consumption c_t and to purchase a financial portfolio consisting of k assets with prices $S_t = (S_{1,t}, \dots, S_{k,t})^\top$

$$\omega_t = c_t + q_t^\top S_t, \quad (2.19)$$

where $q_t = (q_{1,t}, \dots, q_{k,t})^\top$ denotes the vector of asset holdings. The agent's current consumption therefore equals the difference between the income and the financial wealth, i.e. $c_t = \omega_t - q_t^\top S_t$. Given the current portfolio choice q_t , the next period consumption includes the future income and all asset payoffs, that is $c_{t+1} = \omega_{t+1} + q_t^\top S_{t+1}$, where S_{t+1} denotes the asset prices including all corresponding payoffs at time $t+1$.

A representative agent maximizes the following expected time separable and state-dependent utility

$$u(c_t, c_{t+1}) = u(c_t) + \beta_1 \mathbb{E}_t[u(c_{t+1})] \mathbf{I}\{c_t \in [0, x)\} + \beta_2 \mathbb{E}_t[u(c_{t+1})] \mathbf{I}\{c_t \in [x, \infty)\} \quad (2.20)$$

with given reference point x at which the agent may switch between two preference specifications given potentially different (non-negative) impatience parameters β_1 and β_2 , see, e.g., Grith et al. [2011]. We denote the conditional expectation operator by $\mathbb{E}_t[\bullet] = \mathbb{E}[\bullet | \mathcal{F}_t]$, given the information set \mathcal{F}_t up to time t .

It turns out that the maximization of the expected utility (2.20) over the consumption stream is equivalent to the choice of the optimal portfolio holding q_t

$$\begin{aligned} \max_{c_t, c_{t+1}} u(c_t, c_{t+1}) &= \max_{q_t} [u(\omega_t - S_t^\top q_t) + \\ &\quad + \beta_1 \mathbb{E}_t[u(\omega_{t+1} + q_t^\top S_{t+1})] \mathbf{I}\{(\omega_{t+1} + q_t^\top S_{t+1}) \in [0, x)\} + \\ &\quad + \beta_2 \mathbb{E}_t[u(\omega_{t+1} + q_t^\top S_{t+1})] \mathbf{I}\{(\omega_{t+1} + q_t^\top S_{t+1}) \in [x, \infty)\}]. \end{aligned} \quad (2.21)$$

The first order conditions for (expected) utility maximization therefore imply the following fundamental asset pricing formula

$$S_t = \mathbb{E}_t \left[\left\{ \beta_1 \frac{u'(c_{t+1})}{u'(c_t)} \mathbf{I}\{c_t \in [0, x)\} + \beta_2 \frac{u'(c_{t+1})}{u'(c_t)} \mathbf{I}\{c_t \in [x, \infty)\} \right\} S_{t+1} \right]. \quad (2.22)$$

2.4.2 Pricing Kernel

The multiplicative term related to S_{t+1} in (2.22) denotes the consumption based intertemporal pricing kernel or the stochastic discount factor that measures the intertemporal rate of consumption substitution. Under logarithmic utility, i.e. with $u(c_t) = \log c_t$ and $u(c_{t+1}) = \log c_{t+1}$, the intertemporal rate of consumption substitution becomes an inverse function of consumption growth

$$\frac{u'(c_{t+1})}{u'(c_t)} = \left(\frac{c_{t+1}}{c_t} \right)^{-1}. \quad (2.23)$$

In practice one observes poor consumption data quality and therefore the consumption based intertemporal pricing kernel is related to data that approximate or influence agent's consumption, see, e.g. Cochrane [2001]. Consumption growth is therefore in practice approximated by the simple market gross return $r_{m,t+1} = S_{m,t+1}/S_{m,t}$, where the observed stock market index value at time t is denoted by $S_{m,t}$. Finally, the state-dependent pricing kernel (for fixed reference point x) is given by

$$\mathcal{K}_\theta(r_{m,t+1}) = \beta_1 r_{m,t+1}^{-1} \mathbf{I}\{r_{m,t+1} \in [0, x)\} + \beta_2 r_{m,t+1}^{-1} \mathbf{I}\{r_{m,t+1} \in [x, \infty)\}, \quad (2.24)$$

with parameter vector $\theta = (\beta_1, \beta_2)^\top$.

2.4.3 Moment Conditions

The asset pricing equation (2.22) reads as

$$S_t = \mathbb{E}_t [\mathcal{K}_\theta(r_{m,t+1}) S_{t+1}], \quad (2.25)$$

or expressed in return terms it can be interpreted as the expectation of k (conditional) moment conditions

$$\mathbb{E}_t [\mathcal{K}_\theta(r_{m,t+1}) R_{t+1} - 1_k] = 0_k, \quad (2.26)$$

where $R_{t+1} = (S_{1,t+1}/S_{1,t}, \dots, S_{k,t+1}/S_{k,t})^\top$ denotes the simple gross asset return vector. The k -dimensional vectors of ones and zeros are denoted by 1_k and 0_k , respectively. An optimal asset allocation therefore implies that for each asset, the expected value of the cross-product between the pricing kernel and the simple gross asset return equals one.

2.5 Generalized Method of Moments

In estimation of the pricing kernel given in (2.25) we utilize the Generalized Method of Moments (GMM) approach, as proposed by Hansen [1982]. Based on the k (conditional) moment conditions from the asset pricing equation (2.26)

$$g(\theta) \stackrel{\text{def}}{=} \mathcal{K}_\theta(r_{m,t+1}) R_{t+1} - 1_k, \quad \mathbb{E}_t [g(\theta)] = 0_k, \quad (2.27)$$

2 Theoretical Modelling

we define the sample moment function as

$$g_n(\theta) \stackrel{\text{def}}{=} n^{-1} \sum_{t=0}^{n-1} \{\mathcal{K}_\theta(r_{m,t+1}) R_{t+1} - 1_k\}, \quad (2.28)$$

over the data sample of size n .

2.5.1 Parameter Estimation

The parameter vector θ is estimated by two techniques:

- (i) Iterated GMM estimation based on a two-step efficient GMM estimation procedure by Hansen and Singleton [1982], see, e.g., Ferson and Foerster [1994] - in the first step, using a feasible weighting matrix (e.g. the identity matrix of order k) we obtain the estimate

$$\tilde{\theta}_n \stackrel{\text{def}}{=} \arg \min_{\theta} \{g_n^\top(\theta) g_n(\theta)\}. \quad (2.29)$$

The resulting consistent optimal weighting matrix is given by

$$\widetilde{W}_n = n^{-1} \sum_{t=0}^{n-1} g(\tilde{\theta}_n) g(\tilde{\theta}_n)^\top. \quad (2.30)$$

Secondly, based on the optimal weighting matrix \widetilde{W}_n the feasible efficient GMM estimate solves

$$\widehat{\theta}_n \stackrel{\text{def}}{=} \arg \min_{\theta} \{g_n^\top(\theta) \widetilde{W}_n^{-1} g_n(\theta)\} \quad (2.31)$$

and leads to the following consistent optimal weighting matrix

$$\widehat{W}_n = n^{-1} \sum_{t=0}^{n-1} g(\widehat{\theta}_n) g(\widehat{\theta}_n)^\top. \quad (2.32)$$

The second step is iterated until parameter convergence. As a rule of thumb, we stop when the estimated parameters do not differ on the fourth digit.

- (ii) GMM estimation with Hansen-Jagannathan (HJ) weighting matrix

$$\widetilde{W}_n = n^{-1} \sum_{t=0}^{n-1} R_t R_t^\top, \quad (2.33)$$

see, e.g., Jagannathan and Wang [1996] and Hansen and Jagannathan [1997]. The estimate follows directly from (2.31). This technique may provide better finite sample properties of the GMM estimate as the matrix is not a function of the model parameters, i.e. it may lead to more robust results, see, e.g., Cochrane [2001].

2.5.2 Hypothesis Testing

The GMM modeling framework allows us to test for the non-monotonicity of the pricing kernel. We employ the so-called “ D -test”, as proposed by Newey and West [1987]. The corresponding test statistic is given by

$$D = ng_n^\top(\tilde{\theta}_n) \widetilde{W}_n^{-1} g_n(\tilde{\theta}_n) - ng_n^\top(\check{\theta}_n) \check{W}_n^{-1} g_n(\check{\theta}_n) \xrightarrow{\mathcal{L}} \chi_j^2, \quad (2.34)$$

with j denoting the number of imposed parameter restrictions. The estimated parameter vectors using two methods are denoted by $\tilde{\theta}_n$ and $\check{\theta}_n$. The associated weighting matrices are labeled by \widetilde{W} and \check{W} , respectively.

3 Data

It is the theory that decides what we can observe.
Albert Einstein

3.1 Stock Markets around the Globe

In estimating pricing kernels we select six from the worldwide largest ten stock markets, namely the Australian Securities Exchange (Australia - AUS), Deutsche Börse (Germany - GER), the Tokyo Stock Exchange (Japan - JPN), the SIX Swiss Exchange (Switzerland - SUI), the London Stock Exchange (United Kingdom - UK) and the New York Stock Exchange (United States - US). At each exchange, the data collected from Datastream and EcoWin include daily stock index values, interest rates, as well as closing stock prices of the largest 20 companies by market capitalization augmented on 31 May 2012, whose stocks were continuously traded during the sample period from 1 January 1990 until 31 May 2012.

In modelling and forecasting liquidity supply we focus on limit order book data for four companies at the Australian Securities Exchange from 8 July to 16 August 2002, see Section 3.2. The local multiplicative error model is applied to transaction data of five large NASDAQ stocks in the period from 2 January to 31 December 2008, see Section 3.3.

3.1.1 Descriptive Statistics

Denote the index value at given stock market at time t by $S_{m,t}$ and the closing stock prices of 20 blue chips by $S_t = (S_{1,t}, \dots, S_{20,t})^\top$. The variables of interest in estimating the pricing kernel are the monthly overlapping simple market gross return $R_{m,t+1} = S_{m,t+1}/S_{m,t-20}$, as well as the monthly overlapping simple stock gross returns $R_{t+1} = (S_{1,t+1}/S_{1,t-20}, \dots, S_{20,t+1}/S_{20,t-20})^\top$. The time series of the market returns are displayed in Figure 3.1 and descriptive statistics thereof are summarized in Table 3.1.

We find left-skewed and leptokurtic distributions, thus the distributions have more probability mass around the center and in the tails than the normal distribution. In financial time series kurtosis is typically larger than 3 due to the frequent appearance of outliers. The average index return is close to one in all cases with a dispersion of about 4-5%. High market volatility is observed during the stock market distress period from 1998-2003 and during the present financial crisis which started in fall 2008.

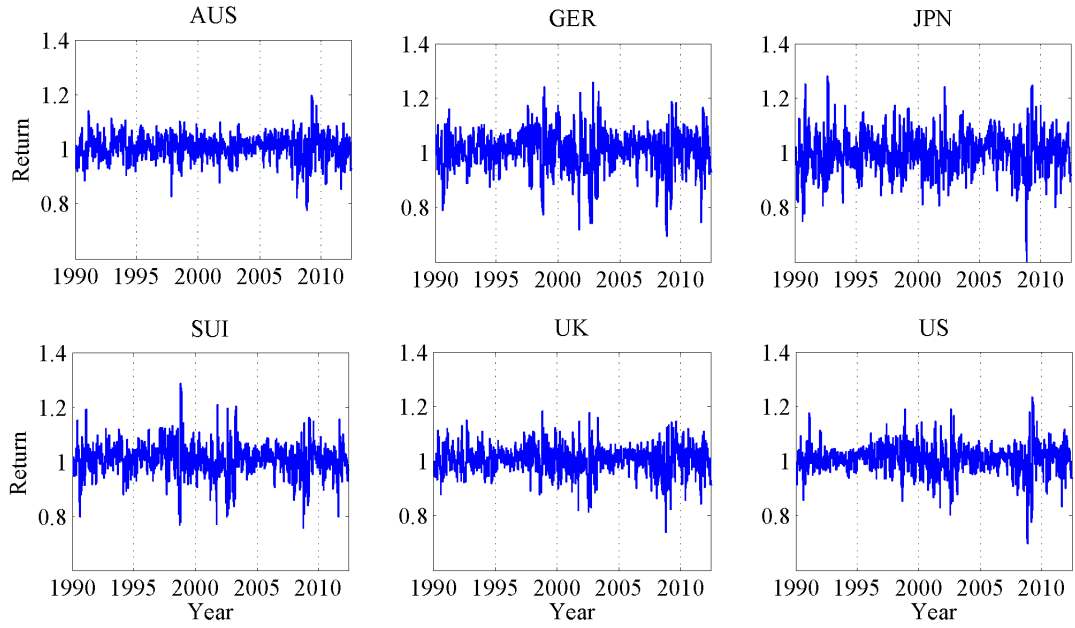


Figure 3.1: Market returns for selected stock market indices on major stock markets from 1 January 1990 until 31 May 2012 (5827 observations per series).

3.1.2 Portfolio Selection

At each stock exchange we select among the most frequently traded stocks a portfolio of 20 largest blue chips. The selected market portfolios share similar worldwide performance (results not reported in the paper) and focus country-wise on the most dominating industries. Our empirical results are therefore comparable across countries. In order to discuss robustness of our results, we split the initial pool of 20 stocks per market into six equally-weighted (sub-)portfolios per criteria selected. Current finance literature suggest to group stocks firstly according to their book/market (B/M) value into three groups (high - 30%, medium - 40% and low - 30% of companies), see, e.g., Fama and French

Country	Index	Mean	St. deviation	Skewness	Kurtosis
Australia	S&P/ASX 200	1.004	0.041	-0.48	4.54
Germany	DAX 30	1.007	0.062	-0.65	4.84
Japan	NIKKEI 225	0.997	0.064	-0.14	4.44
Switzerland	SMI	1.006	0.051	-0.41	5.20
United Kingdom	FTSE 100	1.004	0.045	-0.53	4.88
United States	S&P 500	1.006	0.046	-0.63	6.48

Table 3.1: Descriptive statistics of monthly returns of the selected market indices from 1 January 1990 until 31 May 2012 (5827 observations per series).

3.1 Stock Markets around the Globe

[1992]. Secondly, the stocks are divided into two groups according to one of the following criteria:

- (i) *Size* - based on the market capitalization we distinguish between large and small companies, see, e.g., studies that employ the Fama-French three factor model by Fama and French [1993].
- (ii) *Momentum* - stocks with the highest (lowest) past 12-month return are expected to yield high (low) return in the following month, as firstly shown by Jegadeesh [1990]. We consider therefore the average 12-month overlapping return to categorize companies.
- (iii) *Beta* - in the capital asset pricing model (CAPM), developed independently by various authors in the sixties, the “beta” factor of stock i is given by the ratio $\text{Cov}(R_{i,t}, R_{m,t}) / \text{Var}(R_{m,t})$, $t = 1, \dots, n$. Stocks with larger beta tend to be more volatile and therefore riskier.

The performance of all considered portfolios is illustrated in Tables 3.2, 3.3 and 3.4 for companies with high, medium and low book-to-market ratio, respectively.

Country	Size		Momentum		Beta	
	High	Low	High	Low	High	Low
Australia	1.0101	1.0071	1.0114	1.0058	1.0079	1.0093
Germany	1.0049	1.0028	1.0051	1.0026	1.0026	1.0051
Japan	0.9999	1.0020	1.0040	0.9979	1.0017	1.0002
Switzerland	1.0059	1.0065	1.0080	1.0043	1.0059	1.0065
United Kingdom	1.0101	1.0091	1.0124	1.0067	1.0095	1.0096
United States	1.0070	1.0110	1.0110	1.0070	1.0103	1.0078

Table 3.2: Average return of selected portfolios for companies with a high book-to-market ratio from 1 January 1990 until 31 May 2012.

Country	Size		Momentum		Beta	
	High	Low	High	Low	High	Low
Australia	1.0090	1.0107	1.0116	1.0081	1.0090	1.0106
Germany	1.0093	1.0064	1.0104	1.0053	1.0089	1.0069
Japan	1.0045	1.0036	1.0054	1.0027	1.0033	1.0048
Switzerland	1.0092	1.0083	1.0110	1.0065	1.0069	1.0106
United Kingdom	1.0095	1.0117	1.0135	1.0078	1.0128	1.0085
United States	1.0095	1.0132	1.0154	1.0072	1.0149	1.0078

Table 3.3: Average return of selected portfolios for companies with a medium book-to-market ratio from 1 January 1990 until 31 May 2012.

Country	Size		Momentum		Beta	
	High	Low	High	Low	High	Low
Australia	1.0077	1.0089	1.0107	1.0058	1.0107	1.0058
Germany	1.0127	1.0112	1.0150	1.0088	1.0127	1.0112
Japan	1.0065	1.0026	1.0065	1.0026	1.0041	1.0050
Switzerland	1.0105	1.0111	1.0125	1.0091	1.0120	1.0097
United Kingdom	1.0087	1.0121	1.0150	1.0058	1.0113	1.0095
United States	1.0182	1.0093	1.0182	1.0093	1.0182	1.0093

Table 3.4: Average return of selected portfolios for companies with a low book-to-market ratio from 1 January 1990 until 31 May 2012.

Empirical evidence shows that portfolios of high momentum stocks achieve worldwide best results irrespectively upon the book-to-market ratio. Volatility of a stock (here measured by the beta factor) is negatively related to the book-to-market (B/M) ratio. This means that the best performance of high (low) B/M stocks is achieved for low (high) volatile stocks. The influence of company's size on the B/M performance varies over markets. We distinguish between countries with neutral (Germany), mixed (Switzerland), positive (Australia and UK) and negative effect (Japan and US). For example, on a German market the largest companies outperform small ones in any case. For countries with positive (negative) tendency, portfolios of large (small) companies are more desirable given high B/M value.

3.2 Trading at the Australian Stock Exchange (ASX)

The Australian Stock Exchange (ASX) is a continuous double auction electronic market. The continuous auction trading period is preceded and followed by a call auction. Normal trading takes place continuously on all stocks between 10:09 a.m. and 4:00 p.m. from Monday to Friday. Any buy (sell) order entered that has a price that is greater than (less than) or equal to existing queued buy (sell) orders, will be executed immediately. If an order cannot be executed completely, the remaining volume enters the queues as a limit order. Limit orders are queued in the buy and sell queues according to a strict price-time priority order. Orders can be entered, deleted and modified without restriction.

For order prices below 10 cents, the minimum tick size is 0.1 cents, for order prices above 10 cents and below 50 cents it is 0.5 cents, whereas for orders priced 50 cents and above it is 1 cent. There might be orders which are entered with an undisclosed or hidden volume if the total value of the order exceeds AUD 200,000. As this applies only to a small fraction of the posted volumes, we safely neglect the occurrence of hidden volume in our empirical study. For more details on the data, see Hall and Hautsch [2007] using the same data base as well as the official description of the trading rules of the Stock Exchange Automated Trading System (SEATS) on the ASX on www.asxonline.com.

We select four companies traded at the ASX covering the period from 8 July to 16 August 2002 (30 trading days), namely Broken Hill Proprietary Limited (BHP), National

3.2 Trading at the Australian Stock Exchange (ASX)

Australia Bank Limited (NAB), MIM and Woolworths (WOW). The number of market and limit orders in the period under review for the selected stocks is given in Table 3.5.

Orders	BHP	NAB	MIM	WOW
Market orders				
(i) buy	28,030	16,304	4,115	7,260
(ii) sell	16,755	15,142	2,789	6,464
Limit orders				
(i) buy (bid side)	50,012	28,850	9,551	13,234
- changed	8,009	7,561	1,637	3,203
- cancelled	5,202	4,725	2,044	1,951
(ii) sell (ask side)	32,053	25,953	6,474	11,318
- changed	6,891	6,261	1,862	3,164
- cancelled	4,692	3,863	1,178	1,554

Table 3.5: Total number of market and limit orders for selected stocks traded at the ASX from 8 July to 16 August 2002.

There are more buy orders than sell orders implying that the bid side of the limit order book was changing more frequently than the ask side. BHP and NAB are significantly more actively traded than MIM and WOW shares. Aggregated over all stocks, 20.08% (23.98%) of all bid (ask) limit orders have been changed (after posting), whereas 13.70% (14.89%) have been cancelled. For both traded as well as posted quantities we find that on average sell volumes are higher than buy volumes (not reported here). Liquidity variations on the bid side are again higher than that of the ask side. This finding might be explained by the fact that during the analyzed period the market generally went down creating more sell activities than buy activities.

