

Härdle, Wolfgang Karl; Hautsch, Nikolaus; Mihoci, Andrija

Working Paper

Modelling and forecasting liquidity supply using semiparametric factor dynamics

CFS Working Paper, No. 2009/18

Provided in Cooperation with:

Center for Financial Studies (CFS), Goethe University Frankfurt

Suggested Citation: Härdle, Wolfgang Karl; Hautsch, Nikolaus; Mihoci, Andrija (2009) : Modelling and forecasting liquidity supply using semiparametric factor dynamics, CFS Working Paper, No. 2009/18, Goethe University Frankfurt, Center for Financial Studies (CFS), Frankfurt a. M., <https://nbn-resolving.de/urn:nbn:de:hebis:30-70742>

This Version is available at:

<https://hdl.handle.net/10419/43198>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

No. 2009/18

**Modelling and Forecasting Liquidity Supply
Using Semiparametric Factor Dynamics**

Wolfgang Karl Härdle, Nikolaus Hautsch,
and Andrija Mihoci





Center for Financial Studies

The *Center for Financial Studies* is a nonprofit research organization, supported by an association of more than 120 banks, insurance companies, industrial corporations and public institutions. Established in 1968 and closely affiliated with the University of Frankfurt, it provides a strong link between the financial community and academia.

The CFS Working Paper Series presents the result of scientific research on selected topics in the field of money, banking and finance. The authors were either participants in the Center's Research Fellow Program or members of one of the Center's Research Projects.

If you would like to know more about the *Center for Financial Studies*, please let us know of your interest.

A handwritten signature in black ink, appearing to read "Jan Krahnen".

Prof. Dr. Jan Pieter Krahnen



Modelling and Forecasting Liquidity Supply Using Semiparametric Factor Dynamics*

Wolfgang Karl Härdle¹, Nikolaus Hautsch²,
and Andrija Mihoci³

September 15, 2009

Abstract:

We model the dynamics of ask and bid curves in a limit order book market using a dynamic semiparametric factor model. The shape of the curves is captured by a factor structure which is estimated nonparametrically. Corresponding factor loadings are assumed to follow multivariate dynamics and are modelled using a vector autoregressive model. Applying the framework to four stocks traded at the Australian Stock Exchange (ASX) in 2002, we show that the suggested model captures the spatial and temporal dependencies of the limit order book. Relating the shape of the curves to variables reflecting the current state of the market, we show that the recent liquidity demand has the strongest impact. In an extensive forecasting analysis we show that the model is successful in forecasting the liquidity supply over various time horizons during a trading day. Moreover, it is shown that the model's forecasting power can be used to improve optimal order execution strategies.

JEL-Classifications: C14, C32, C53, G11

Keywords: Limit Order Book, Liquidity Risk, Semiparametric Model, Factor Structure, Prediction

* We are very grateful to Anthony Hall for providing us the data. For helpful comments and discussions we thank Joachim Grammig, Jeffrey Russell and the participants of the 2009 Humboldt-Copenhagen Conference on Financial Econometrics in Berlin, the 2009 annual conference of the Society for Financial Econometrics (SoFie) in Geneva, the 2009 European Meeting of the Econometric Society in Barcelona, the International Conference on Price, Liquidity and Credit Risk in Konstanz, 2008, as well as the 4th World Congress of the International Association for Statistical Computing in Yokohama, 2008. Furthermore, we are grateful to Szymon Borak for helping us with the implementation of the Dynamic Semiparametric Factor Model in MATLAB. Moreover, financial support from the Deutsche Forschungsgemeinschaft via SFB 649 "Ökonomisches Risiko", Humboldt-Universität zu Berlin, is gratefully acknowledged.

1 Humboldt-Universität zu Berlin, C.A.S.E. - Center for Applied Statistics and Economics, Spandauer Str. 1, 10178 Berlin, Germany, and National Central University, Department of Finance, No. 300, Jhongda Rd., Jhongli City, Taoyuan County 32001, Taiwan (R.O.C.)

2 School of Business and Economics as well as C.A.S.E. - Center for Applied Statistics and Economics, Humboldt-Universität zu Berlin, Quantitative Products Laboratory (QPL), Berlin, and Center for Financial Studies (CFS), Frankfurt. Address: Spandauer Str. 1, 10178 Berlin, Germany, tel: +49 (0)30 2093 5711, email: nikolaus.hautsch@wiwi.hu-berlin.de

3 Humboldt-Universität zu Berlin, C.A.S.E. - Center for Applied Statistics and Economics, Spandauer Str. 1, 10178 Berlin, Germany, and Faculty of Economics & Business, Zagreb, Croatia

1 Introduction

Due to technological progress in the organization of trading systems and exchanges, electronic limit order book trading has become the dominant trading form for equities. Open limit order books provide important information on the current liquidity supply in the market as reflected by the offered price-quantity relationships on both sides of the market. Whereas in traditional market maker markets or over-the-counter markets only the corresponding best ask and bid prices are observable, the limit order book displays also limit prices and volumes behind the market. Demand and supply schedules provide valuable information on traders' price expectations in the spirit of the seminal paper by Glosten (1994), reflect the current implied costs of trading as well as demand and supply elasticities. Such information is useful for investors to optimally design order execution strategies.