The original dataset contains all limit order book records as well as the corresponding order curves represented by the underlying price-volume combinations. The latter is the particular object of interest for the remainder of the analysis. The underlying limit order book data contains identification attributes regarding $r = 1, \dots, R$ different orders as well as quantities demanded and offered for different price levels $j = 1, \dots, J$, at any time point $t = 1, \dots, T$. At any t , we observe $J = 101$ price levels on a fixed minimum tick size grid originating from the best bid and ask quote.

Since the order book dynamics are found to be very persistent, we choose a sampling frequency of five minutes without losing too much information on the liquidity supply. To remove effects due to market opening and closure, the first 15 minutes and last 5 minutes are discarded. At each trading day, starting at 10:15 and ending at 15:55, we select per stock 69 price-quantity vectors, in total $T = 2070$ vectors over the whole sample period. Denote by $\tilde{Y}_{t,j}^b$ and $\tilde{Y}_{t,j}^a$ the pending bid and ask volumes at bid and ask limit prices $\tilde{S}_{t,j}^b$ and $\tilde{S}_{t,j}^a$, respectively at time point t .

The best bid price at time t is defined as the highest buy price $\tilde{S}_{t,101}^b$, and similarly, the best ask price at t as the lowest sell price $\tilde{S}_{t,1}^a$. The corresponding quantities at

3 Data

best bid and ask prices are then $\tilde{Y}_{t,101}^b$ and $\tilde{Y}_{t,1}^a$, respectively, yielding the mid-quote price to be defined as $\tilde{S}_t^* = (\tilde{S}_{t,101}^b + \tilde{S}_{t,1}^a)/2$. The absolute price deviations from the best bid and ask price at level j and time t are given by $\check{S}_{t,j}^b = \tilde{S}_{t,j}^b - \tilde{S}_{t,101}^b$ and $\check{S}_{t,j}^a = \tilde{S}_{t,j}^a - \tilde{S}_{t,1}^a$, respectively and constitute a fixed price grid. To measure spreads between individual price levels in *relative* terms, i.e., in relation to the prevailing best bid and ask price, we define so-called 'relative price levels' as $S_{t,j}^b = \tilde{S}_{t,j}^b/\tilde{S}_{t,101}^b$ and $S_{t,j}^a = \tilde{S}_{t,j}^a/\tilde{S}_{t,1}^a$, respectively.

In order to investigate to which extent order book information might reveal information to predict high-frequency returns, we regress 1 min and 5 min mid-quote returns, respectively, on lagged order imbalances

$$\tilde{Y}_{t-1,j}^b / (\tilde{Y}_{t-1,j}^b + \tilde{Y}_{t-1,j}^a) \text{ and } \tilde{Y}_{t-1,j}^a / (\tilde{Y}_{t-1,j}^b + \tilde{Y}_{t-1,j}^a),$$

respectively, for $j = 1, \dots, 101$. Figure 3.2 shows the implied R^2 values in dependence of the number of included imbalance levels. It turns out that order book imbalances indeed reveal short-term predictability. Even levels far apart from the market have still distinct prediction power pushing the R^2 to values of approximately 10%. These findings show that the order book itself reveals predictive content for future price movements which could be exploited in trading strategies.




Figure 3.2: Coefficients of determination (R^2) implied by linear regression of 1 min (red) and 5 min (blue) mid-quote returns on lagged order imbalances for selected stocks traded at the ASX from 8 July to 16 August 2002 (30 trading days). The horizontal axis depicts the number of included imbalance levels.

In order to account for intra-day seasonality effects, we adjust the order volumes correspondingly. To avoid to seasonally adjust all individual volume series separately, we assume that the seasonality impact on quoted volumes at all levels is identical and is well captured by the seasonalities in market depth on the best bid and ask levels $\tilde{Y}_{t,101}^b$ and $\tilde{Y}_{t,1}^a$, respectively. Assuming a multiplicative impact of the seasonality factor, the seasonally adjusted quantities are computed for both sides of the market at price level

j , and time t as

$$Y_{t,j}^b = \frac{\tilde{Y}_{t,j}^b}{s_t^b} \quad (3.1)$$

$$Y_{t,j}^a = \frac{\tilde{Y}_{t,j}^a}{s_t^a}, \quad (3.2)$$

with s_t^b and s_t^a representing the seasonality components at time t for the bid and the ask side, respectively.

The non-stochastic seasonal trend factors s_t^b and s_t^a are specified parametrically using a flexible Fourier series approximation as proposed by Gallant [1981] and are given by

$$s_t^b = \delta^b \cdot \bar{t} + \sum_{m=1}^{M^b} \{ \delta_{c,m}^b \cos(\bar{t} \cdot 2\pi m) + \delta_{s,m}^b \sin(\bar{t} \cdot 2\pi m) \} \quad (3.3)$$

$$s_t^a = \delta^a \cdot \bar{t} + \sum_{m=1}^{M^a} \{ \delta_{c,m}^a \cos(\bar{t} \cdot 2\pi m) + \delta_{s,m}^a \sin(\bar{t} \cdot 2\pi m) \}. \quad (3.4)$$

Here δ^b , δ^a , $\delta_{c,m}^b$, $\delta_{c,m}^a$, $\delta_{s,m}^b$ and $\delta_{s,m}^a$ are coefficients to be estimated, and \bar{t} denotes a normalized time trend mapping the time of the day on a $(0, 1]$ intervals. The polynomial orders M^b and M^a are selected according to the Bayes information criterion (BIC). For all stocks we select $M^b = M^a = 1$, except for the bid side for BHP ($M^b = 2$). The resulting intra-day seasonality patterns for both sides of all order book markets are plotted in Figure 3.3.

Liquidity supply increases for all stocks before market closure. We attribute this finding to traders' pressure and willingness to close positions overnight. Posting aggressive limit orders on the best levels (or even within the spread) maximizes the execution probability and avoids crossing the spread. Moreover, weak evidence for a 'lunch time dip' is presented which, however, is only observed for the more liquid stocks (NAB and BHP). For the less liquid stocks the amount of posted volume nearly monotonically increases over the course of the day.

3.3 Trading at the NASDAQ Stock Market

We use transaction data of five large companies traded at NASDAQ: Apple Inc. (AAPL), Cisco Systems, Inc. (CSCO), Intel Corporation (INTC), Microsoft Corporation (MSFT) and Oracle Corporation (ORCL). These companies account for approximately one third of the market capitalization within the technology sector. Our variable of interest is the one-minute cumulative trading volume, reflecting high-frequency liquidity demand, covering the period from 2 January to 31 December 2008 (250 trading days with continuous trading activity). To remove effects due to market opening, the first 30 minutes of each trading session are discarded. Hence, at each trading day, we analyze data from 10:00 to 16:00. Descriptive statistics of daily and one-minute cumulated trading volume of the

3 Data

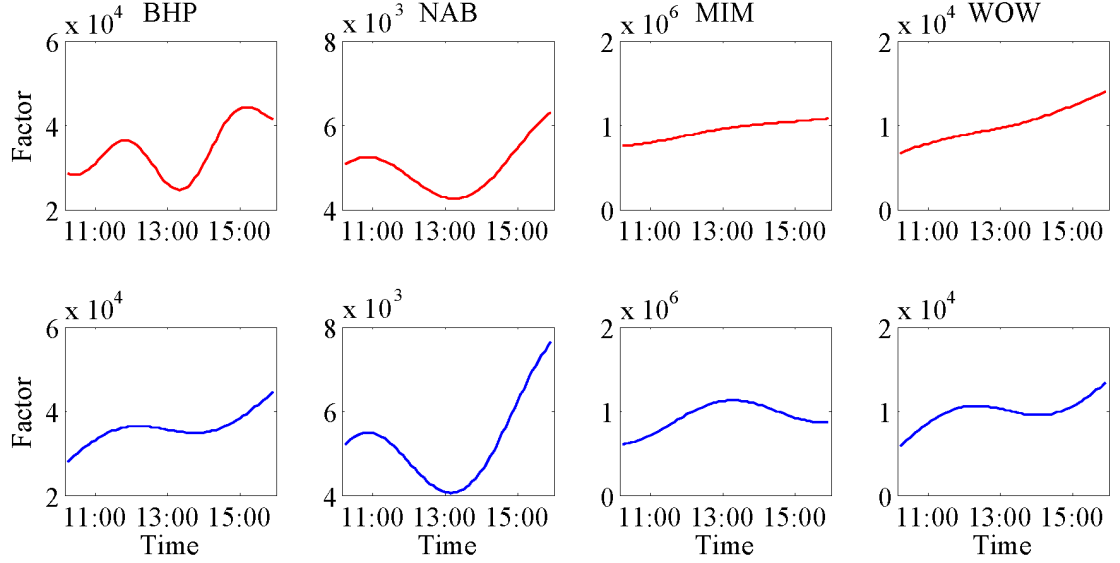


Figure 3.3: Estimated intra-day seasonality factors for quantities offered at best bid prices (red) and for quantities supplied at best ask prices (blue) across selected stocks traded at the ASX from 8 July to 16 August 2002 (30 trading days).

five analyzed stocks are shown in Table 3.6.

We find right-skewed distributions with higher dispersions on the high-frequency level than on the daily level. The Ljung-Box (LB) tests statistics indicate a strong serial dependence as the the null hypothesis of no autocorrelations (among the first 10 lags) is clearly rejected on any reasonable significance level.

Denote the one-minute cumulative trading volume by \check{y}_i . Assuming a multiplicative impact of intra-day periodicity effects, we compute seasonality adjusted volumes by

$$y_i = \check{y}_i s_i^{-1}, \quad (3.5)$$

with s_i representing the intraday periodicity component at time point i . Typically, seasonality components are assumed to be deterministic and thus constant over time. To capture slowly moving ('long-term') components, we estimate the periodicity effects on the basis of 30-days rolling windows, see, e.g., Engle and Rangel [2008]. Seasonality effects could be captured directly within the local adaptive framework presented below avoiding to fix the length of the rolling window on an ad hoc basis. As our focus is on (pure stochastic) short-term variations in parameters rather than on (more deterministic) periodicity effects, we decide to remove the former beforehand. This leaves us with non-homogeneity in processes which is not straightforwardly taken into account and allows us to evaluate the potential of a local adaptive approach even more convincingly. The intra-day component s_i is specified via a flexible Fourier series approximation as proposed by

3.3 Trading at the NASDAQ Stock Market

	AAPL	CSCO	INTC	MSFT	ORCL
Daily volume in million					
Minimum	8.7	12.8	12.5	15.3	8.2
25%-quantile	24.3	38.2	41.8	48.7	25.6
Median	30.6	47.7	54.9	64.7	33.3
75%-quantile	39.3	59.4	67.5	81.3	41.9
Maximum	100.4	177.3	227.8	204.8	88.4
Mean	33.4	50.9	58.3	68.7	35.0
Standard deviation	13.4	19.0	24.8	28.0	13.1
LB(10)	651.8	271.9	373.3	537.0	252.8
One-minute volume in 1000 shares					
Minimum	1.5	0.4	0.6	1.6	0.4
25%-quantile	47.3	58.7	63.6	78.6	35.9
Median	75.4	105.7	119.4	141.7	70.1
75%-quantile	118.5	180.8	208.9	242.1	124.4
Maximum	2484.8	3064.9	12231.4	7360.8	3558.2
Mean	92.9	141.4	162.0	190.8	97.1
Standard deviation	68.9	131.7	166.4	183.0	101.1
LB(10)	334076.1	164999.2	142128.8	197173.7	107629.6

Table 3.6: Descriptive statistics and Ljung-Box statistics (based on 10 lags) of daily and one-minute cumulated trading volumes of five large companies traded at NASDAQ between January 2 and December 31, 2008 (250 trading days, 90000 observations per stock).

Gallant [1981].

$$s_i = \delta \cdot \bar{i} + \sum_{m=1}^M \{ \delta_{c,m} \cos(\bar{i} \cdot 2\pi m) + \delta_{s,m} \sin(\bar{i} \cdot 2\pi m) \}. \quad (3.6)$$

The coefficients to be estimated are denoted by δ , $\delta_{c,m}$ and $\delta_{s,m}$, and $\bar{i} \in (0, 1]$ denotes a normalized intraday time trend defined as the number of minutes from opening until i divided by the length of the trading day, i.e. $\bar{i} = i/360$. The order M is selected according to the Bayes Information Criterion (BIC) within each 30-day rolling window. To avoid forward-looking biases in the forecasting study, at each observation the seasonality component is estimated using previous data only. The sample of seasonality standardized cumulative one-minute trading volumes accordingly covers the period from 14 February to 31 December 2008, corresponding to 220 trading days and 79,200 observations per stock. In nearly all cases, $M = 6$ is selected. We observe that the estimated daily seasonality factors change mildly in their level reflecting slight long-term movements. The intraday shape is rather stable.

Figure 3.4 displays the intra-day periodicity components associated with the lowest and largest monthly volumes, respectively, observed through the sample period. We observe the well-known (asymmetric) U-shaped intraday pattern with high volumes at the opening and before market closure. Before closure, it is evident that traders intend

3 Data




Figure 3.4: Estimated intra-day periodicity components for cumulative one-minute trading volumes (in units of 100,000 and plotted against the time of the day) of selected companies at NASDAQ on 2 September (blue, lowest 30-day trading volume) and 30 October 2008 (red, highest 30-day volume).

to close their positions creating high activity.

4 Applications

An economist is someone who sees something that works in practice and wonders if it would work in theory.
Ronald Reagan

This chapter summarizes empirical results concerning structural and adaptive modelling of financial series. We start with the modelling and forecasting of the (high-dimensional) liquidity supply. The predictions are based on an intimate relationship between the order book's shape and the current stock market conditions. Note that forecasting liquidity supply implies prediction of the LOB curves (quantity), as well as the market conditions (price) over time. All financial and economic applications are consequently related to liquidity supply prediction. One of the highlights is a proposed trading strategy suitable for reducing transaction costs in order splitting, i.e. cost related to trading relatively large market orders over the course of a trading day.

The second part is dominated by a case study of (short-term) trading volume predictability. Recall, using our localized MEM methodology one may model and forecast any (positive) financial process, such as, e.g., durations, trading volumes, bid-ask spreads, volatilities or trading costs. Two prediction methods are statistically compared, namely the local parametric approach (LPA) and a 'standard' approach with ad-hoc selected estimation window. Results favor the LPA technique. In practice, short-term volume forecasts are used for trading cost optimization, particularly in trading strategies related to the execution of large orders. This aspect is yet to be explored.

The last part is devoted to market microstructure modelling, where empirical evidence for the existence of the EPK paradox on the leading stock markets is provided. The results are discussed in a dynamic context, across countries and over different economic scenarios.

4.1 Modelling and Forecasting Liquidity Supply using Semiparametric Factor Dynamics

For each stock individually (BHP, NAB, MIM and WOW) we employ a two step modelling procedure. In the first step, the shape of order book curves is modelled in dependence of relative price deviations from the best bid price and best ask price, $S_{t,j}^b$ and $S_{t,j}^a$, respectively. The dynamics of the estimated factor loadings is in the second step analyzed jointly with the best bid/ask quotes and the bid-ask spread using a vector error correction (VEC) specification. Due to the estimation complexity, the seasonal trend coefficients in (3.1) and (3.2) are not estimated jointly with the unknown parameters (matrix A) and the factor loadings from (2.5).

4 Applications

In the modelling part we focus on the cross-dependency between the bid and ask side of the market, the relations between the LOB and the quotes, as well as on the impact of the bid-ask spread on liquidity supply. The predictability of order book's shape is related to various covariates, e.g., the past trading volume, past log returns and past (realized) volatility.

4.1.1 Limit Order Book Modelling using the DSFM

We focus on two implementation methods of the DSFM:

- (i) Separated approach: the bid side $Y_{t,j}^b \in \mathbb{R}^{101}$ and the ask side $Y_{t,j}^a \in \mathbb{R}^{101}$ are analyzed separately.
- (ii) Combined approach: both sides of the limit order book are modelled simultaneously, with the bid side reversed, i.e. $(-Y_{t,j}^b, Y_{t,j}^a) \in \mathbb{R}^{202}$.

We impose $K = 20$ and $K = 40$ knots for the B-spline functions in the case of a separated and combined approach, respectively. Using more knots does not result in significant improvements of the explained variance or in the corresponding RMSE, as defined in 2.6 and 2.7.

Up to approximately 95% of the explained variation in order curves can be explained using $L = 2$ factors, see Table 4.1. The marginal contribution of a potentially third factor is very small. A two-factor DSFM specification is therefore sufficient to capture the curve dynamics and is used in the sequel of the analysis.

L	BID				ASK			
	BHP	NAB	MIM	WOW	BHP	NAB	MIM	WOW
Separated								
1	0.925	0.934	0.990	0.916	0.916	0.909	0.946	0.938
2	0.964	0.965	0.996	0.975	0.941	0.948	0.953	0.959
3	0.971	0.976	0.996	0.981	0.941	0.961	0.949	0.964
Combined								
1	0.922	0.522	0.762	0.558	0.546	0.806	0.696	0.944
2	0.921	0.936	0.975	0.914	0.930	0.912	0.951	0.948
3	0.961	0.938	0.977	0.972	0.932	0.950	0.973	0.949

Table 4.1: Explained variance (EV) of estimated order book variations depending on relative prices based on different number of factors L using both DSFM approaches.

In almost all cases the DSFM-Separated approach outperforms the DSFM-Combined approach in terms of a higher proportion of explained variance and lower values of the root mean squared error, see, e.g., Table 4.2. The root mean squared errors for different absolute price levels j , $\check{S}_{t,j}^b$ and $\check{S}_{t,j}^a$, respectively, are compared in Figure 4.1. Again, at almost every price level the DSFM-Separated approach outperforms

4.1 Modelling and Forecasting Liquidity Supply using Semiparametric Factor Dynamics

the DSFM-Combined approach. The remainder of the analysis will therefore rely on the DSFM-Separated approach with two factors.

L	BID				ASK			
	BHP	NAB	MIM	WOW	BHP	NAB	MIM	WOW
Separated								
1	3.49	2.51	0.29	2.10	2.60	3.09	0.81	2.73
2	2.40	1.82	0.19	1.16	2.18	2.32	0.76	2.22
3	2.17	1.52	0.18	0.10	2.18	2.02	0.79	2.07
Combined								
1	3.55	6.75	1.41	4.81	6.03	4.50	1.93	2.59
2	3.57	2.47	0.46	2.13	2.37	3.03	0.78	2.50
3	2.50	2.44	0.44	1.21	2.33	2.29	0.57	2.49

Table 4.2: Root mean squared errors (RMSEs) implied by estimated order book variations depending on relative prices based on different number of factors L using both DSFM approaches.




Figure 4.1: Root mean squared errors (RMSEs) for different absolute price levels, $\check{S}_{t,j}^b$ (red) and $\check{S}_{t,j}^a$ (blue), using the DSFM-Separated (solid) and the DSFM-Combined approach (dashed).

Estimated first and second factor \hat{m}_1 and \hat{m}_2 in dependence of the relative price grids are plotted in Figure 4.2. The first factor captures the overall slope of the curve which is associated with the average trading costs for all volume levels on the corresponding sides of the market. Order curve fluctuations are captured by the second factor around the overall slope and thus can be interpreted as a 'curvature' factor.

We observe that the shape of the second factor looks different for levels close to the best bid/ask quotes than for levels very deep in the book. The shapes of the estimated factors are remarkably similar for all stocks, except for MIM. The shapes of both factors of MIM are quite similar and significantly deviate from those reported for the other stocks. This finding is explained by the fact that liquidity is concentrated on relatively few price levels around the best ask and bid quotes. For higher levels the book flattens out.

4 Applications




Figure 4.2: Estimated first and second factor of the limit order book depending on relative price levels using the DSFM-separated approach with two factors for selected stocks traded at the ASX from July 8 to August 16, 2002 (30 trading days). Red: bid curve, blue: ask curve.

It is unclear a priori whether modelling order book curves based on all $J = 101$ price levels is most appropriate for prediction. For example, predictive information may depend on the distance to the best quotes. Price levels far away from the best quotes can carry important information, but also virtually contain only noise (stale orders). To mitigate this challenges, we base our model selection on in-sample information and evaluate the explained variance when not the full grid of 101 levels but just 25, 50 and 75 levels are employed. The results are reported in Table 4.3.

Employing the full curve ($J = 101$) yields the highest explained variance. This is particularly true for order books of less liquid stocks. It turns out that the factor structure remains unchanged, and therefore we proceed with the analysis of the full order book.

The estimated factor loadings, \hat{Z}_t^b and \hat{Z}_t^a , strongly vary over time reflecting time variations in the shape of the book, see, e.g., Figure 4.3. The series reveal clustering structures indicating a relatively high persistence. This is not surprising, since order book inventories do not change too severely during short time periods. On higher frequencies than 5 minutes, this persistence further increases, ultimately driving the factor loadings toward unit root processes. This behavior is distinct for less frequently traded stocks and less severe for highly active stocks, see, e.g., Hautsch and Huang [2011].

The high persistence is confirmed by autocorrelation functions of \hat{Z}_t^b and \hat{Z}_t^a and corresponding unit root and stationarity tests. According to the Schmidt-Phillips test, see, e.g., Schmidt and Phillips [1992], as shown in Table 4.4 for all processes the null hypothesis of a unit root can be rejected at the 5% significance level. Testing the null

	BID				ASK			
	BHP	NAB	MIM	WOW	BHP	NAB	MIM	WOW
$L = 1, J = 101$	0.925	0.934	0.990	0.916	0.916	0.909	0.946	0.938
$L = 2, J = 101$	0.964	0.965	0.996	0.975	0.941	0.948	0.953	0.959
$L = 3, J = 101$	0.971	0.976	0.996	0.981	0.941	0.961	0.949	0.964
$L = 1, J = 75$	0.912	0.926	0.974	0.912	0.927	0.907	0.935	0.938
$L = 2, J = 75$	0.953	0.960	0.979	0.966	0.950	0.947	0.941	0.961
$L = 3, J = 75$	0.959	0.972	0.980	0.973	0.956	0.960	0.859	0.971
$L = 1, J = 50$	0.908	0.920	0.972	0.914	0.933	0.910	0.906	0.932
$L = 2, J = 50$	0.951	0.958	0.981	0.959	0.961	0.950	0.931	0.962
$L = 3, J = 50$	0.955	0.968	0.982	0.970	0.968	0.961	0.926	0.972
$L = 1, J = 25$	0.895	0.868	0.964	0.901	0.904	0.882	0.852	0.849
$L = 2, J = 25$	0.925	0.904	0.978	0.932	0.942	0.920	0.897	0.892
$L = 3, J = 25$	0.930	0.908	0.978	0.938	0.949	0.922	0.908	0.889

Table 4.3: Explained variance (EV) of estimated order books depending on relative prices based on different number of factors L and price grid points using the DSFM-Separated approach

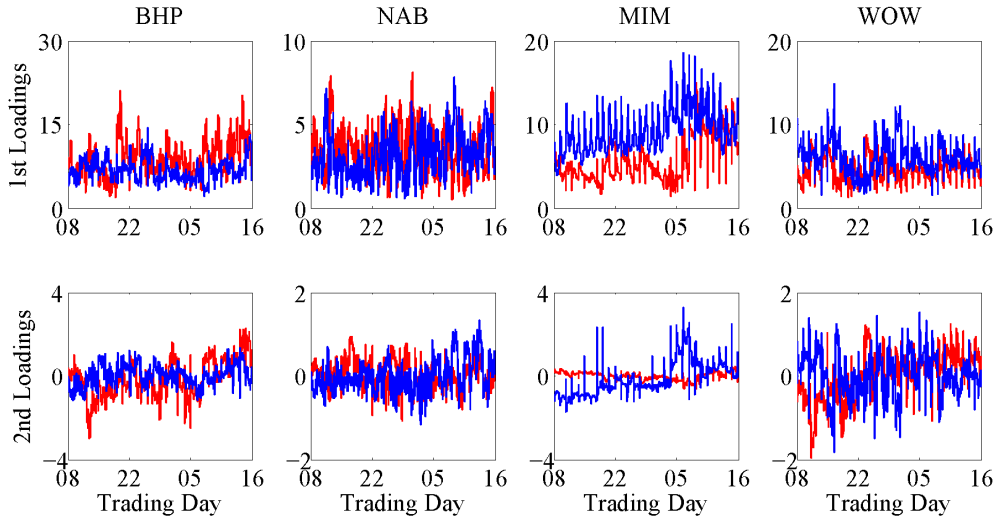


Figure 4.3: Estimated first and second factor loadings of the limit order book depending on relative price levels using the DSFM-separated approach with two factors for selected stocks traded at the ASX from July 8 to August 16, 2002 (30 trading days). Red: bid curve, blue: ask curve.

4 Applications

of stationarity using the KPSS test, see, e.g., Kwiatkowski et al. [1992], implies no rejections for the majority of the processes. In five cases we have to reject stationarity, see, e.g. 4.5. To test for possible cointegration between the factor loadings, we perform Johansen's trace test, see, e.g., Johansen [1991]. No significant evidence is found for common stochastic trends underlying the order book.

Factor Loadings	BID				ASK			
	BHP	NAB	MIM	WOW	BHP	NAB	MIM	WOW
$\hat{Z}_{1,t}$	-74.95	-164.33	-67.16	-158.90	-69.89	-145.47	-111.34	-102.56
$\hat{Z}_{2,t}$	-71.21	-201.53	-53.88	-186.95	-143.59	-159.49	-182.96	-141.29

Table 4.4: Schmidt-Phillips test statistics for estimated factor loadings (H_0 : unit root, critical values are -15.0, -18.10 and -25.20 for significance levels 10%, 5% and 1%, respectively.)

Factor Loadings	BID				ASK			
	BHP	NAB	MIM	WOW	BHP	NAB	MIM	WOW
$\hat{Z}_{1,t}$	0.10	0.06	0.26	0.06	0.16	0.11	0.17	0.09
$\hat{Z}_{2,t}$	0.12	0.05	0.33	0.18	0.17	0.15	0.12	0.12

Table 4.5: KPSS test statistics for estimated factor loadings (H_0 : weak stationarity, critical values are 0.12, 0.15 and 0.22 for significance levels 10%, 5% and 1%, respectively.)

A graphical illustration depicting the estimated vs. the actually observed limit order book curve suggest that the model fits the observed order book very well, see, e.g., Figure 4.4 in the case of NAB. For other stocks we find similar results. The fit is best for price levels close to the best ask and bid quotes, at any chosen trading day and stock. Slight deviations are observed deeply in the book. Note that the last case is less relevant for financial applications.




Figure 4.4: True (solid) and estimated (dashed) limit order book using the DSFM-separated approach with two factors ($EV \approx 95\%$) on 8 July 2002 for NAB. Red: bid curve, blue: ask curve.

4.1.2 Modelling Limit Order Book Dynamics

The limit order book dynamics is now investigated in a multivariate time series modelling context, recall our modelling philosophy: *smooth in space and parametric in time*. The order book dynamics is additionally related to the time evolution of additional covariates.

For each stock (BHP, NAB, MIM and WOW) we focus on the dynamics of the four estimated stationary factor loadings. Denote the first (1) and second (2) factor loadings for the bid (b) and ask side (a) by loadings by $\hat{Z}_{1,t}^b, \hat{Z}_{2,t}^b, \hat{Z}_{1,t}^a$ and $\hat{Z}_{2,t}^a$. Jointly with the best bid and the best ask price returns, consider a (six dimensional) vector of endogenous variables

$$z_t = \left(\hat{Z}_{1,t}^b, \hat{Z}_{2,t}^b, \hat{Z}_{1,t}^a, \hat{Z}_{2,t}^a, \Delta \log \tilde{S}_{t,101}^b, \Delta \log \tilde{S}_{t,1}^a \right)^\top,$$

with $\Delta \log \tilde{S}_{t,101}^b$ and $\Delta \log \tilde{S}_{t,1}^a$ representing the best bid and ask price return, respectively.

The bid-ask spread serves as a cointegration relationship between the two integrated ask and bid price return series, see, e.g., Engle and Patton [2004] and Hautsch and Huang [2011]. With stationary factor loadings, a vector error correction (VEC) specification of order q with the spread as the only cointegration relationship is established

$$z_t = c + \Gamma_1 z_{t-1} + \dots + \Gamma_q z_{t-q} + \gamma \left(\log \tilde{S}_{t-1,101}^b - \log \tilde{S}_{t-1,1}^a \right) + \varepsilon_t. \quad (4.1)$$

Here c denotes a vector with constants, vector $\gamma = (\gamma_1, \dots, \gamma_6)^\top$ collects parameters associated with the lagged bid-ask spread and ε_t represents a white noise error term. The matrices $\Gamma_1, \Gamma_2, \dots, \Gamma_q$ are parameter matrices associated with lagged endogenous variables. Technically, we determine the order q according to the Bayes information criteria (BIC).

A maximum lag order of $q = 4$ is sufficient in all cases. We select the following model orders: BHP and WOW ($q = 3$), NAB ($q = 2$) and MIM ($q = 4$). The estimates of matrix Γ_1 and vector γ for all four stocks are reported here since they contain the most relevant economic information (5% significance is denoted by an asterisk (*)):

$$\begin{bmatrix} 0.95^* & 0.63^* & -0.05 & -0.26^* & 3.03 & -18.08 \\ 0.02^* & 0.79^* & 0.00 & 0.04 & 10.68 & -16.12 \\ 0.04^* & 0.00 & 0.75^* & 0.02 & -59.60 & 67.60 \\ -0.00 & 0.04 & 0.02^* & 0.77^* & -13.99 & 13.55 \\ 0.00^* & 0.00 & -0.00^* & 0.00^* & -0.59 & 0.29 \\ 0.00^* & 0.00 & -0.00^* & 0.00^* & -0.26 & -0.04 \end{bmatrix}, \begin{bmatrix} -95.70 \\ -34.13 \\ 86.83 \\ -13.21 \\ -0.42 \\ 0.02 \end{bmatrix},$$

$$\begin{bmatrix} 0.71^* & 0.16 & -0.04 & -0.21 & 123.78^* & -124.07^* \\ 0.04^* & 0.78^* & -0.00 & 0.07^* & -22.39^* & 21.91 \\ 0.04 & 0.13 & 0.73^* & 0.18 & -88.56^* & 86.91^* \\ -0.03^* & -0.03 & 0.03^* & 0.71^* & 26.03^* & -25.46^* \\ 0.00^* & 0.00^* & -0.00^* & -0.00 & 0.21 & -0.34 \\ 0.00 & 0.00^* & -0.00^* & 0.00 & 0.29 & -0.41 \end{bmatrix}, \begin{bmatrix} -174.41^* \\ 9.26 \\ 47.60 \\ -20.59 \\ -0.85 \\ -0.04 \end{bmatrix},$$

4 Applications

$$\begin{bmatrix} 0.90^* & 1.29^* & -0.00 & 0.55^* & -46.92 & 50.79 \\ 0.00 & 0.93^* & -0.01^* & -0.01 & 1.12 & -1.49 \\ -0.02 & 1.23^* & 0.99^* & 0.48^* & 31.56 & -25.63 \\ 0.00 & 0.04 & 0.03^* & 0.84^* & 6.73 & -5.89 \\ 0.00 & 0.00 & -0.00 & -0.00 & 0.40 & -0.58 \\ 0.00 & 0.00 & -0.00 & -0.00 & 0.90 & -1.09 \end{bmatrix}, \begin{bmatrix} 62.01^* \\ 0.25 \\ -44.50^* \\ -21.66^* \\ -0.28 \\ -0.18 \end{bmatrix} \text{ and } \\
 \begin{bmatrix} 0.74^* & -0.02 & 0.12^* & 0.38^* & 28.87 & -37.11 \\ 0.04^* & 0.82^* & -0.02^* & -0.04 & 2.53 & -3.58 \\ 0.04 & 0.03 & 0.87^* & 0.19^* & -70.61^* & 72.84^* \\ -0.03^* & 0.02 & 0.02^* & 0.83^* & 12.81 & -13.70 \\ -0.00^* & -0.00^* & 0.00 & 0.00^* & 0.02 & -0.15 \\ -0.00^* & -0.00^* & 0.00 & 0.00^* & 0.21 & -0.34 \end{bmatrix}, \begin{bmatrix} -27.14 \\ -6.33 \\ 59.98 \\ -4.04 \\ -0.51 \\ 0.05 \end{bmatrix}.$$

Firstly, empirical results suggest strong own-process dynamics and relatively weak cross-dependencies between the endogenous variables. The market cross-dependencies are most pronounced for less frequently traded stocks (MIM and WOW). Time variations in the liquidity schedule on the one side is almost unaffected by that on the other side due to the quite weak inter-dependencies.

Secondly, quote changes are short-run predictable given the shape of the order book. Changes in the factor loading have a short term impact (up to 5-10 minutes) on the quote changes. High frequently traded stocks (BHP and NAB) have more pronounced impact than less liquid stocks (MIM and WOW). A shock on the bid side results economically in upward rotation the bid curve, i.e. it induces a higher sell pressure. We find that such a shock leads to an instantaneous decrease in the best bid quote followed by a significant increase of the price within the next few minutes, see, e.g. Figure 4.5. The price movements are driven by a growing buy pressure reflected by an increase of bid depth at and behind the market. The impulse responses of ask and bid quotes driven by a shock in the order book slope are plotted in Figure 4.5. Observe that the effects are quite distinct on the bid side and a more neutral on the ask side. Note that the quote predictability holds over comparably short horizons, up to, say 10-15 minutes. For daily order execution strategies these effects are only of limited use, see Section 4.1.4.

Thirdly, slight evidence is found for asymmetric reactions of slope factor loadings on changes of the bid-ask spread. Rising spreads tend to reduce (increase) the order aggressiveness on the bid (ask) side. As the bid and ask curves move apart, the price is therefore (on average) decreasing. The price is expected to increase as the bid-ask spread shrinks. This re-confirms our finding in Chapter 3: liquidity variations on the bid side are higher than those on the ask side with more sell than buy activities.




Figure 4.5: Orthogonalized impulse-response analysis: responses of the best bid quote return to a one standard deviation shock in the estimated first bid factor loadings (upper panel) and response of the best ask quote return to a one standard deviation shock in the estimated first ask factor loadings (lower panel). We employ the DSFM-separated approach with two factors and a VEC specification for selected stocks traded at the ASX from 8 July to 16 August 2002 (30 trading days). The response variable always enters the VEC specification in the first position. 95% confidence intervals are shown with dashed lines.

4.1.3 Drivers of the Order Book Shape

An important finding is that quote changes are short-run predictable given the shape of the order book. Now we analyze to which extent is the order book's shape predictable based on (weakly exogenous) trading variables. Four variables are selected for which we expect to observe the strongest impact on the LOB's shape, namely the past 5-min aggregated trading volume on both sides of the market representing the recent liquidity demand, the best bid/ask price, the past 5-min log mid-quote return as well as the past 5-min volatility.

Representing liquidity demand, we consider the buy and sell trading volumes at time t by the sum of traded quantities from all market orders r , \tilde{Q}_r^b and \tilde{Q}_r^s , over five minutes interval, namely, $\tilde{Q}_t^b = \sum_{r=1}^{R_t^b} \tilde{Q}_r^b$ and $\tilde{Q}_t^s = \sum_{r=1}^{R_t^s} \tilde{Q}_r^s$, with R_t^b and R_t^s denoting the number of buy and sell orders over the interval $(t-1, t]$, respectively. Log returns r_t and volatility V_t are computed as

$$r_t = \log \frac{\tilde{S}_t^*}{\tilde{S}_{t-1}^*} \quad (4.2)$$

$$V_t = r_t^2, \quad (4.3)$$

with mid-quotes \tilde{S}_t^* and \tilde{S}_{t-1}^* at time points t and $t-1$, respectively. The trading volumes and volatility series are seasonally adjusted following the procedure explained above. The necessary DSFM standardization of variables into the $[-1, 1]$ interval is performed based on the minimum and maximum observations. As nonparametric regression becomes computationally cumbersome for a high number of regressors, we include the regressors only individually together with the relative price levels.

The estimated first factors for the bid and the ask side in dependence of the past 5-min sell and buy trading volumes as well as best bid and ask prices are shown in Figures 4.6, 4.7, 4.8 and 4.9, respectively.

Past liquidity demand influences the order book curve. A high trading volume removes a non-trivial part of the pending volume in the book. Variation of the factor's shape is then induced either by the complete absorption of price levels close to best quotes or by the distributional change of the pending volumes across the (relative) price levels.

One observes that the curve flattens in the area of high volumes, as well as a decaying pattern if the volume sizes decline. The maximum slope (and thus the highest level of liquidity supply) is observed for magnitudes of the standardized volume between -1 and 0 , i.e., comparably small trading volumes, most likely due to the boundary effect of non-parametric regression or the standardization procedure. Note that because of the curse-of-dimensionality we cannot simultaneously control for other variables. For example, small trading volumes can indicate the occurrence of market imbalances or might be associated with wide spreads. Investors would in both cases be forced to post limit orders rather than market orders. That could help to explain the decaying shape of the factors after having observed small trading volumes.

To evaluate whether the inclusion of regressors further increases the model's goodness-

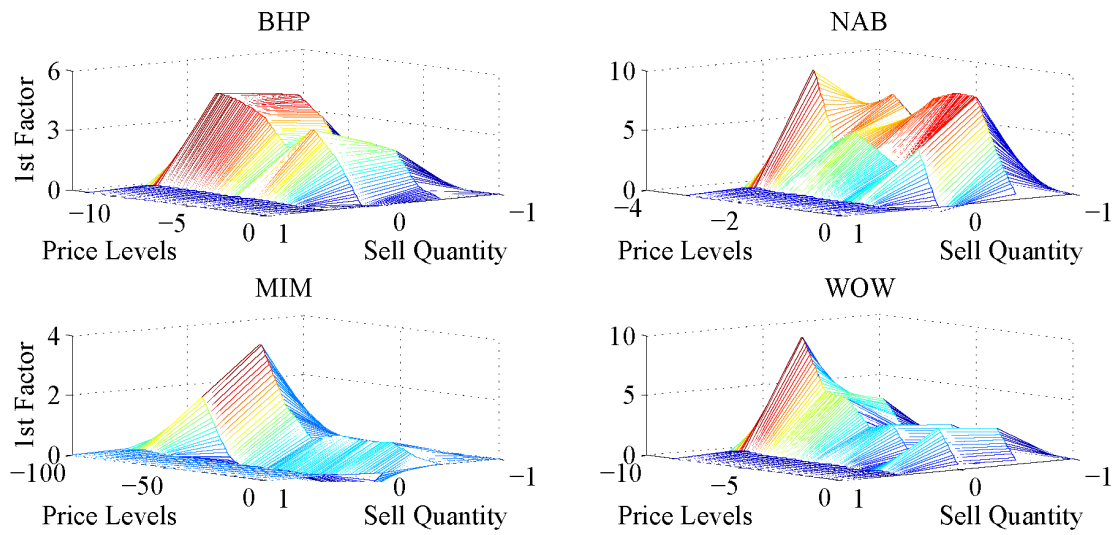


Figure 4.6: Estimated first factors of the bid side with respect to relative price levels and the past log traded sell volume using the DSFM-Separated approach with two factors for selected stocks traded at the ASX from 8 July to 16 August 2002 (30 trading days).

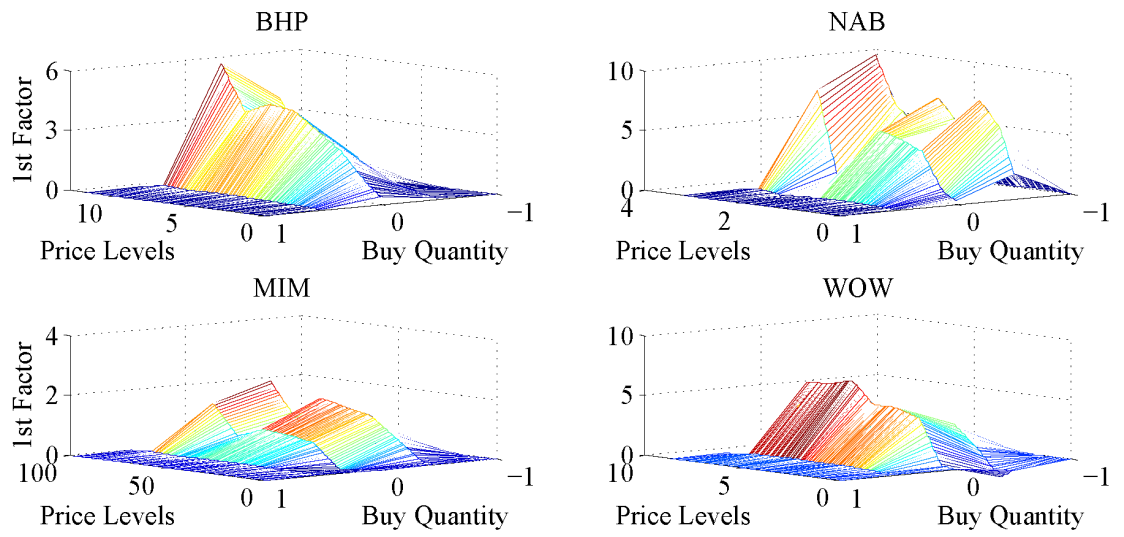


Figure 4.7: Estimated first factors of the ask side with respect to relative price levels and the past log traded buy volume using the DSFM-Separated approach with two factors for selected stocks traded at the ASX from 8 July to 16 August 2002 (30 trading days).

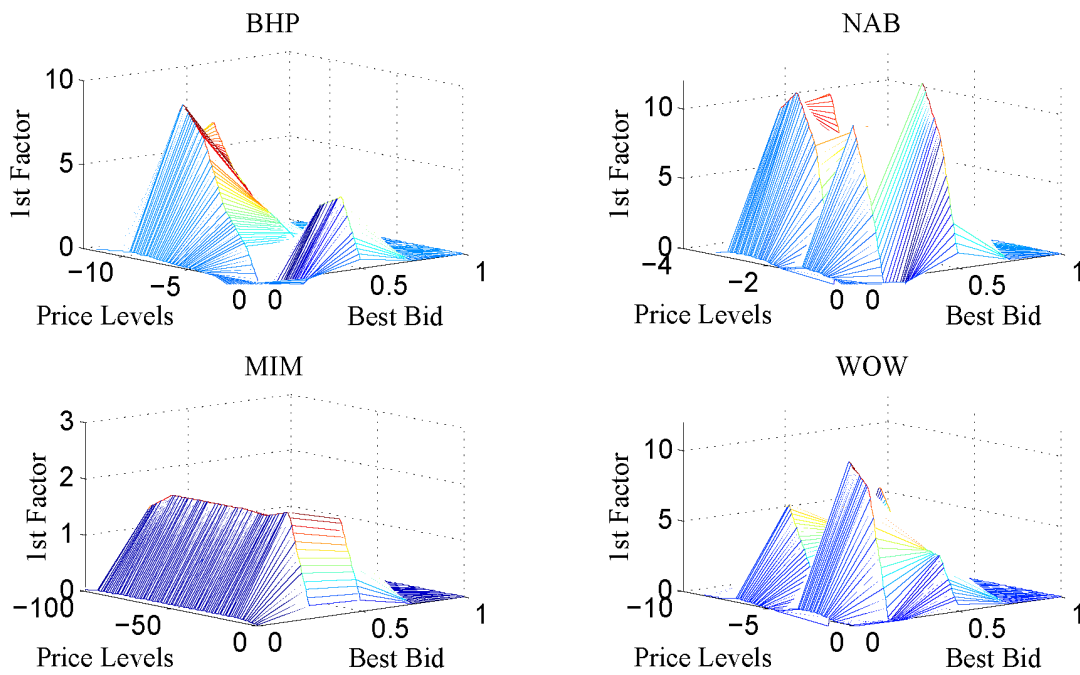


Figure 4.8: Estimated first factors of the bid side with respect to relative price levels and the best bid price using the DSFM-Separated approach with two factors for selected stocks traded at the ASX from 8 July to 16 August 2002 (30 trading days).

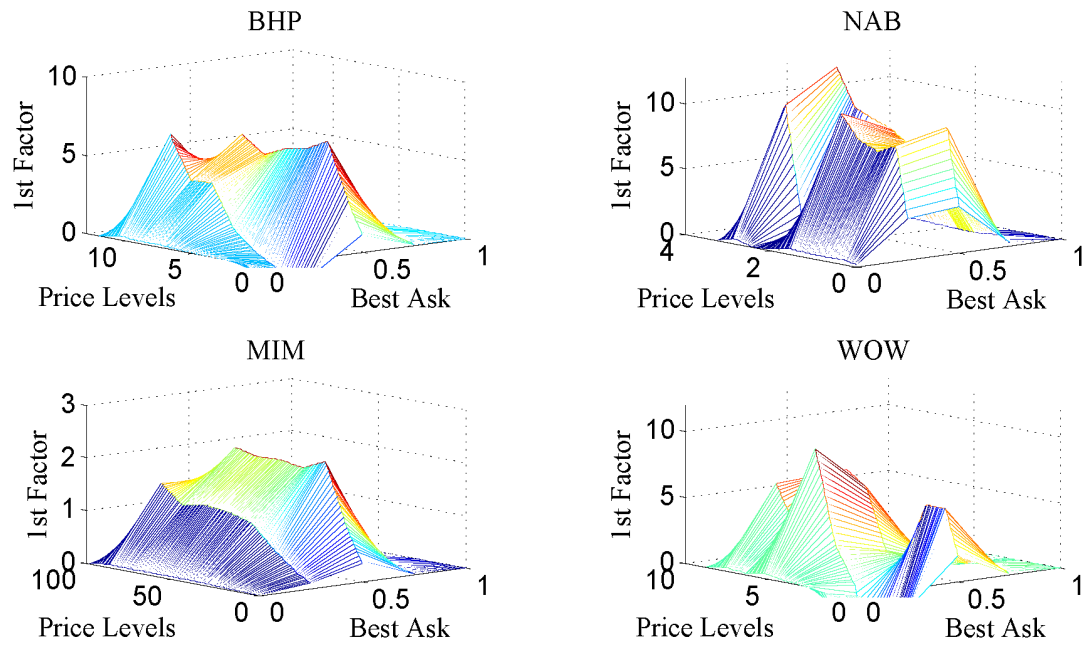


Figure 4.9: Estimated first factors of the ask side with respect to relative price levels and the best ask price using the DSFM-Separated approach with two factors for selected stocks traded at the ASX from 8 July to 16 August 2002 (30 trading days).

4.1 Modelling and Forecasting Liquidity Supply using Semiparametric Factor Dynamics

of-fit, we calculated the corresponding RMSEs. Comparing the results with that for the basis model shows that the included regressors yield higher estimation errors, see, e.g., Tables 4.2 (basis model), 4.6 (buy volume), 4.7 (sell volume) and 4.8 (log-returns). The range of the RMSEs for best bid/ask quotes for the DSFM-Separated approach with two factors is [3.78, 11.19]. Inclusion of additional regressors generates more noise overcompensating a possibly higher explanatory power.

L	BID				ASK			
	BHP	NAB	MIM	WOW	BHP	NAB	MIM	WOW
Separated								
1	18.17	28.25	18.99	13.47	11.53	28.82	17.66	17.77
2	18.45	159.52	22.01	15.16	12.09	217.03	122.93	18.72
3	19.75	610.53	52.38	15.29	12.45	564.95	97.05	21.13
Combined								
1	17.92	93.34	11.01	12.84	55.25	82.63	24.82	44.68
2	18.02	367.58	11.07	12.74	55.11	352.75	25.30	44.46
3	18.21	312.20	7.36	12.78	55.21	299.23	20.03	44.25

Table 4.6: Root mean squared error (RMSE) of the estimated limit order book data for all selected stocks based on the traded buy quantities evaluated for different number of factors L using both DSFM approaches

L	BID				ASK			
	BHP	NAB	MIM	WOW	BHP	NAB	MIM	WOW
Separated								
1	26.69	16.73	27.28	12.50	20.97	16.20	51.20	17.01
2	42.95	45.59	193.50	12.95	37.00	19.97	61.58	17.32
3	23.82	49.84	222.66	18.61	30.79	22.77	154.01	21.72
Combined								
1	23.10	14.95	10.17	12.58	60.54	46.43	23.66	44.35
2	22.42	17.63	16.01	12.69	59.84	48.62	27.18	44.25
3	22.76	15.90	19.26	13.29	59.44	46.07	28.02	45.00

Table 4.7: Root mean squared error (RMSE) of the estimated limit order book data for all selected stocks based on the traded sell quantities evaluated for different number of factors L using both DSFM approaches

We conclude that due to the lower dimensionality of the regressors in comparison with that of the limit order book the modelling performance declines. As the included regressors do not reveal any variation across the levels of the book, the explanatory variables can improve only the the model's dynamic fit. It is evident that the dynamic fit is not sufficient to obtain an overall reduction of modelling errors.

4.1.4 Forecasting Liquidity Supply

In this section we analyze the model's forecasting performance in a realistic setting mimicking the situation in financial applications. We assume that an investor observes the limit order book at 5-minute snapshots together with the history over the past 10 trading days. During a trading day an investor updates the limit order book every 5 minutes and has to produce forecasts for all (5 minutes) intervals of the remainder of the trading day. This forecasts may be used for optimal order execution during the course of a day. The investor does not exceed beyond the end of the trading day in order to avoid overnight effects. The forecasting horizon h therefore subsequently declines if we approach market closure.

Starting at 10:30, an investor produces multi-step forecasts for all remaining $h = 66$ intervals during the day. At 15:50, we are left with a horizon of $h = 1$. Since quotes, according to our results above, are only predictable over very short (virtually irrelevant) horizons for our forecasting study, we do not explicitly incorporate this information here.

The model is re-estimated every five minutes based on past information over a fixed window of 10 trading days. The estimation period includes the recent observation. We do not produce forecasts for the first two weeks of our sample and thus focus on the period between 22 July and 16 August 2002. We thereby cover the period of 20 trading days. In accordance with our in-sample results the DSFM-Separated approach based on two factors without additional regressors is selected as underlying specification.

We evaluate our model's performance against the naive forecast. Using the naive forecasting approach, we assume that the investor has no appropriate prediction model and takes the last observed limit order book as the forecast for the remainder of the trading day. The models are compared based on the predicted volume $\hat{Y}_{t'+h,j}$ at a given time point t' from 22 July at 10:25 until 16 August 2002, at 15:50, $t' = 693, \dots, 2069 = T - 1$, over a forecasting horizon $1 \leq h \leq 66$, and over the absolute price level j . Our investor can use the following two approaches in order to forecast liquidity supply:

- (i) DSFM approach: Firstly, the factors and factor loadings are estimated using the DSFM-Separated approach with two factors. We impose $K = 20$ knots for the B-spline basis functions and consider the past 690 observed (de-seasonalized) limit order book curves. At time point t' , relative price levels $S_{t'-691:t',j}^b$ and $S_{t'-691:t',j}^a$ as well as the de-seasonalized observed bid and ask sides $Y_{t'-691:t',j}^b$ and $Y_{t'-691:t',j}^a$ enter the estimation. There are 66 estimates for the bid (ask) side per day for each stock, in total 1320 estimates over 20 days.

Secondly, because short-term quote return predictability does not enter our forecasting setup, we focus only on the forecasts the liquidity supply. A simple 4-dimensional VAR(p) model is employed for the four time-varying factor loadings, $\hat{Z}_{1,t}^b$, $\hat{Z}_{2,t}^b$, $\hat{Z}_{1,t}^a$ and $\hat{Z}_{2,t}^a$. According to the BIC, a maximum lag order $p = 4$ is sufficient when the entire time series (30 trading days) is fitted by a VAR(p) model. The following VAR(p) models are selected: BHP and MIM - VAR(4), NAB - VAR(2) and WOW - VAR(3). We forecast the factor loadings over the forecasting period $\hat{Z}_{t'+h}$ using this specifications. The predicted factor loadings together with the

estimated time-invariant factors $\hat{m}_l, l = 0, \dots, 2$ are then used to predict liquidity supply.

- (ii) Naive approach: Among all historical 690 limit order book curves (10 trading days), only the last one at time t' , $(Y_{t',j}^b, Y_{t',j}^a)$, is selected as the h -step ahead order book forecast.

Using the root mean squared prediction error (RMSPE) we evaluate the performance of the competing forecasting approaches. The RMSPE is a version of the in-sample RMSE (2.7), i.e. the sum over the sampling periods t and the sample size T are replaced by the forecasting horizons h and H , respectively.

Note that since future quotes and relative price grids are not predicted by the model, we impose the assumption that quotes follow random walk processes and that the bid-ask spread remains constant. Future quotes are therefore predicted using the current observation while the predicted future relative price grid remains constant.

4.1.5 Forecasting Results

We obtain the squared residuals after every forecasting procedure and then base the evaluation of our results on the root mean squared prediction error (RMSPE). At the same absolute price level, we record the differences between the true observed limit order book curves and the predicted ones for both approaches. For example, predicted limit order book curves and the true ones for each stock on 22 July 2002, at 11:00 and 15:00 are shown in Figure 4.10.

The predictive performance of the competing forecasting methods is illustrated by displaying the RMSPEs. For each forecasting horizon h during a trading day the RMSPEs implied by the DSFM as well as by the naive model are shown in Figure 4.11. We summarize our forecasting performance findings in the sequel.

Overall, the DSFM method outperforms the naive approach. This is a very good achievement because the naive forecasting method is a serious competitor given the high persistence in liquidity supply. Focusing on the average RMSPEs, averaged over all forecasting horizons and both order book sides, one observes that the DSFM performance is significantly higher than that of the benchmark. The average RMSPEs are reported in Table 4.9.

The DSFM outperforms the naive approach especially on the bid side of the market. Recall, during the sample period one observes a downward market with high trading activity on the bid side. The forecasting result is therefore explained by the fact that the DSFM, stipulated under the philosophy *smooth in space and parametric in time*, successfully captures the dynamics of high-dimensional objects evolving over time. This is particularly true for the more active bid side.

Our DSFM approach is clearly better than the naive model over relatively short forecasting horizons, say, up to 1 to 2 hours. We contribute this finding to the modelling flexibility. For longer horizons, the DSFM outperformance over the naive method is diminishing. Even for trading strategies over relatively longer horizons during a trading day it is justified to use our DSFM approach.

4 Applications

L	BID				ASK			
	BHP	NAB	MIM	WOW	BHP	NAB	MIM	WOW
Separated								
1	10.21	12.90	52.99	5.28	9.92	9.69	173.67	13.39
2	17.27	12.49	1104.46	9.51	15.76	12.68	7750.16	16.40
3	29.47	14.14	8456.95	6.78	16.37	11.59	2139.35	31.40
Combined								
1	17.20	13.03	151.01	4.99	44.59	25.44	146.74	34.32
2	20.01	13.16	97.09	5.08	45.29	25.23	107.25	34.21
3	16.82	14.77	57.39	5.97	43.83	25.33	56.20	34.60

Table 4.8: Root mean squared error (RMSE) of the estimated limit order book data for all selected stocks based on the intraday log-return evaluated for different number of factors L using both DSFM approaches




Figure 4.10: Predicted limit order book curves (dashed) and the true ones (solid) on July 22, 2002, at 11:00 (upper panels) and 15:00 (lower panels) at different absolute price levels in AUD. The predicted curve using the naive approach is shown with black solid line.

Approach	BID				ASK			
	BHP	NAB	MIM	WOW	BHP	NAB	MIM	WOW
Naive	7.11	7.59	6.03	6.08	6.50	5.96	5.83	6.19
DSFM	7.18	5.10	4.84	5.33	5.56	5.46	5.63	5.45

Table 4.9: Average root mean squared prediction errors (RMPSEs) of both limit order book sides implied by the DSFM-separated approach with two factors and the naive model for selected stocks traded at the ASX in the period from 22 July to 6 August 2002 (20 forecasting days).

4.1.6 Financial and Economic Applications

The DSFM successfully predicts liquidity supply over various forecasting horizons during a day. The forecasting results are now applied to three practical examples. The first example is devoted to an order execution strategy, whereas the second one deals with demand and supply elasticity forecasts. Our last example shows how the DSFM framework can be used to simultaneously predict the shape and the position of the limit order book in the short run.

EXAMPLE 1. (Trading Strategy)

Suppose that an institutional investor decides to optimize the trading costs resulting from buying or selling shares on a stock market. We assume that the investor decides to buy (sell) a certain number of shares v over the course of a trading day, starting from 10:30 until 15:40. For BHP, NAB and WOW we select the size of the traded quantity to be 5 or 10 times the average volume at the best bid or ask level. In the case of MIM the liquidity supply is more concentrated at the first few levels and the book is very thin for higher levels. The size of the traded quantity is respectively set to 2 and 5 times the average depth at the first level. The following volumes have been selected in the case of a high (a) and very high (b) liquidity demand:

- (a) BHP - 175,000 shares; NAB - 25,000 shares; WOW - 50,000 shares; MIM - 1,860,000 shares
- (b) BHP - 350,000 shares; NAB - 50,000 shares; WOW - 100,000 shares; MIM - 4,650,000 shares.

We assume furthermore that trading is performed every 5 minutes through the day, i.e. at every of the 63 possible trading time points. This assumption relaxes the computational burden. The investor makes consequently a trading decision at 10:30 and does not monitor the stock market anymore over the course of a trading day. The decision is to adopt one of the following two order execution strategies:

- (i) Splitting the buy (sell) order of size v proportionally over the trading day, i.e. 'trading' market orders of size $v/63$ at each of the 63 time points.

4 Applications

- (ii) Placing orders at those m (5 minute interval) time points throughout the day where the DSFM-based predicted implied trading costs c of the volume v are smallest (among all 63 possible periods). The orders are thus not proportionally split over the trading day, i.e. the volume v is split over the m time points according to the relative proportions of expected trading costs. At interval i , $w_i \cdot v$ shares are 'traded' at the stock market, with $w_i = c_i / \sum_{j=1}^m c_j$ for $i = 1, \dots, m$.

Strategy (i) can be seen as a special case of strategy (ii) if $m = 63$ and the volume v is equally split over the course of a trading day. In the other extreme case, when the entire volume is traded only once in a day ($m = 1$), it is required to severely 'walk up' the limit order book. Note that 'walking up' the book might be very costly because the trader would face prices that differ from the best quotes.

The trading cost predictions using the DSFM framework are computed based on the predicted order book shape at each possible 'trading' time point and the effective costs to buy or sell v shares at the bid and ask quotes prevailing at 10:25. This follows from our random walk assumption about the price dynamics.

In our setup we fix the quantity v and do not optimize it. One possibility to optimize the order size would be to minimize the predicted trading costs for relative proportions of v at different execution time points. This order size optimization lies beyond the scope of the current study. The quantity v corresponds here to the maximally possible order size at a trading moment. Our strategy selects therefore those trading points where trading the entire quantity v is expected to be cheapest. We cover the hypothetical (limiting) case of putting all weight w_i on a single time point, i.e., the so-called 'one-shot' order execution strategy.

We consider 20 forecasting days covering the period from 22 July to 16 August 2002 to implement and evaluate the competing forecasting strategies. The average percentage reduction in trading costs of strategy (ii) in excess of the strategy (i), i.e. that of equal-splitting, is shown in Figure 4.12. We select $m = 1, \dots, 63$ and consequently include both extreme cases, namely the 'one-shot' order execution ($m = 1$) and the proportional splitting strategy ($m = 63$).

Our DSFM yields better results than the proportional splitting strategy. Its trading gains are on average 10 basis points higher than those implied by the 'naive' strategy. The DSFM is quite successful in predicting the time points where the market is deep enough to absorb large orders. The results are robust relative to the trading volume size v , i.e., similar patterns are found for quantities that equal to 5 or 10 times the mean posted first level volume.

The gain curves show a non-monotonic behavior. For a very small number of daily orders m an investor would prefer a proportional trading strategy, as the costs of 'walking up' the book are quite high due to the price impact. When the number of orders m increases, the gains become increasing and positive. The benefits converge to zero as the DSFM method approaches the proportional splitting strategy. Financial benefits in the extreme case ($m = 63$) are driven by the used non-equal weighting scheme. Interestingly, an investor would benefit from trading a small number of orders per day for the MIM stock. This is obviously induced by the extremely deep order book's shape

4.1 Modelling and Forecasting Liquidity Supply using Semiparametric Factor Dynamics




Figure 4.11: Root mean squared prediction errors (RMSPEs) implied by the DSFM-separated approach with two factors for the bid side (red) as well as the ask side (blue) and by the naive approach (black) for all intra-day forecasting horizons (in hours) for selected stocks traded at the ASX. Prediction period: July 22 to August 16, 2002 (20 trading days).




Figure 4.12: Average percentage gains by reduced transaction costs compared to an equal-splitting strategy when buying (blue) and selling (red) shares based on m DSFM-predicted time points per day. Upper panel: Daily volumes corresponding to 5 (2) times the average first level market depth for BHP, NAB, WOW (MIM). Lower panel: Daily volumes corresponding to 10 (5) times the average first level market depth for BHP, NAB, WOW (MIM). Prediction period: 22 July to 16 August 2002 (20 trading days).

4 Applications

at the first levels. 'One-shot' executions of large volumes in case of MIM are therefore quite beneficial.

The selling strategy outperforms the buying activity. This is in accordance with our forecasting findings above. There are substantial differences between the sell-based and the buy-based gains for BHP, particularly in the case of a low or moderate number of trading points.

In summary, the DSFM performs reasonably well and is promising for more elaborate applications. Note that the reported results do not include transaction fees. A proportional splitting strategy certainly induces more transaction costs than a (single) market order. Our DSFM approach would thus be even more beneficial than a 'naive' trading execution strategy. Note that trading cost optimization as well as the future returns predictability power of the DSFM framework would definitely lead to an improvement of the proposed trading strategy.

EXAMPLE 2. (Demand and Supply Elasticity)

Assume that an investor decides to optimize the marginal trading costs. From an economic point of view, the investor focuses on the elasticities of the order book curves. At best bid ($\tilde{S}_{t',101}^b$) and best ask prices ($\tilde{S}_{t',1}^a$) the elasticities are computed as

$$\hat{E}_{t'+h}^d = \frac{\hat{Y}_{t'+h,1}^b - \hat{Y}_{t'+h,101}^b}{\hat{Y}_{t'+h,101}^b} / \frac{\tilde{S}_{t',1}^b - \tilde{S}_{t',101}^b}{\tilde{S}_{t',101}^b}, \quad (4.4)$$

$$\hat{E}_{t'+h}^s = \frac{\hat{Y}_{t'+h,101}^a - \hat{Y}_{t'+h,1}^a}{\hat{Y}_{t'+h,1}^a} / \frac{\tilde{S}_{t',101}^a - \tilde{S}_{t',1}^a}{\tilde{S}_{t',1}^a}, \quad (4.5)$$

for the bid and the ask side, respectively. The limit order book represents excess supply, i.e., the volume that is observed above the equilibrium trading volume which is traded at the equilibrium market price. In our work we focus on elasticities computed at the bid or the ask side. This market sides constitute only a part of the entire stock market demand and supply curves. The elasticities reported here are thus related to the marginal trading costs of excess supply.

The setup follows the forecasting framework discussed above. Shortly, we assume that an investor aims to predict the elasticities for the bid and the ask curves for all 5-min intervals over the course of a trading day for horizons $h = 1, \dots, 66$. Last 10 days of data is used for the estimation of the DSFM parameters and the price is assumed to follow a random walk. Last observed bid and ask quotes are therefore used for elasticity prediction. For illustration, the predictions of bid and ask elasticities for all trading days at 10:30 and stocks are shown in Figure 4.13.

Marginal trading costs are varying significantly over time. Since predicted elasticities reveal daily trend patterns one may use this information to improve trading strategies. Consider the NAB stock on 24 July and 30 July 2002. The bid elasticities are in absolute terms increasing on the first day, and decreasing through the second day. A trader would therefore earn if the shares are sold late on 24 July and early on 30 July provided that




Figure 4.13: Predicted demand and supply elasticities at best bid (red) and best ask prices (blue) using the DSFM-separated approach with two factors for selected stocks traded at the ASX from 22 July to 2 August 2002 (upper panels, 10 trading days) and from 5 August to 16 August 2002 (lower panels, 10 trading days).

the price does not change significantly during both trading days. The ask elasticities show converse patterns over the selected days. It would be beneficial to buy shares early on 24 July and late on 30 July, given unchanged prices. Setting a dynamic trading strategy lies beyond the scope of the paper.

EXAMPLE 3. (Predicting the Limit Order Book: Shape and Location)

Suppose that an investor aims to predict the shape of the limit order book and its location for the BHP stock on July 22 at 15:00 within the next 20 minutes. Using the DSFM-Separated approach with two factors the investor estimates the corresponding factors and factor loadings. A VEC model specification is then fitted to a multivariate time series containing the four estimated factor loadings, the best bid and the best ask quotes. As an additional regressor, the investor considers the lagged bid-ask spread.

The predicted and observed limit order book curves are shown in Figure 4.14. Six levels on each side of the book are shown. One observes better forecasts of the limit order book's shape in the short run (i.e., within the first 5-10 minutes) than in the long run. This findings support our impulse-response analysis results conducted earlier. The best price fit is at 10 minutes. An investor can therefore use the proposed DSFM modelling approach and successfully forecast liquidity supply: the shape and the location of the limit order book.

4 Applications




Figure 4.14: Predicted (dashed) and realized (solid) limit order book curves for BHP on 22 July 2002, between 15:05-15:20 using the DSFM-Separated approach with two factors.

4.2 Local Adaptive Multiplicative Error Models for High-Frequency Forecasts

The goal of the local parametric approach (LPA) is to find a length of a data window over which one can safely apply a parametric model. Our objective is to localize a multiplicative error model (MEM). Using a sequential testing procedure (i.e. the local change point (LCP) detection test) one finds an *interval of homogeneity* that is used to adaptively estimate the model parameters.

4.2.1 Parameter Dynamics

Prior to the implementation of the local MEM, we now conduct a study about the parameter dynamics. The analysis is accomplished on a rolling window basis where we vary the window length. The results of this section will be used in the modelling setup of the local MEM approach. We discuss the time evolution and the distribution of MEM parameters, as well as a possible tradeoff between parameter variability and the modelling bias. Our study results with reasonable parameter constellations that are used for simulation of critical values for the LCP detection test.

The dynamics of MEM parameters is studied based on data windows of lengths of 1 hour, 2 hours, 3 hours, 1 trading day (6 hours), 2 trading days (12 hours) and 1 trading week (30 hours). Since non-trading periods are removed, the estimation windows contain data that (potentially) cover more trading days. The EACD(1, 1) and the WACD(1, 1) model are applied to all five stocks (AAPL, CSCO, INTC, MSFT and ORCL) at each minute from 22 February to 31 December 2008 covering the period of 215 trading days. Note that the first 30 days in 2008 are used in seasonality adjustments and additional 5 days are used to obtain the first 'weekly' estimate (i.e., an estimate using one trading week of data). In total, we estimate 4,664,000 parameter vectors - 6 data window lengths, 5 stocks, 2 models, 77400 minutes. Estimated EACD parameters for INTC for estimation windows of one-day (six trading hours) and one-week (30 trading hours) are shown in Figure 4.15.

One observes that estimated parameters ($\tilde{\omega}$, $\tilde{\alpha}$ and $\tilde{\beta}$) as well as persistence levels ($\tilde{\alpha} + \tilde{\beta}$) clearly vary over time. Parameter estimates are less (more) volatile for longer

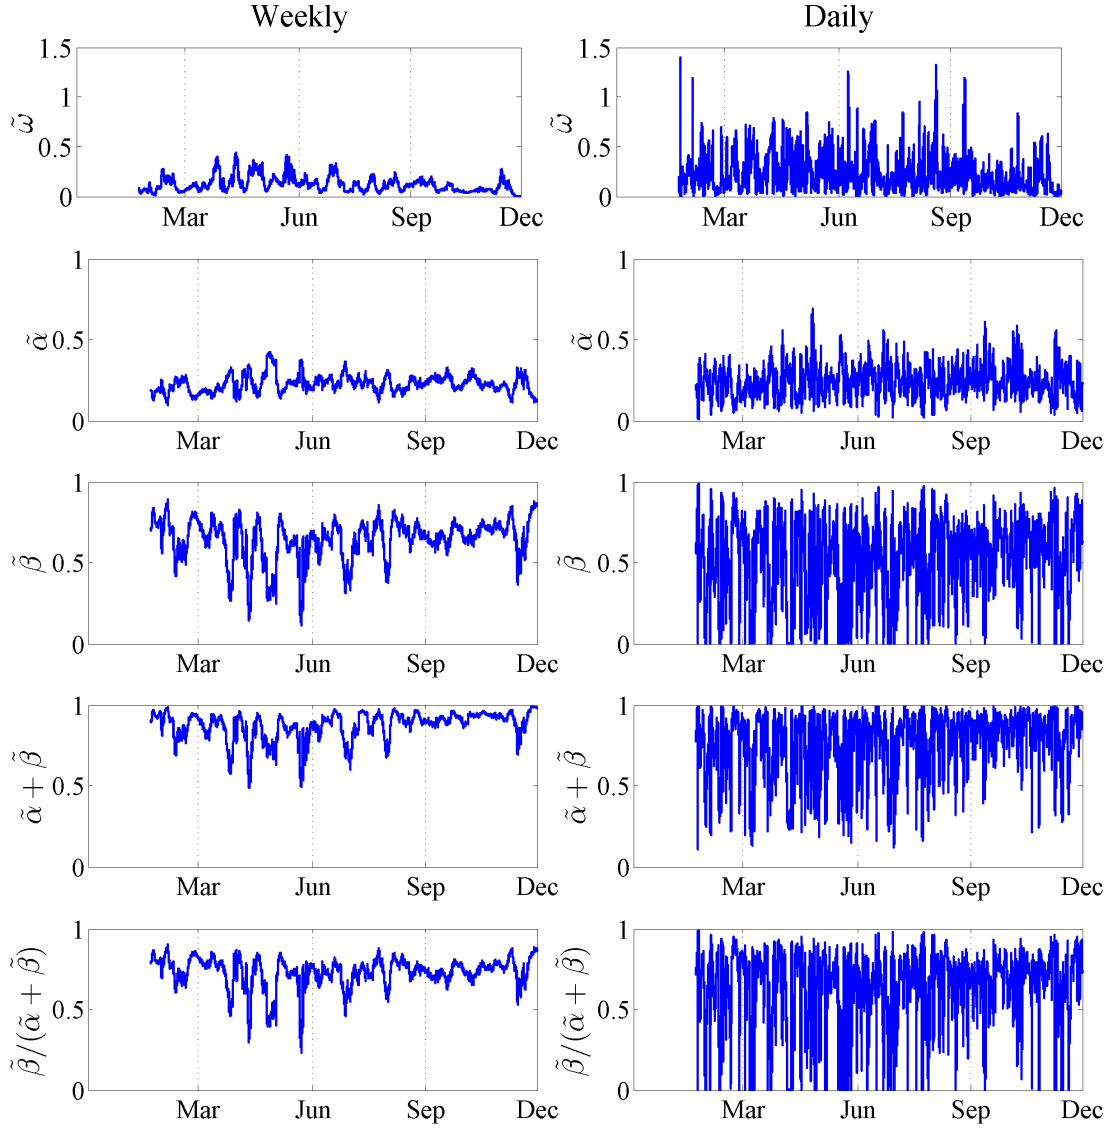


Figure 4.15: Time series of estimated 'weekly' (left panel, rolling windows covering 1800 observations) and 'daily' (right panel, rolling windows covering 360 observations) EACD(1,1) parameters and functions thereof based on seasonally adjusted one-minute trading volumes for Intel Corporation (INTC) at each minute from 22 February to 31 December 2008 (215 trading days). First 35 days are used for initialization. Based on 154,800 individual estimations.

4 Applications

(shorter) estimation windows. This finding describes our bias-variance tradeoff: precise estimates over comparably long estimation periods result in a high modelling bias and vice versa. A local MEM successfully deals with this tradeoff.

Empirical densities of estimated $\text{EACD}(1, 1)$ and $\text{WACD}(1, 1)$, are shown in Figures 4.16 and 4.17, respectively. The densities exhibit different shapes with respect to the estimation window size. There is a lower dispersion of a weekly estimated parameter, as compared to the daily one. Results for time periods between 1 hour and 1 week support this empirical finding. Modelling over long (short) time intervals thus inflates (decreases) the modelling bias and reduces (increases) the parameter variability.

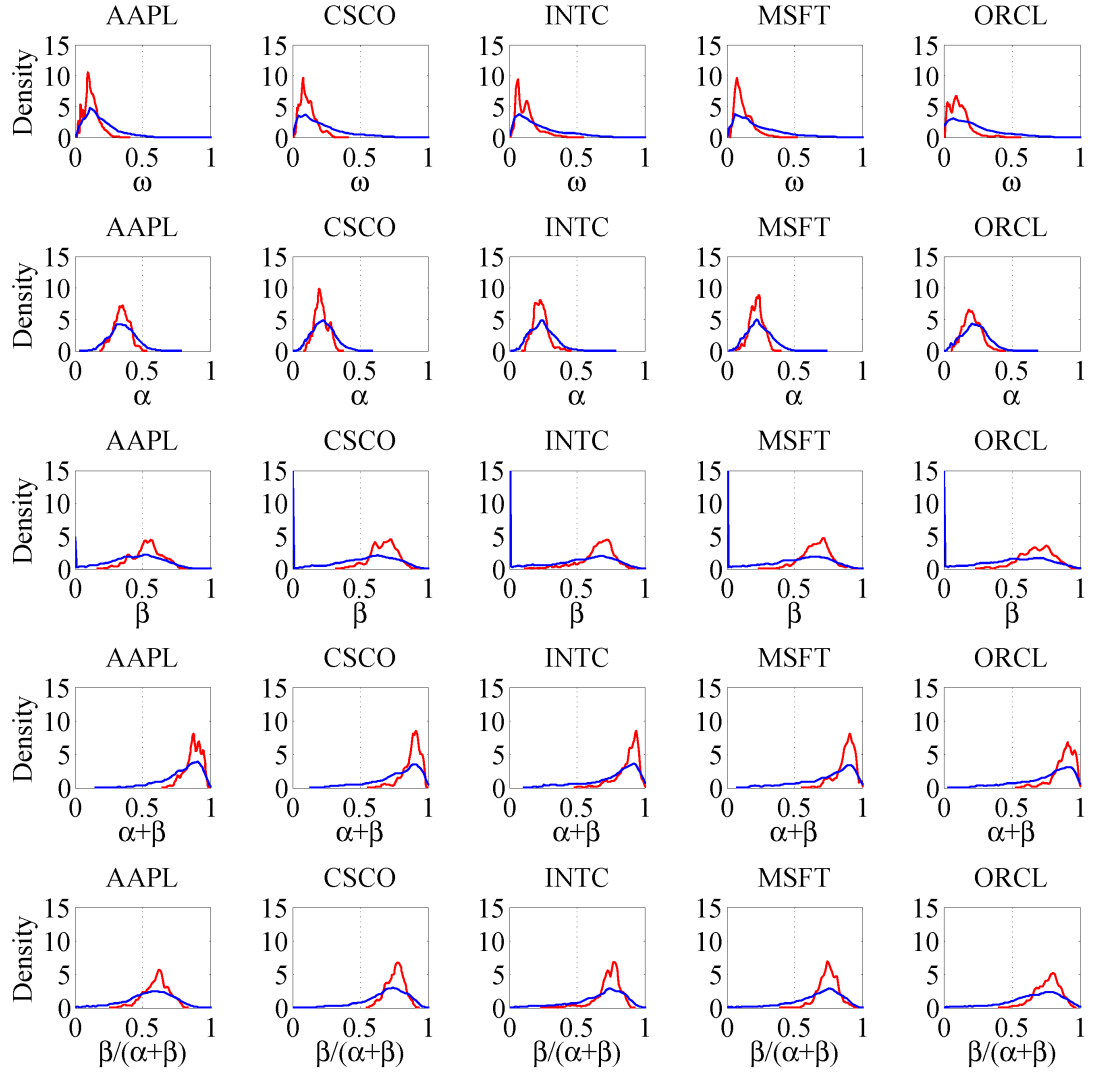


Figure 4.16: Kernel density plots (Gaussian kernel with optimal bandwidth) of estimated $\text{EACD}(1, 1)$ parameters for seasonally adjusted trading volumes over weekly (red) and daily window (blue).

4.2 Local Adaptive Multiplicative Error Models for High-Frequency Forecasts

For all five stocks we observe that estimated MEM parameters and functions thereof substantially vary over time. Parameter variation is intimately related to the lengths of underlying local estimation windows. The MEM parameter variations are reflected in their empirical distributions. Quartiles of the estimated persistence $(\tilde{\alpha} + \tilde{\beta})$ level across all five stocks are summarized in Table 4.10.

Estimation window	EACD(1, 1)			WACD(1, 1)		
	Low	Moderate	High	Low	Moderate	High
1 week	0.85	0.89	0.93	0.82	0.88	0.92
2 days	0.77	0.86	0.92	0.74	0.84	0.91
1 day	0.68	0.82	0.90	0.63	0.79	0.89
3 hours	0.54	0.75	0.88	0.50	0.72	0.87
2 hours	0.45	0.70	0.86	0.42	0.67	0.85
1 hour	0.33	0.58	0.80	0.31	0.57	0.80

Table 4.10: Quartiles of estimated persistence levels $(\tilde{\alpha} + \tilde{\beta})$ for all five stocks at each minute from 22 February to 31 December 2008 (215 trading days) and six lengths of local estimation windows based on EACD and WACD specifications. We label the first quartile as 'low', the second quartile as 'moderate' and the third quartile as 'high'.

Quartiles of the estimated persistence parameter for a given data window size are associated with 'low' (25% quantile), 'moderate' (median) and 'high' (75% quantile) persistence level. One observes that the estimated persistence increases with diminishing variability as the length of the underlying data window includes more observations. This finding suggests that persistence may be more reliably estimated over relatively long data intervals.

Empirical evidence suggests that MEM parameters, their variability and their distribution properties change over time. Results depend upon the length of the underlying data window. Here longer estimation windows increase the parameter precision (i.e. decrease parameter variability) and enlarge the misspecification risk. The modelling bias is thereby inflated as parametric models are fitted over long data stretches. Standard time series analysis selects large estimation windows in order to obtain relatively precise estimates. As the 'price' of assuming parameter homogeneity over long intervals is rather high, the LPA strikes a balance between parameter variability and the modelling bias. As discussed above, the aim of the LPA is to find the longest possible data interval over which parameter homogeneity hypothesis cannot be rejected. The target of the LCP detection test is thus to find the *interval of homogeneity*.

The LCP sequential testing procedure requires as an input a set of critical values that here depend upon the 'true' parameter. Our approach is to select reasonable parameter constellations. 'True' parameter candidates are selected using a data driven approach, i.e. we focus on those parameters that can most likely be estimated from financial data.

According to the computed quartiles of the 'weekly' persistence $(\tilde{\alpha} + \tilde{\beta})$ levels we

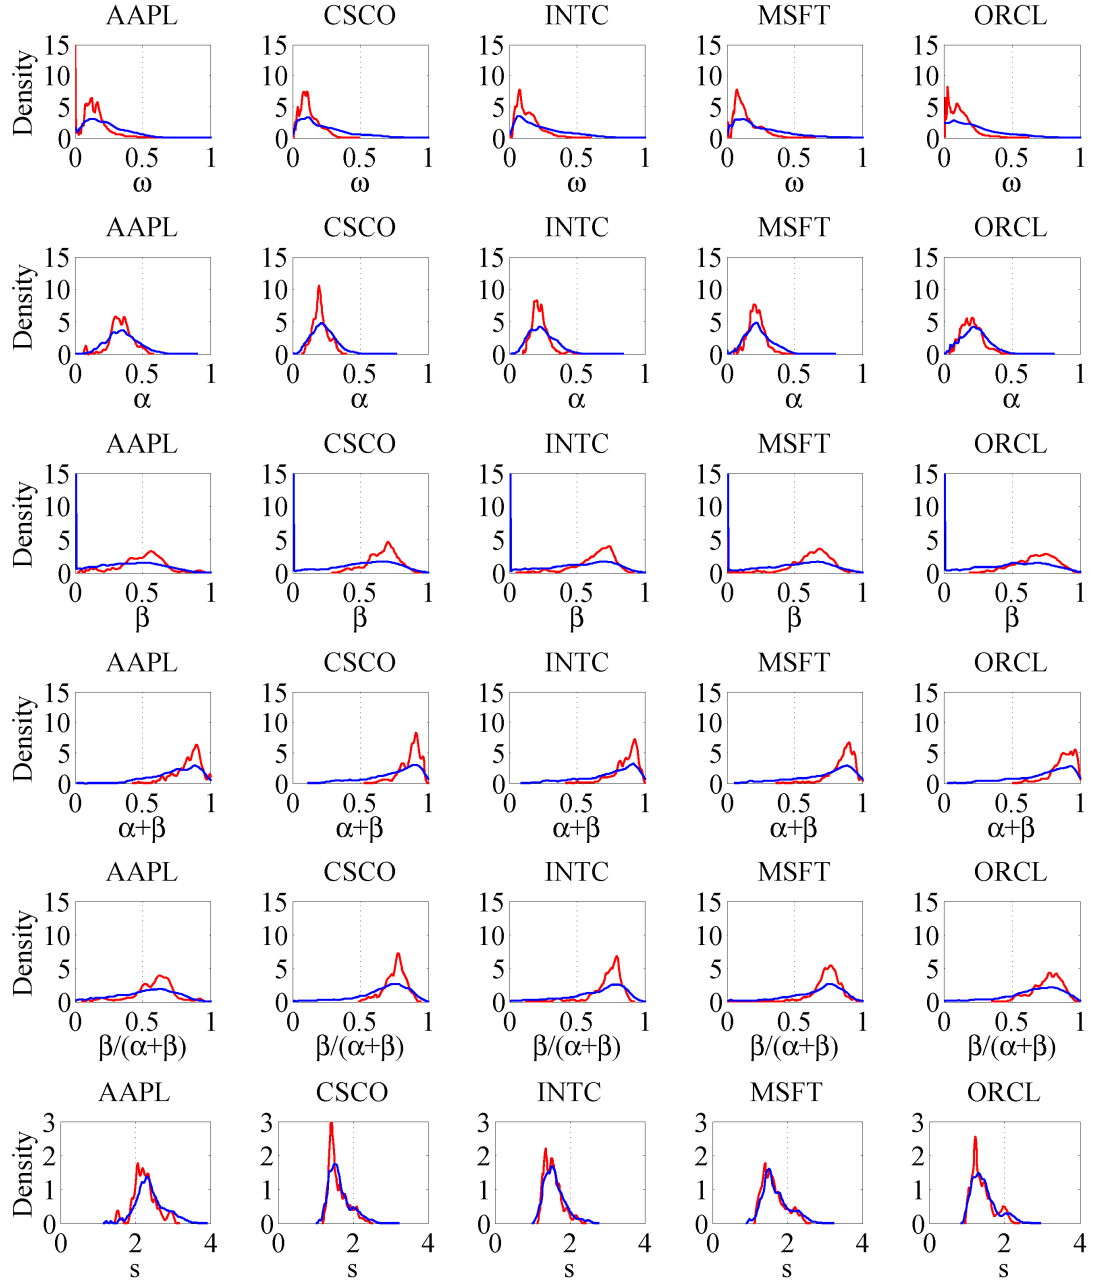


Figure 4.17: Kernel density plots (Gaussian kernel with optimal bandwidth) of estimated WACD(1, 1) parameters for seasonally adjusted trading volumes over weekly (red) and daily window (blue).

4.2 Local Adaptive Multiplicative Error Models for High-Frequency Forecasts

discriminate between three persistence levels (low, medium or high), see the first row of Table 4.10. The range for classification of a parameter vector into the medium persistence group is then $[0.87, 0.91]$ in the case of EACD and $[0.85, 0.90]$ for the WACD model. For example, an EACD parameter vector with estimated persistence of 0.90 would be classified into the medium persistence level group as 0.90 is closer to the median than to the third quartile.

Conditional on a persistence group we then distinguish between different magnitudes of $\tilde{\alpha}$ relative to $\tilde{\beta}$. Similarly to the above procedure we form three groups based on the quartiles of the ratio $\tilde{\beta}/(\tilde{\alpha} + \tilde{\beta})$. The resulting categories are labeled as low, mid or high ratio, see Table 4.11. As an example for classification, consider the case where $\tilde{\alpha} + \tilde{\beta} = 0.20 + 0.71 = 0.91$. The parameter vector is then classified into the medium persistence and high ratio group ($0.20/0.91 = 0.22$ is closer to $0.19/0.89 = 0.21$ than to $0.23/0.89 = 0.26$).

Model	Low Persistence			Moderate Persistence			High Persistence		
	Low	Mid	High	Low	Mid	High	Low	Mid	High
EACD, $\tilde{\alpha}$	0.28	0.22	0.18	0.30	0.23	0.19	0.31	0.24	0.20
EACD, $\tilde{\beta}$	0.56	0.62	0.67	0.59	0.66	0.71	0.62	0.68	0.73
WACD, $\tilde{\alpha}$	0.28	0.21	0.17	0.30	0.23	0.18	0.32	0.24	0.19
WACD, $\tilde{\beta}$	0.54	0.60	0.65	0.58	0.65	0.70	0.60	0.68	0.74

Table 4.11: Quartiles of 774,000 estimated ratios $\tilde{\beta}/(\tilde{\alpha} + \tilde{\beta})$ (based on estimation windows covering 1800 observations) for all five stocks at each minute from 22 February to 31 December 2008 (215 trading days) and both model specifications (EACD and WACD) conditional on the persistence level (low, moderate or high). We label the first quartile as 'low', the second quartile as 'mid' and the third quartile as 'high'. The shape parameter for the WACD model equals the median value in all cases ($\tilde{s} = 1.57$).

Our classification procedure enables us to identify nine groups of parameter constellations that are used for the simulation of critical values, see Section 2.3.4.

4.2.2 Adaptive Estimation

In applying the local MEM at each time point (our sample covers 77,400 minutes) we select the curve of critical values according to the previously introduced classification scheme. At every time point we therefore compute the (past) 'weekly' persistence level and the value of the ratio $\tilde{\beta}/(\tilde{\alpha} + \tilde{\beta})$. In the example above ($\tilde{\alpha} + \tilde{\beta} = 0.20 + 0.71 = 0.91$) our selected curve corresponds to the medium-low (persistence-ratio) group. If for example at some other point we have $\tilde{\alpha} = 0.32$ and $\tilde{\beta} = 0.53$, then the selected curve comes from the low-low parameter vector group.

The adaptive estimate $\hat{\theta}$ equals to the QMLE at the *interval of homogeneity*, see Section 2.3.3. The adaptive choice of intervals for all five stocks on 2 February 2008 is shown in Figure 4.18.

4 Applications




Figure 4.18: Estimated length of *intervals of homogeneity* n_k (in hours) for seasonally adjusted one-minute cumulative trading volumes of selected companies in case of a modest ($r = 0.5$, blue) and conservative ($r = 1$, red) modelling risk level. We use the interval scheme with $K = 13$ and $\rho = 0.25$. Underlying model: EACD(1,1). NASDAQ trading on 22 February 2008.

The results are here presented for the local MEM with the underlying EACD model specification, parameters $\rho = 0.25$ and $K = 13$ as well as for both risk scenarios, i.e. the modest and the conservative risk case with $r = 0.5$ and $r = 1$, respectively. The adaptive estimation selects an estimation interval between 1.5 and 3.5 hours during the selected day given the modest risk case. According to the conservative risk case, the adaptive estimate would be based on a interval length of around 4 hours. In the modest risk case one selects shorter intervals than in the conservative case.

4.2.3 Empirical Findings

Understanding of economic risk plays an important role in practice. A trader aims to control the price or the volume risk, i.e. the exposure of its portfolio (wealth) to market changes, such as, e.g., the price volatility dynamics or the evolution of the traded quantity. A local MEM successfully deals with both kinds of risks. For example, using a GARCH framework, one better models the dynamics of price volatility than a 'standard' approach, see, e.g., Čížek et al. [2009]. Applying the framework to MEM for intra-day volume modelling, we control the volume risk. Our benefits arise while outperforming a 'standard' approach with ad-hoc selected estimation windows.

The LPA is applied to periodicity adjusted 1-min trading volumes for all five stocks (AAPL, CSCO, INTC, MSFT and ORCL) at each minute from 22 February to 31 December 2008 (215 trading days, 77400 trading minutes). We focus on the EACD and the WACD model, interval schemes with $K = 8$ and $K = 13$, two risk levels (modest, $r = 0.5$, and conservative, $r = 1$) as well as on two significance levels, $\rho \in \{0.25, 0.50\}$. The distributions of the length of the *interval of homogeneity* are shown in Figure 4.19.

A modest risk approach ($r = 0.5$) leads to shorter data intervals than a conservative risk case ($r = 1$). If a trader desires a more precise estimation procedure, the advice is to take (relatively) longer estimation intervals, such as, e.g., 4-5 hours. An estimation that favors a smaller modelling bias is obtained by setting $r = 0.5$, i.e., we suggest to estimate the model over (relatively) shorter intervals, for example, including at most 2-3

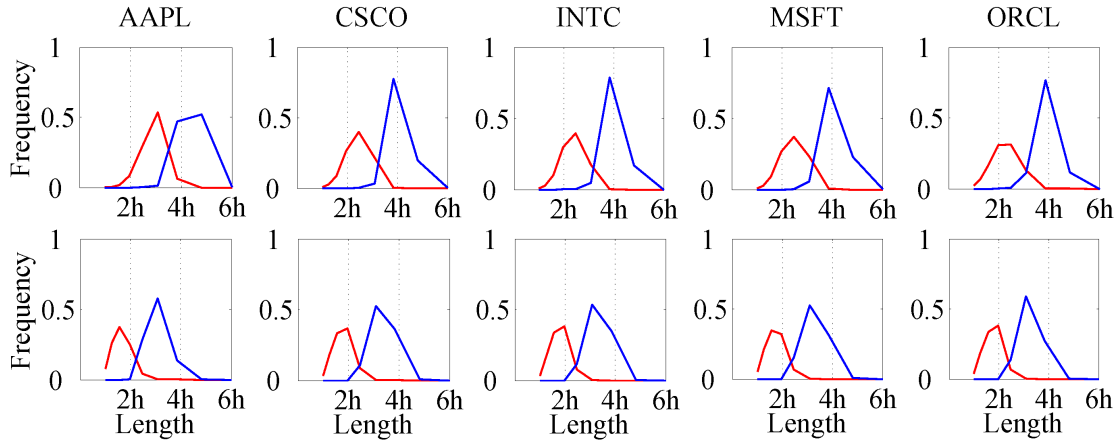


Figure 4.19: Distribution of estimated interval length n_k (in hours) for seasonally adjusted trading volumes of selected companies in case of modest ($r = 0.5$, red) and conservative modelling risk ($r = 1$, blue), using an EACD (upper panel) and a WACD model (lower panel) from 22 February to 31 December 2008 (215 trading days). We select 13 estimation windows based on significance level $\rho = 0.25$.

hours of data.

Applying a local MEM, the maximum length of the *interval of homogeneity* is 6 hours in all cases. Recall, the longest investigated interval of homogeneity I_K includes 1800 observations (i.e. 30 trading hours or 1 week of data). The 'weekly' estimation is certainly not appropriate for modelling of high-frequency data. Since the results are quite robust across different schemes, we recommend a modelling horizon of up to one trading day of data. The 'right' length is varying across time, see, e.g., 4.18. An ad hoc selected (fixed) interval is likewise not appropriate in financial time series modelling.

Concerning on the model complexity, one observes that a relatively longer (shorter) estimation intervals should be used while using the local EACD (WACD) model specification. This finding is possibly due to the variability of the shape parameter of the WACD model. The results are quite robust across other employed scenarios.

Focusing on the evolution of the *interval of homogeneity* over a course of a typical trading day, we show averages of the selected *intervals of homogeneity* at each minute through our sample, see, e.g., Figure 4.20. After seasonal adjustment of the data, one sees slightly shorter intervals in the opening and before closure. We attribute this finding to the fact that estimation windows in the morning hours include (significant) portion of data from the previous day. A possible overnight effect therefore changes the parameter dynamics. Induced trading activity by the end of the day similarly influences the parameter evolution.




Figure 4.20: Average estimated interval length n_k (in hours) over the course of a trading day for seasonally adjusted trading volumes of selected companies in case of modest ($r = 0.5$, red) and conservative modelling risk ($r = 1$, blue), using an EACD (upper panel) and a WACD model (lower panel) from 22 February to 31 December 2008 (215 trading days). We select 13 based on significance level $\rho = 0.25$.

4.2.4 Forecasting Trading Volumes

Economic benefits of the local MEM arise with better volume risk control. Presented evidence on the (in)homogeneity of the MEM parameters suggests that a LPA may yield better volume forecasts in the short run than a 'standard' benchmark approach in a rolling window exercise. Note that in the 'standard' modelling, the length of the estimation window is fixed a priori.

At each trading minute from 22 February to 22 December 2008 (210 trading days, 75600 minutes) the trading volume is predicted over all horizons $h = 1, 2, \dots, 60$ min during the next hour. Multi-step-ahead forecasts of seasonally adjusted volume are computed using a recursive scheme using the currently estimated MEM parameters and initialized based on the observed data that stem from the current estimation window. Following the multiplicative data structure (3.5), we multiply the seasonally adjusted volume forecasts with the corresponding estimated (past) seasonal factors.

The local MEM is based on the LPA approach with $r \in \{0.5, 1\}$ and $\rho \in \{0.25, 0.5\}$. We consider two modelling specifications (EACD and WACD) as well as two interval selection schemes ($K = 8$ and $K = 13$). Denote the h -step LPA volume prediction by \hat{y}_{i+h} and the resulting prediction error $\hat{\varepsilon}_{i+h} = \check{y}_{i+h} - \hat{y}_{i+h}$, where \check{y}_{i+h} represents the observed trading volume.

Current econometric literature suggests to use several months of data in modelling and forecasting of high-frequency data. Here the 'standard' approach is based on a fixed estimation window covering (only) 1800 observations (i.e. 30 trading hours or 1 week of data). To make the 'standard' approach even more competitive we set the interval to

4.2 Local Adaptive Multiplicative Error Models for High-Frequency Forecasts

cover just 360 observations (6 hours or 1 trading day), see Section 4.2.3. The predictions are denoted by \tilde{y}_{i+h} and the corresponding prediction errors by $\tilde{\varepsilon}_{i+h} = \check{y}_{i+h} - \tilde{y}_{i+h}$.

With the idea of regressing the observed trading volume on the predicted volume time series, we test for the unbiasedness and the efficiency of the forecasting methods, see, e.g., Mincer and Zarnowitz [1969]. We consider the following three linear regression models at fixed forecasting horizon h with unknown parameters $\hat{\gamma}$'s, $\tilde{\gamma}$'s and η 's to be estimated

$$\check{y}_{i+h} = \hat{\rho}_0 + \hat{\rho}_1 \hat{y}_{i+h} + \hat{\varepsilon}_{i+h} \quad (4.6)$$

$$\check{y}_{i+h} = \tilde{\rho}_0 + \tilde{\rho}_1 \tilde{y}_{i+h} + \tilde{\varepsilon}_{i+h} \quad (4.7)$$

$$\check{y}_{i+h} = \eta_0 + \eta_1 \hat{y}_{i+h} + \eta_2 \tilde{y}_{i+h} + \varepsilon_{i+h}, \quad (4.8)$$

where $\hat{\varepsilon}$, $\tilde{\varepsilon}$ and ε represent white noise processes. A forecasting method is said to be unbiased if the regression constant in models (4.6) and (4.7) is not significant, and efficient if the corresponding slope equals one. Moreover, the coefficient of determination R^2 reflects the strength of correlation between forecast and outcome. If the LPA technique would cover all forecasting information, the parameters from the encompassing regression (4.8) equal to $\eta_1 = 1$ and $\eta_2 = 0$.

The strength of correlation between forecast and outcome favors the LPA technique, see, e.g., Figure 4.21. The LPA predictive power dominates across all stocks, except for the INTC stock given modest risk case ($r = 0.5$). Low risk power r leads to stronger relationship in the short run. Best forecasting results are therefore expected over relatively short forecasting horizons.




Figure 4.21: Coefficient of determination R^2 of forecasting regressions (4.6) and (4.7) computed at fixed horizon $h = 1, \dots, 60$ for the EACD model from 22 February to 22 December 2008 (210 trading days). Results are shown for the LPA technique with 'significance' level $\rho = 0.25$ for two risk levels, i.e., $r = 0.5$ (solid red line) and $r = 1$ (dashed red line), as well as for two specifications of the standard method, i.e., ad hoc selected estimation window of 360 (solid blue line) and 1800 observations (dashed blue line).

Both forecasting methods share similar unbiasedness and efficiency results allowing us to compare their predictive accuracy. One observes that the probability of rejecting the unbiasedness or efficiency hypothesis increases with the forecasting horizon, see, e.g., a brief overview of test results in Table 4.12. The total (average) number of non-rejections

4 Applications

across all stocks is quite similar. In a small number of cases for the AAPL stock the forecasts are statistically unbiased (regression constant equals zero), as opposed to the CSCO and ORCL stocks where the same null is not rejected too often.

	EACD					WACD				
	AAPL	CSCO	INTC	MSFT	ORCL	AAPL	CSCO	INTC	MSFT	ORCL
Unbiasedness										
$r = 0.5, \rho = 0.25$	0	6	6	5	11	0	3	0	3	4
$r = 0.5, \rho = 0.50$	0	6	6	5	12	0	3	0	3	4
$r = 1.0, \rho = 0.25$	0	51	19	0	23	0	20	15	40	39
$r = 1.0, \rho = 0.50$	0	51	19	0	25	0	20	16	43	41
Stan., 1 week	0	0	0	0	59	0	0	0	16	58
Stan., 1 day	0	41	26	16	13	0	44	26	8	45
Efficiency										
$r = 0.5, \rho = 0.25$	3	3	4	4	20	1	1	1	2	4
$r = 0.5, \rho = 0.50$	3	3	4	4	21	2	1	1	2	4
$r = 1.0, \rho = 0.25$	0	36	18	0	0	0	4	14	38	47
$r = 1.0, \rho = 0.50$	0	37	19	0	0	0	4	15	38	46
Stan., 1 week	0	0	0	0	59	0	0	0	0	59
Stan., 1 day	0	46	40	12	0	0	35	23	23	25

Table 4.12: Number of non-rejections of the unbiasedness or the efficiency null hypothesis given horizon $h = 1, \dots, 60$ of five large companies traded at the NASDAQ from 22 February to 22 December 2008 (210 trading days) using an EACD(1,1) and a WACD(1,1) model at a 5% significance level. We specify four tuning parameter constellations when using the LPA technique and two ad hoc selected window lengths when employing the standard method (1 week or 1 day). Maximum number of non-rejections is 60 in each case.

After estimating the encompassing regressions (4.8) we observe that the LPA technique statistically covers all forecasting information at least once in every case, see, e.g., Table 4.13. This finding is most likely due to the outstanding short-term forecasting performance of the LPA technique, i.e. the 'full forecasting information' hypothesis is not rejected for short forecasting horizons. Best results are interestingly observed for the EACD case.

Results show that the local MEM outperforms the 'standard' approach with ad hoc fixed estimation window. Now we compare both methods using a Diebold and Mariano [1995] testing framework. Define the loss differential d_h between the squared prediction errors of both methods over a fixed horizon h and for n observations as $d_h = \{d_{i+h}\}_{i=1}^n$, with $d_{i+h} = \hat{\varepsilon}_{i+h}^2 - \tilde{\varepsilon}_{i+h}^2$. Testing whether the local MEM yields qualitatively lower predictions errors is based on the test statistic

$$T_{ST,h} = \left\{ \sum_{i=1}^n \mathbf{I}(d_{i+h} > 0) - 0.5n \right\} / \sqrt{0.25n}. \quad (4.9)$$

The statistic is approximately $N(0, 1)$ distributed. Our sample covers $n = 75600$ trading minutes (corresponding to 210 trading days). Quantitative forecasting superiority of the

4.2 Local Adaptive Multiplicative Error Models for High-Frequency Forecasts

	EACD					WACD				
	AAPL	CSCO	INTC	MSFT	ORCL	AAPL	CSCO	INTC	MSFT	ORCL
$r = 0.5, \rho = 0.25$	5	3	2	2	6	1	3	1	1	0
$r = 0.5, \rho = 0.25$	5	3	2	2	6	1	2	1	1	0
$r = 0.5, \rho = 0.25$	11	23	12	0	11	1	9	5	12	8
$r = 0.5, \rho = 0.25$	11	24	12	0	12	1	9	6	12	8

Table 4.13: Number of non-rejections of the 'forecasting information coverage' null hypothesis ($H_0 : \eta_1 = 1$ and $H_0 : \eta_2 = 0$) given horizon $h = 1, \dots, 60$ of five large companies traded at the NASDAQ from 22 February to 22 December 2008 (210 trading days) using an EACD(1,1) and a WACD(1,1) model with a significance level of 5%. We specify four tuning parameter constellations when comparing the LPA technique to the standard method with an ad hoc selected window length of 1800 observations.

local MEM is performed by testing the null $H_0 : E[d_h] = 0$. Corresponding test statistic is given by

$$T_{DM,h} = \bar{d}_h / \sqrt{2\pi \hat{f}_{d_h}(0) / n} \xrightarrow{\mathcal{L}} N(0, 1). \quad (4.10)$$

The average loss differential is computed by $\bar{d}_h = n^{-1} \sum_{i=1}^n d_{i+h}$. A consistent estimate of the spectral density of the loss differential at frequency zero, $\hat{f}_{d_h}(0)$, may be computed as

$$\hat{f}_{d_h}(0) = (2\pi)^{-1} \sum_{m=-(h-1)}^{h-1} \hat{\gamma}_{d_h}(m), \quad (4.11)$$

$$\hat{\gamma}_{d_h}(m) = n^{-1} \sum_{i=|m|+1}^n (d_{i+h} - \bar{d}_h) (d_{i+h-|m|} - \bar{d}_h), \quad (4.12)$$

see, e.g., Diebold and Mariano [1995]. We display the Diebold-Mariano test statistics $T_{DM,h}$ against the forecasting horizon h in Figure 4.22. The underlying LPA is based on the EACD model with significance level $\rho = 0.25$.

Quantitatively, our local MEM provides smaller forecasting errors than a 'standard' approach, i.e., the fixed-window based forecast is worse than the LPA in any case. The 'standard' approach performs poorly if the data interval includes many observations (here one week of data). The data intervals are obviously too large and thus too restrictive in practice. This misspecification leads to significantly worse predictions, see, e.g., Figure 4.22.

Qualitatively, the local MEM produces smaller (squared) forecasting errors in all cases, see, e.g., Table 4.14 that includes the test statistics $T_{ST,h}$ given in (4.9). The largest (i.e., least negative) statistics across all 60 forecasting horizons support our previous findings. The prediction accuracy is robust against the underlying LPA tuning parameters.

Even after reducing the number of observations to 360 (i.e., 1 day of data) while using the 'standard' method according to our suggestion, the LPA statistically outperforms




Figure 4.22: Test statistic $T_{DM,h}$ across all 60 forecasting horizons for five large companies traded at NASDAQ from 22 February to 22 December 2008 (210 trading days). The red curve depicts the statistic based on a test of the LPA against a fixed-window scheme using 360 observations (6 trading hours). The blue curve depicts the statistic based on a test of the LPA against a fixed-window scheme using 1800 observations (30 trading hours). The upper panel shows the results for the 'modest risk case' ($r = 0.5$) and the lower panel shows the results for the 'conservative risk case' ($r = 1$) given a significance level of $\rho = 0.25$.

4.2 Local Adaptive Multiplicative Error Models for High-Frequency Forecasts

	EACD					WACD				
	AAPL	CSCO	INTC	MSFT	ORCL	AAPL	CSCO	INTC	MSFT	ORCL
1 week										
$r = 0.5, \rho = 0.25$	-38.9	-28.6	-24.1	-33.8	-31.4	-22.6	-25.7	-20.2	-26.7	-26.6
$r = 0.5, \rho = 0.50$	-38.7	-28.7	-24.2	-33.8	-31.4	-22.7	-25.5	-20.3	-26.7	-26.6
$r = 1.0, \rho = 0.25$	-40.5	-31.4	-23.3	-39.1	-32.8	-27.9	-30.8	-21.5	-31.3	-29.8
$r = 1.0, \rho = 0.50$	-40.4	-31.3	-23.3	-39.0	-32.9	-28.1	-30.8	-21.5	-31.5	-29.7
1 day										
$r = 0.5, \rho = 0.25$	-10.8	-6.0	-13.1	-5.7	-15.1	-6.4	-3.5	-6.1	-4.9	-12.6
$r = 0.5, \rho = 0.50$	-10.6	-6.0	-12.8	-5.5	-15.0	-6.3	-3.2	-6.2	-4.8	-12.7
$r = 1.0, \rho = 0.25$	-6.9	-8.6	-8.7	-4.4	-12.9	-4.1	-5.1	-6.5	-4.2	-11.5
$r = 1.0, \rho = 0.50$	-7.1	-8.6	-8.8	-4.4	-13.0	-3.9	-5.2	-6.5	-4.1	-11.4

Table 4.14: Largest (in absolute terms) test statistic $T_{ST,h}$ across all 60 forecasting horizons as well as EACD and WACD specifications for five large companies traded at NASDAQ from 22 February to 22 December 2008 (210 trading days). We compare LPA-implied forecasts with those based on rolling windows using a priori fixed lengths of one week and one day, respectively. Negative values indicate lower squared prediction errors resulting from the LPA. According to the Diebold-Mariano test (4.10), the average loss differential is significantly negative in all cases (significance level 5%).

the 'standard' approach. One sees that the fixed-window setting with shorter intervals significantly outperforms the setting with 1 week of data, as shown in Figure 4.22 (quantitative performance) and Table 4.14 (qualitative accuracy). A 'standard' approach with an ad hoc selected (relatively) long estimation intervals fails in short-term forecasting.

Our local MEM predicts trading volumes fairly well over short horizons. The strongest overperformance is visible over horizons between two and four minutes, see, e.g., Figure 4.23. Depicting the ratio of RMSPEs

$$\sqrt{n^{-1} \sum_{i=1}^n \hat{\varepsilon}_{i+h}^2} / \sqrt{n^{-1} \sum_{i=1}^n \tilde{\varepsilon}_{i+h}^2},$$

in dependence of the length of the forecasting horizon, LPA again exhibit superior forecasting power. The prediction performance shrinks slowly with the horizon. The strongest advantage is achieved over horizons of up to 20 minutes. One still expects financial gains implied by the local MEM in trading strategies within the next one hour.

A better forecasting performance of the LPA method is documented over time, see, e.g., Figure 4.24. During most of the trading days in 2008 the ratio lies below one. The striking outperformance is visible over the last months covering the financial crisis. Since market distress periods affect the trading activity, the MEM parameters may change more often than during a calm (summer) period. The results again do not critically depend on the LPA modelling setup, demonstrating again the LPA strength for understanding high-frequency dynamics.

Forecasting performance per intraday time favor the LPA method as well. The stan-

4 Applications




Figure 4.23: Ratio between the RMSPEs of the LPA and of a fixed-window approach (covering 6 trading hours) over the sample from 22 February to 22 December 2008 (210 trading days). Upper panel: EACD model, lower panel: WACD model.




Figure 4.24: Ratio between the RMSPEs of the LPA and of a fixed-window approach (covering 6 trading hours) over the sample from 22 February to 22 December 2008 (210 trading days). Upper panel: Results for underlying (local) EACD model. Lower panel: Results for underlying (local) WACD model.

standard approach is clearly outperformed, for illustration see Figure 4.25. The best LPA overperformance is interestingly achieved at the end of a typical trading day, usually after 14:00, irrespectively upon the tuning parameters. The local MEM exhibits again better performance than a fixed-window approach.




Figure 4.25: Ratio between the RMSPEs of the LPA method and the RMSPE of the standard approach over the course a typical trading day using an EACD (upper panel) and WACD (lower panel) model from 22 February to 22 December 2008 (210 trading days).

4.2.5 Financial Applications

We illustrate the usefulness of the local MEM by introducing a trading strategy based on predicted volume series using the LPA and the 'standard' method. The example shows how a trader may control volume risk. A more sophisticated strategy would include predictions of the volume weighted average price (VWAP), see, e.g., Berkowitz et al. [1988] and Fuh et al. [2010].

EXAMPLE 4. (Trading Strategy)

Suppose that a trader decides to buy (sell) certain number of shares at each minute from 10:00-15:00. There are two methods for predicting the future one-minute trading volume at each observation i :

- (i) 'Standard' method - predicted volume \tilde{y}_{i+1} ,
- (ii) Local MEM - predicted volume \hat{y}_{i+1} .

Suppose that technical analysis suggest the trader to buy (sell) stocks when the adaptively predicted volume is larger (smaller) than the standard method volume. Let fur-

4 Applications

thermore the number of traded shares equals the difference $\hat{y}_{t+1} - \tilde{y}_{t+1}$. Cash flow is calculated at each trading hour (assumed transaction fees and taxes are set to be 10%).

Empirical evidence shows that the investment strategy provides positive financial benefits over the entire year, see Figure 4.26. This is in particular true for AAPL, CSCO, MSFT and their joint (equally weighted) portfolio. Except for May and September, the strategy achieves moderate/high financial gains. The results confirm that the local MEM outperforms the 'standard' method and can be used in modelling of volume risk.

4.2 Local Adaptive Multiplicative Error Models for High-Frequency Forecasts

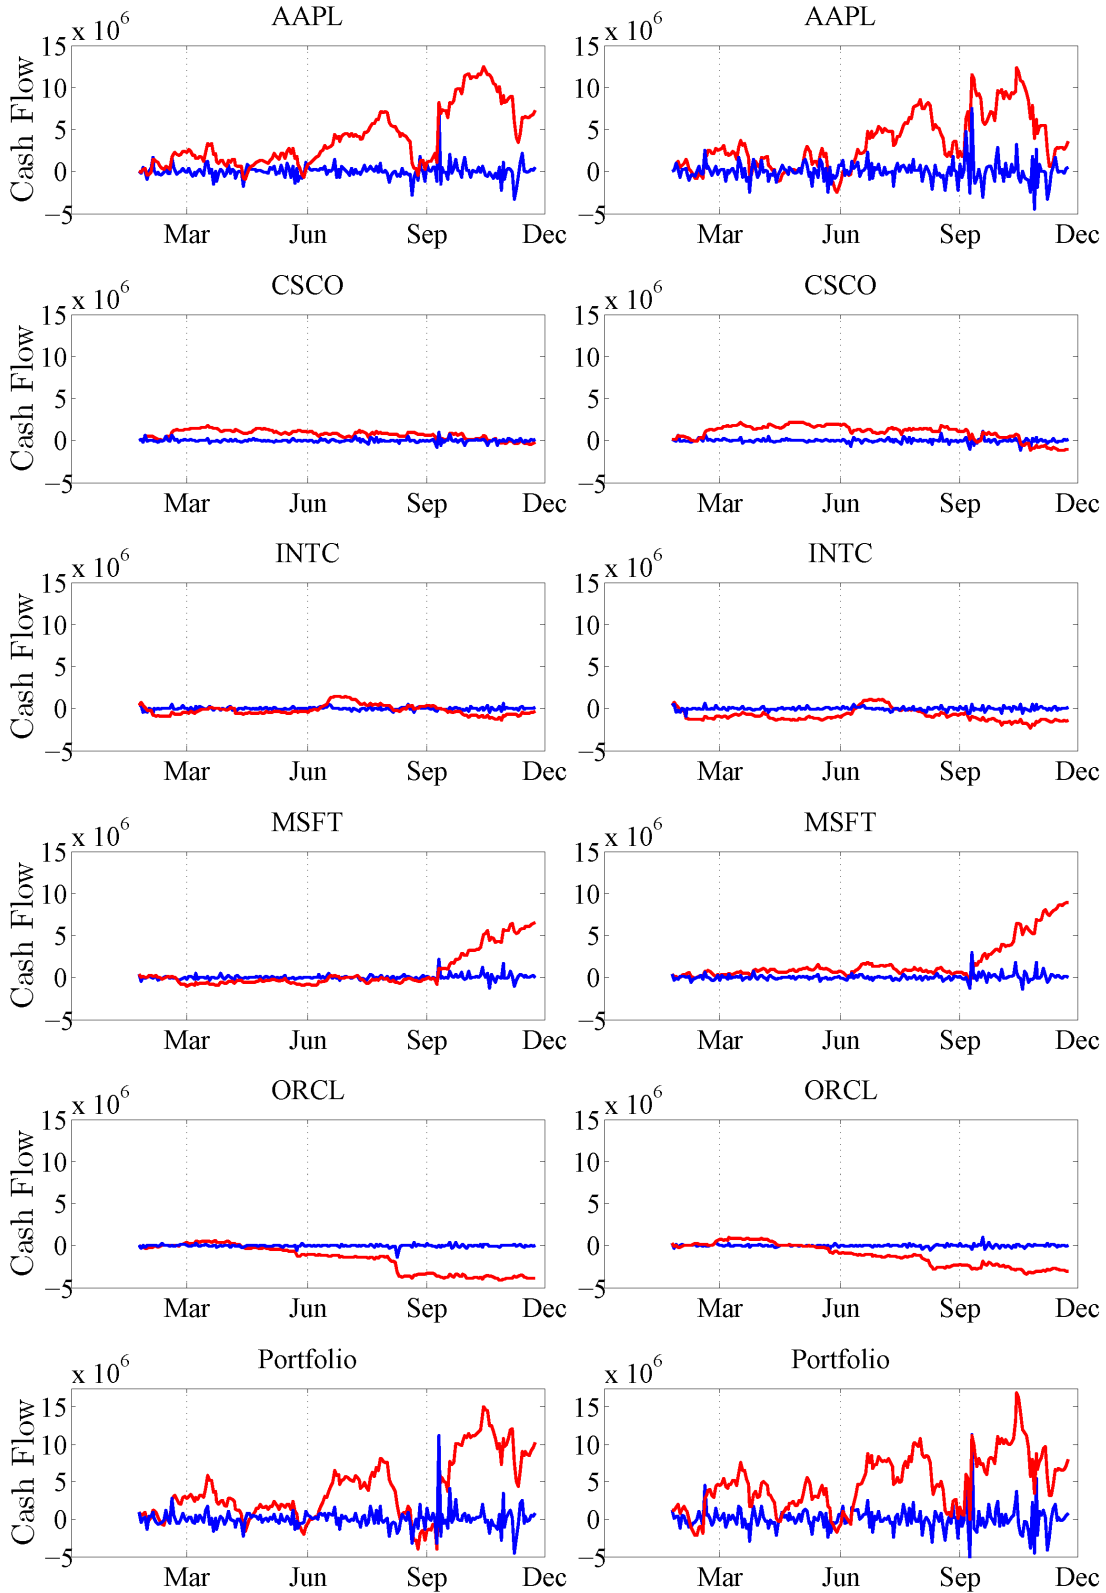


Figure 4.26: Daily cash flow in USD (blue) and cumulated daily cash flow in USD (red) of the investment strategy. The investor uses an EACD (left panel) and WACD model (right panel) from 22 February to 22 December 2008 (210 trading days).

4.3 Cross Country Evidence for the EPK Paradox

Understanding the market microstructure plays an important role in economics. Focusing on asset pricing, our primary goal is to study the preference dynamics and to provide statistical evidence for the existence of the EPK paradox on stock markets. The economic framework of a state-dependent utility is now applied to market data at six worldwide largest markets (AUS, GER, JPN, SUI, UK and US) in the period from 1 January 1990 until 31 May 2012.

4.3.1 Modelling Setup

Recall the pricing kernel (PK) specification (2.24) under state-dependent preferences

$$\mathcal{K}_\theta(r_{m,t+1}) = \beta_1 r_{m,t+1}^{-1} \mathbf{I}\{r_{m,t+1} \in [0, x)\} + \beta_2 r_{m,t+1}^{-1} \mathbf{I}\{r_{m,t+1} \in [x, \infty)\}, \quad (4.13)$$

with simple market gross return $r_{m,t+1}$, parameter vector $\theta = (\beta_1, \beta_2)^\top$ and fixed unknown reference point x . This specification allows us to study the estimation quality and the parameter dynamics. It is a suitable framework to provide evidence for the non-monotonicity of the empirical pricing kernel on stock markets.

Consider the following three scenarios:

- (a) State-dependent and unconstrained estimation with $\beta_1, \beta_2 > 0$,
- (b) State-dependent and constrained estimation with $\beta_1 > \beta_2 > 0$ and
- (c) State-independent estimation with $\beta_1 = \beta_2 = \beta = 0$.

In the implementation of the above scenarios we use two estimation techniques, namely the iterated GMM and the GMM based on the Hansen-Jagannathan (HJ) weighting matrix, see Section 2.5.1. The estimation techniques treat the unknown change point x as endogenous. Using a grid search method we optimize the objective function simultaneously in θ and x in all cases. Three different estimation window lengths are used, i.e., we select $n \in \{250 \text{ (1 year)}, 500 \text{ (2 years)}, 1250 \text{ (5 years)}\}$.

4.3.2 Estimation Quality

The in-sample performance of both estimation techniques is compared based on the optimal value of the objective function. It is sufficient to focus on the average reported objective function value for a fixed scenario and a given stock market. The resulting estimation quality indicators (averages) are given in Tables 4.15, 4.16 and 4.17 over the estimation windows including $n = 250$, $n = 500$ and $n = 1250$ observations, respectively.

The general observation is that cases allowing for non-monotonicity of the pricing kernel, i.e., scenarios (a) and (b), favor the iterated GMM estimation procedure. The specification (c) leads to a monotone decreasing pricing kernel and is (on average) best estimated by the GMM technique with the HJ weighting matrix. One furthermore observes that in all cases the (average) optimal objective function value decreases with

4.3 Cross Country Evidence for the EPK Paradox

	GMM with HJ matrix			Iterated GMM		
	(a)	(b)	(c)	(a)	(b)	(c)
AUS	4.05	4.25	4.45	3.55	3.95	4.55
GER	2.96	3.09	3.34	2.66	2.96	3.39
JPN	2.44	2.55	2.77	2.15	2.38	2.80
SUI	3.35	3.48	3.72	2.96	3.21	3.79
UK	2.63	2.73	2.91	2.40	2.58	2.98
US	3.33	3.50	3.68	3.01	3.41	3.76

Table 4.15: Average optimal objective function value across six worldwide largest markets for two competing estimation techniques and three scenarios: (a) $\beta_1, \beta_2 > 0$, (b) $\beta_1 > \beta_2 > 0$ and (c) $\beta_1 = \beta_2 = \beta = 0$. The estimation window covers $n = 250$ observations (1 year).

	GMM with HJ matrix			Iterated GMM		
	(a)	(b)	(c)	(a)	(b)	(c)
AUS	1.19	1.31	1.39	1.08	1.25	1.42
GER	0.85	0.91	1.00	0.81	0.90	1.01
JPN	0.68	0.71	0.80	0.66	0.71	0.81
SUI	1.01	1.06	1.15	0.88	0.96	1.17
UK	0.84	0.89	0.95	0.79	0.86	0.97
US	0.91	0.96	1.01	0.84	0.95	1.03

Table 4.16: Average optimal objective function value across six worldwide largest markets for two competing estimation techniques and three scenarios: (a) $\beta_1, \beta_2 > 0$, (b) $\beta_1 > \beta_2 > 0$ and (c) $\beta_1 = \beta_2 = \beta = 0$. The estimation window covers $n = 500$ observations (2 years).

	GMM with HJ matrix			Iterated GMM		
	(a)	(b)	(c)	(a)	(b)	(c)
AUS	0.34	0.38	0.39	0.29	0.38	0.40
GER	0.25	0.26	0.30	0.23	0.25	0.31
JPN	0.17	0.18	0.22	0.16	0.18	0.22
SUI	0.25	0.26	0.29	0.21	0.24	0.29
UK	0.28	0.29	0.32	0.26	0.28	0.32
US	0.27	0.29	0.31	0.25	0.29	0.31

Table 4.17: Average optimal objective function value across six worldwide largest markets for two competing estimation techniques and three scenarios: (a) $\beta_1, \beta_2 > 0$, (b) $\beta_1 > \beta_2 > 0$ and (c) $\beta_1 = \beta_2 = \beta$. The estimation window covers $n = 1250$ observations (5 years).

the sample size. This implies that the parameter precision increases with the number of observations.

4.3.3 Parameter Dynamics and Reference Point Analysis

The estimated parameters for cases (b) and (c) are illustrated in Figures 4.27 and 4.28, respectively. Parameter estimates indicate that preferences change over time, even in the 'standard' case (c) with a monotone decreasing pricing kernel. Results clearly provide (descriptive) evidence for the existence of a pricing kernel paradox as one observes more trading days with different estimated parameters β_1 and β_2 . The 'standard' case is thus of limited use in practice.

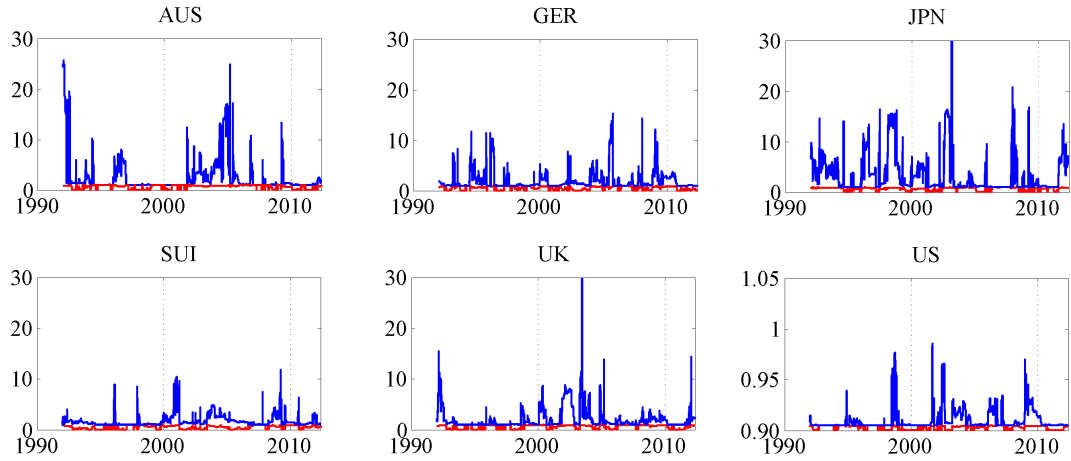


Figure 4.27: Time series of the estimated parameters β_1 (red) and β_2 (blue) across six worldwide largest stock markets. We employ the iterated GMM estimation technique with $n = 500$ (2 years).

The distributions of optimal reference points shows interesting patterns across the stock markets, see, e.g., Figure 4.29. One selects values slightly above one with high probability. This results supports empirical evidence from the option market literature. There are two groups of countries that share similar distribution patterns. Countries with relatively high percentage of optimal reference points lying just above one (AUS, SUI and UK) potentially belong to the same group. For the remaining countries (GER, JPN and US) the optimal reference points near one are located below and above one in roughly similar proportions.

4.3 Cross Country Evidence for the EPK Paradox




Figure 4.28: Time series of the estimated parameter β across six worldwide largest stock markets. We employ the GMM estimation technique with the HJ weighting matrix and with $n = 500$ (2 years).

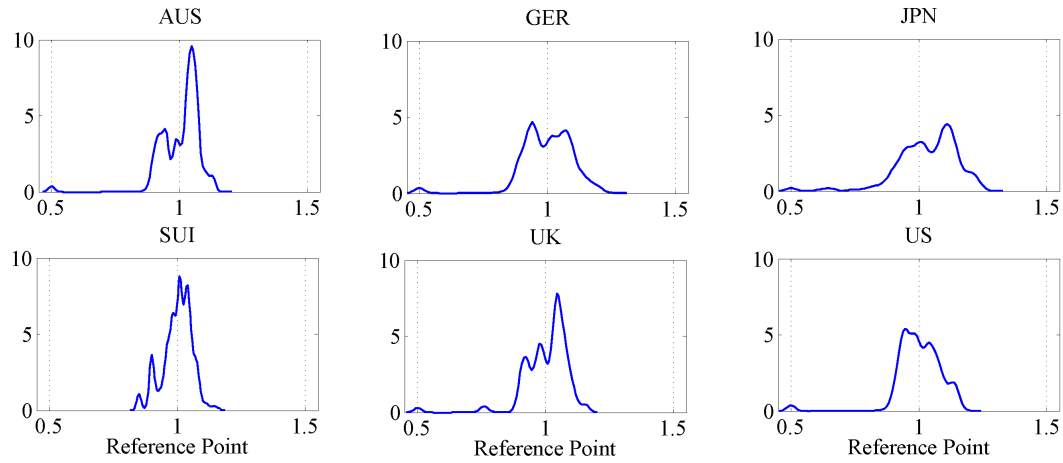


Figure 4.29: Kernel density plots (Gaussian kernel with optimal bandwidth) of optimal reference point x . We employ the iterated GMM estimation technique with $n = 500$ (2 years).

4.3.4 Empirical Pricing Kernels across Stock Markets

Results of the D -test statistically indicate the existence of the EPK paradox on stock markets, see, e.g., Table 4.18. This evidence is furthermore supported by a graphical representation of the empirical pricing kernel fitted to the average parameter values and average optimal reference point at a selected market, see, e.g., Figure 4.30. This findings suggest that modern statistical techniques lead to a better understanding of the market microstructure.

	Iterated GMM			GMM with HJ matrix		
	1 year	2 years	5 years	1 year	2 years	5 years
AUS	76.32	79.49	67.76	68.64	69.88	70.21
GER	89.94	88.99	81.76	81.55	84.27	86.30
JPN	84.22	83.02	83.15	83.60	84.67	76.93
SUI	92.06	88.47	87.14	85.21	79.77	80.62
UK	82.13	86.43	79.26	86.20	73.61	81.32
US	78.16	75.92	74.85	70.44	52.64	54.81

Table 4.18: Percentage of rejections of the null hypothesis of the D -test ($H_0 : \beta_1 = \beta_2 = \beta$) as indicator for the existence of the EPK paradox across the worldwide largest six stock markets. We employed two GMM estimation techniques.




Figure 4.30: Empirical pricing kernels across six worldwide largest stock markets for two scenarios, i.e., (a) $\beta_1, \beta_2 > 0$ (blue) and (c) $\beta_1 = \beta_2 = \beta$ (red). The empirical pricing kernels are fitted to average parameter estimates and the average value of the optimal reference point. We employ the iterated GMM estimation technique in (a) and the GMM estimation technique with the HJ weighting matrix in (c) with $n = 500$ (2 years).

5 Conclusions

*If all the economists in the world were laid
end to end, they would still not reach a conclusion.
George Bernard Shaw*

5.1 Modelling and Forecasting Liquidity Supply

The dynamic semiparametric factor model (DSFM) is here used for modelling and forecasting of excess supply on stock markets, i.e., to capture the dynamics of high-dimensional bid and ask curves on a limit order book market. We extend and rewrite the work by Härdle et al. [2012a] while applying the DSFM proposed by Fengler et al. [2007], Brüggemann et al. [2008], Park et al. [2009] and Cao et al. [2009]. The idea behind the DSFM, stipulated under the philosophy 'smooth in space and parametric in time', is to capture the order curve's spatial structure using a factor decomposition which is estimated nonparametrically. The order book's dynamics is in the following step modelled using a vector error correction (VEC) specification applied to the corresponding factor loadings with the bid-ask spread as the (only) cointegration relationship. Due to the modelling flexibility, one successfully reduces the high dimension of the book (25, 50, 75 or 101 price grids on each market side) as well as extracts the relevant information concerning the order book dynamics.

The model is applied to limit order book data of four stocks traded at the Australian Securities Exchange (ASX) from 8 July to 16 August 2002 (30 trading days). After removing the seasonal pattern in the data, we contest two DSFM implementation methods. The DSFM-Separated (each market side analyzed separately) approach outperforms the DSFM-Combined approach (bid side reversed), explaining up to 95% of in-sample variations. Including more knots does not lead to significant improvements of the explained variance (EV) or in the corresponding RMSE. A two-factor DSFM-Separated specification is sufficient to capture the curve dynamics and is used in the sequel of the analysis. The first factor captures the overall slope of the curve which is associated with the average trading costs and the second factor the order curve fluctuations. The shape of the second factor looks different for levels close to the best bid/ask quotes than for levels very deep in the book, i.e., the second factor is indeed responsible for order book's curvature. Modelling the full bid or ask curve (dimension $J = 101$) yields the highest explained variance.

The factor loadings vary considerably over time and are (statistically) strongly persistent. On higher frequencies than 5 minutes the factor loadings are driven toward unit root processes. No significant evidence is found for common stochastic trends (cointegration) underlying the order book factor loadings. A vector error correction (VEC)

specification of order q (maximum lag order of $q = 4$ is sufficient) with the spread as the only cointegration relationship is established. The processes follow a strong own-process dynamics and show relatively weak cross-dependencies between the endogenous variables. The market cross-dependencies are most pronounced for less frequently traded stocks (MIM and WOW). Quote changes are short-run predictable given the shape of the order book. Changes in the factor loading have a short term impact (up to 5-10 minutes) on the quote changes. High frequently traded stocks (BHP and NAB) have more pronounced impact than less liquid stocks. Asymmetric reactions of slope factor loadings on changes of the bid-ask spread are reported. Rising spreads tend to reduce (increase) the order aggressiveness on the bid (ask) side.

The DSFM approach successfully predicts the liquidity supply over various forecasting horizons during a trading day in an realistic out-of-sample forecasting exercise. It moreover outperforms a naive forecasting approach. In a trading strategy order execution costs can be reduced if orders are optimally placed according to predictions of liquidity supply. Optimal order placement in periods of high liquidity results in smaller transaction costs than in the case of a proportional splitting over time. Our flexible approach allows us to estimate and to predict future (excess) demand and supply elasticities, as well as to forecast simultaneously the shape and the location of the limit order book.

These results show that the DSFM approach is suitable for modelling and forecasting of liquidity supply. Since it is computationally tractable, it can serve as a valuable building block for automated trading models.

5.2 Localizing Multiplicative Error Models

A local adaptive multiplicative error model (MEM) has been proposed by Härdle et al. [2012b] for modelling and forecasting of high-frequency variables. By providing more implementation details we enrich their study in our work. The local MEM relies on the local parametric approach (LPA) introduced by Spokoiny [1998] which has been gradually introduced to time series literature. It addresses the tradeoff between parameter variability and the modelling bias. The length of the *interval of homogeneity* is chosen by a sequential testing procedure. By estimating the MEM over the *interval of homogeneity* one obtains the adaptive estimate used for financial data predictions.

The proposed approach is applied to the high-frequency series of one-minute cumulative trading volumes based on several NASDAQ blue chip stocks. One observes that the MEM parameters and their distribution clearly vary over time. The length of the *interval of homogeneity* varies between one and six hours. A conservative modelling approach would select on average around 3-4 hours of data, whereas a more modest (risk) approach suggest to take 2-3 hours. The local MEM provides significantly better out-of-sample forecasts than competing 'standard' approaches using a priori fixed lengths of estimation intervals. The findings are quite robust with respect to the choice of underlying tuning parameters.

Our findings can be carried over to other (persistent) financial processes that exhibit similar stochastic properties as cumulative trading volumes. For example, one can con-

sider duration data, trade counts, bid-ask spreads, transaction costs, market depth or volatilities. Adaptive techniques play an important role in high-frequency forecasting. One may gain deeper insights into the local variation of the model parameters and structural relationships and use the information to achieve economic and financial benefits.

5.3 Cross Country Evidence for the EPK Paradox

Following a state-dependent utility approach, see, e.g., Grith et al. [2011], we provide worldwide evidence for the existence of the EPK paradox on stock markets. The results are quite robust across countries and for the underlying framework specifications. Since estimated EPK parameters exhibit a time-varying pattern, one may systematically study the preference dynamics. One could relate the present study to portfolio selection and then evaluate the implied economic and financial benefits. As shown in this work, it is possible to achieve positive results while applying modern statistical and econometric methods.

Bibliography

- H. J. Ahn, K. H. Bae, and K. Chan. Limit orders, depth, and volatility: evidence from the stock exchange of Hong Kong. *Journal of Finance*, 56:767–788, 2001.
- A. Alfonsi, A. Fruth, and A. Schied. Optimal execution strategies in limit order books with general shape functions. *Quantitative Finance*, 10:143 – 157, 2010.
- R. Almgren and N. Chriss. Optimal Execution of Portfolio Transactions. *Journal of Risk*, 3:5–39, 2000.
- S. Berkowitz, D. Logue, and E. Noser. The Total Cost of Transactions on the NYSE. *Journal of Finance*, 43:97–112, 1988.
- D. Bertsimas and A. W. Lo. Optimal Control of Execution Costs. *Journal of Financial Markets*, 1:1–50, 1998.
- B. Biais, P. Hillion, and C. Spatt. An empirical analysis of the limit order book and the order flow in the Paris bourse. *Journal of Finance*, 50:1655–1689, 1995.
- R. Bloomfield, M. O’Hara, and G. Saar. The ”make or take” decision in an electronic market: evidence on the evolution of liquidity. *Journal of Financial Economics*, 75: 165–200, 2005.
- C. T. Brownlees, F. Cipollini, and G. M. Gallo. Intra-daily Volume Modeling and Prediction for Algorithmic Trading. *Journal of Financial Econometrics*, 9(3):489–518, 2011.
- R. Brüggemann, W. Härdle, J. Mungo, and C. Trenkler. VAR modelling for Dynamic Semiparametric Factors of Volatility Strings. *Journal of Financial Econometrics*, 5 (2):189–218, 2008.
- J. Cao, W. Härdle, and J. Mungo. A Joint Analysis of the KOSPI 200 Option and ODAX Option Markets Dynamics. Discussion Paper 019, Collaborative Research Center 649 ”Economic Risk”, Humboldt-Universität zu Berlin, 2009.
- G. C. Chacko, J. W. Jurek, and E. Stafford. The Price of Immediacy. *Journal of Finance*, 63:1253–1290, 2008.
- Y. Chen, W. Härdle, and U. Pigorsch. Localized Realized Volatility. *Journal of the American Statistical Association*, 105(492):1376–1393, 2010.

Bibliography

- P. Čížek, W. K. Härdle, and V. Spokoiny. Adaptive pointwise estimation in time-inhomogeneous conditional heteroscedasticity models. *Econometrics Journal*, 12:248–271, 2009.
- J. H. Cochrane. *Asset Pricing*. Princeton University Press, Princeton, New Jersey, 2001.
- H. Degryse, F. Jong, M. Ravenswaaij, and G. Wuyts. Aggressive Orders and the Resiliency of a Limit Order Market. *Review of Finance*, 9:201 – 242, 2005.
- F.X. Diebold and R. Mariano. Comparing predictive accuracy. *Journal of Business and Economic Statistics*, 13:253–265, 1995.
- R. F. Dittmar. Nonlinear Pricing Kernels, Kurtosis Preference, and Evidence from the Cross Section of Equity Returns. *Journal of Finance*, 57(1):369–403, 2002.
- R. F. Engle. Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation. *Econometrica*, 50(4):987–1007, 1982.
- R. F. Engle. New Frontiers for ARCH Models. *Journal of Applied Econometrics*, 17:425–446, 2002.
- R. F. Engle and R. Ferstenberg. Execution Risk. *Journal of Portfolio Management*, 33:34–45, 2007.
- R. F. Engle and A. J. Patton. Impacts of trades in an error-correction model of quote prices. *Journal of Financial Markets*, 7:1–25, 2004.
- R. F. Engle and J. G. Rangel. The Spline-GARCH Model for Low-Frequency Volatility and Its Global Macroeconomic Causes. *Review of Financial Studies*, 21:1187–1222, 2008.
- R. F. Engle and J. R. Russell. Autoregressive Conditional Duration: A New Model for Irregularly Spaced Transaction Data. *Econometrica*, 66(5):1127–1162, 1998.
- E. F. Fama and K. R. French. The Cross-Section of Expected Stock Returns. *Journal of Finance*, 47(2):427–465, 1992.
- E. F. Fama and K. R. French. Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics*, 33:3–56, 1993.
- M. R. Fengler, W. Härdle, and E. Mammen. A Dynamic Semiparametric Factor Model for Implied Volatility String Dynamics. *Journal of Financial Econometrics*, 5(2):189–218, 2007.
- W. E. Ferson and S. R. Foerster. Finite sample properties of the Generalized Method of Moments in tests of conditional asset pricing models. *Journal of Financial Economics*, 36:29–55, 1994.
- C. D. Fuh, H. W. Teng, and R. H. Wang. On-line VWAP Trading Strategies. *Sequential Analysis*, 29(3):292–310, 2010.

- A. R. Gallant. On the bias of flexible functional forms and an essentially unbiased form. *Journal of Econometrics*, 15:211–245, 1981.
- R. Garvey and F. Wu. Intraday time and order execution quality dimensions. *Journal of Financial Markets*, 12:203–228, 2009.
- R. Y. Goyenko, C. W. Holden, and C. A. Trzcinka. Do liquidity measures measure liquidity? *Journal of Financial Economics*, 92:153–181, 2009.
- M. D. Griffiths, B. F. Smith, D. A. S. Turnbull, and R. W. White. The costs and determinants of order aggressiveness. *Journal of Financial Economics*, 56:65–88, 2000.
- M. Grith, V. Krätschmer, and W. Härdle. A Microeconomic Explanation of the EPK Paradox. Resubmitted to *Journal of Financial Econometrics*. 20.09.2011, Collaborative Research Center 649 "Economic Risk", Humboldt-Universität zu Berlin, 2011.
- A. D. Hall and N. Hautsch. Order aggressiveness and order book dynamics. *Empirical Economics*, 30:973–1005, 2006.
- A. D. Hall and N. Hautsch. Modelling the buy and sell intensity in a limit order book market. *Journal of Financial Markets*, 10(3):249–286, 2007.
- L. P. Hansen. Large Sample Properties of Generalized Method of Moments Estimators. *Econometrica*, 50(4):1029–1054, 1982.
- L. P. Hansen and R. Jagannathan. Assessing Specification Errors in Stochastic Discount Factor Models. *Journal of Finance*, 52(2):557–590, 1997.
- L. P. Hansen and K. J. Singleton. Generalized Instrumental Variables Estimation of Nonlinear Rational Expectations Models. *Econometrica*, 50(5):1269–1286, 1982.
- W. Härdle, N. Hautsch, and A. Mihoci. Modelling and Forecasting Liquidity Supply using Semiparametric Factor Dynamics. *Journal of Empirical Finance*, 19(4):610–625, 2012a.
- W. Härdle, N. Hautsch, and A. Mihoci. Local Adaptive Multiplicative Error Models for High-Frequency Forecasts. Discussion Paper, submitted to *Journal of Applied Econometrics*. 24.04.2012, manuscript ID, MS 8151 2012-31, Collaborative Research Center 649 "Economic Risk", Humboldt-Universität zu Berlin, 2012b.
- J. Hasbrouck. Trading Costs and Returns for U.S. Equities: Estimating Effective Costs from Daily Data. *Journal of Finance*, 64:1445–1477, 2009.
- J. Hasbrouck and G. Saar. Technology and liquidity provision: The blurring of traditional definitions. *Journal of Financial Markets*, 12:143–172, 2009.
- N. Hautsch. *Econometrics of Financial High-Frequency Data*. Springer, Berlin, 2012.
- N. Hautsch and R. Huang. The Market Impact of a Limit Order. *Journal of Economic Dynamics and Control* - forthcoming, 2011.

Bibliography

- B. Hollifield, R. A. Miller, and P. Sandås. Empirical analysis of limit order markets. *Review of Economic Studies*, 71:1027–1063, 2004.
- R. Jagannathan and Z. Wang. The Conditional CAPM and the Cross-Section of Expected Returns. *Journal of Finance*, 51(1):3–53, 1996.
- N. Jegadeesh. Evidence of Predictable Behaviour of Security Returns. *Journal of Finance*, 45(3):881–898, 1990.
- S Johansen. Estimation and hypothesis testing of cointegration vectors in gaussian vector autoregressive models. *Econometrica*, 59(6):1551–1580, 1991.
- T. C. Johnson. Volume, liquidity and liquidity risk. *Journal of Financial Economics*, 87:388–417, 2008.
- D. Kwiatkowski, P. C. B. Phillips, P. Schmidt, and Y. Shin. Testing the null hypothesis of stationarity against the alternative of a unit root: How sure are we that economic time series have a unit root? *Journal of Econometrics*, 54:159–178, 1992.
- J. Large. Measuring the Resiliency of an Electronic Limit Order Book. *Journal of Financial Markets*, 10:1–25, 2007.
- W.-M. Liu. Monitoring and limit order submission risks. *Journal of Financial Markets*, 12:107–141, 2009.
- S. Manganelli. Duration, volume and volatility impact of trades. *Journal of Financial Markets*, 8:377–399, 2005.
- D. Mercurio and V. Spokoiny. Statistical inference for time-inhomogeneous volatility models. *The Annals of Statistics*, 32(2):577–602, 2004.
- J. Mincer and V. Zarnowitz. *The Evaluation of Economic Forecasts - in Economic Forecasts and Expectations*. National Bureau of Economic Research, New York, 1969.
- C. R. Nelson and A. F. Siegel. Parsimonious Modelling of Yield Curves. *Journal of Business*, 60:473–489, 1987.
- W. K. Newey and K. D. West. Hypothesis Testing with Efficient Method of Moments Estimation. *International Economic Review*, 28(3):777–787, 1987.
- A. Obizhaeva and J. Wang. Optimal Trading Strategy and Supply/Demand Dynamics. Working Paper 11444, NBER, 2005.
- B. Park, E. Mammen, W. Härdle, and S. Borak. Time Series Modelling With Semiparametric Factor Dynamics. *Journal of the American Statistical Association*, 104(485): 284–298, 2009.
- A. Ranaldo. Order aggressiveness in limit order book markets. *Journal of Financial Markets*, 7:53–74, 2004.

- P. Schmidt and P. C. B. Phillips. LM tests for a unit root in the presence of deterministic trends. *Oxford Bulletin of Economics and Statistics*, 54:257–287, 1992.
- U. Schweri. *Is the Pricing Kernel U-Shaped*. PhD thesis, University of Zurich, Swiss Finance Institute, 2010.
- V. Spokoiny. Estimation of a function with discontinuities via local polynomial fit with an adaptive window choice. *The Annals of Statistics*, 26(4):1356–1378, 1998.
- V. Spokoiny. Multiscale local change point detection with applications to Value-at-Risk. *The Annals of Statistics*, 37(3):1405–1436, 2009.
- H. Tong. *Non-linear Time Series: A Dynamical System Approach*. Oxford University Press, Oxford, 1990.
- M. Y. Zhang, J. R. Russell, and R. S. Tsay. A nonlinear autoregressive conditional duration model with applications to financial transaction data. *Journal of Econometrics*, 104:179–207, 2001.

Bibliography

Selbständigkeitserklärung

Ich bezeuge durch meine Unterschrift, dass meine Angaben über die bei der Abfassung benutzten Hilfsmittel, über die mir zuteil gewordene Hilfe sowie über frühere Begutachtungen meiner Dissertation in jeder Hinsicht der Wahrheit entsprechen.

Berlin, 10. August 2012

Andrija Mihoci