In this paper, we propose a dynamic semiparametric factor approach to model and to forecast the dynamics of liquidity supply in a limit order book. The liquidity supply is captured by the demand schedule on both sides of the market as represented by the posted order volumes on a price grid around the prevailing best ask and bid quotes. The paper's major idea is to capture the shape of the high-dimensional ask and bid curves by a lower-dimensional factor structure which is estimated non-parametrically. The curves' dynamic behavior is driven by time-varying factor loadings which are modelled parametrically employing a vector autoregressive (VAR) approach.

Due to the availability of data on limit order book markets, the empirical analysis thereof has become an important field in empirical finance and high-frequency econometrics. A dominant part of empirical research on limit order book markets is devoted to the analysis of traders' order submission strategies and implications thereof for liquidity and volatility dynamics, such as, e.g., Biais et al. (1995), Griffiths et al. (2000), Ahn et al. (2001), Rinaldo (2004), Hollifield et al. (2004), Bloomfield et al. (2005), Hall and Hautsch (2006, 2007) or Hasbrouck and Saar (2009). A further major issue is the analysis of market transparency and its impact on liquidity, see, e.g., Baruch (2005), Boehmer et al. (2005), Comerton-Forde and Tang (2009) and Eom et al. (2007). A central aspect in this literature is to analyze the question of how to optimally balance risks and gains of a trader's decision whether to post a market order or a limit order. As recently illustrated by Chacko et al. (2008), a limit order can be ultimately seen as an American option and transaction costs are rents that a monopolistic market maker




Figure 1: Limit order book for selected stocks traded at the ASX on July 8, 2002 at 10:15. Red: bid curve, blue: ask curve.

extracts from impatient investors who trade via aggressive limit orders or market orders. Consequently, the analysis of liquidity risks (see, e.g., Johnson (2008), Liu (2009), Garvey and Wu (2009) or Goyenko et al. (2009)) and transaction costs (see, e.g. Chacko et al. (2008) and Hasbrouck (2009)) are in the central focus of recent literature.

In financial practice, the question of how to reduce the costs of trading by optimally splitting a large order over time (e.g., the course of a trading day) is of high relevance, see, e.g. Engle and Ferstenberg (2007). Optimal splitting naturally requires to predict future liquidity demand and supply. While there is some (though little) empirical evidence on the predictability of intraday volume (see, e.g., Hautsch (2008) and Brownlees et al. (2009)), the dynamic behavior of liquidity supply is widely unknown.

This paper aims to address this question and to gain deeper insights into the question to which extent the offered ask and bid volume can be modelled and predicted. The ultimate objective of interest is the number of shares pending on a relative price grid around the best ask and bid curves. Given the objective to capture not only the volume around the best quotes but also pending quantities 'behind' the market, the underlying problem becomes inherently high-dimensional. A typical graphical snapshot of ask and bid curves for four stocks traded at the Australian Securities Exchange (ASX) in 2002, is given by Figure 1. The curse of dimensionality applies immediately as soon as time variations of the order curve shapes have to be taken into account. As shown by Figure 1 and as illustrated in more detail in the sequel of the paper, order volume is not necessarily only concentrated around the best quotes but can be substantially dispersed over a wider range of price levels. This is a typical scenario for moderately liquid markets as that of the ASX. In such a context, the dynamic modelling of all volume levels individually becomes complicate and intractable.

Consequently, we suggest reducing the high dimensionality of the order book by means of a factor decomposition using 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). Accordingly, we model the shape of the book in terms of underlying

1 Introduction

latent factors which are defined on the price grid space around the best ask or bid quote and can depend on additional explanatory variables capturing, e.g., the state of the market. In order to avoid specific functional forms for the shape of the curves, the factors as well as the corresponding loadings are estimated nonparametrically using B-splines. Then, in a second step, we model the multivariate dynamics of the factor loadings using a VAR model.

Using this framework we aim answering the following research questions: (i) How many factors are required to model order book curves reasonably well? (ii) How does the shape of the factors look like? (iii) How do the dynamics of the estimated factor loadings look like? (iv) Does the shape of the order book curves depend on past price movements, past trading volume as well past volatility? (v) How successful is the model in predicting future liquidity supply and can it be used to improve order execution strategies?

Using limit order book data from four stocks traded at the ASX covering two months in 2002, we show that approximately 95% of the order book variations observed on 5-min intervals can be explained by two underlying time-varying factors. While the first factor captures the overall slope of the curves, the second one is associated with its curvature. It turns out that recent liquidity demand represented by the cumulative buy/sell trading observed over the past 5 minutes has an effect of the shape of the curve but does not induce a higher explanatory power. Similar evidence is shown for the impact of past returns and corresponding (realized) volatility. Furthermore, it is shown that the factor loadings follow highly persistent though stationary dynamics suggesting an underlying vector autoregressive structure.

To evaluate the model's forecasting power, we perform an extensive out-of-sample forecasting analysis which is in line with a typical scenario in financial practice. In particular, at every 5-min interval during a trading day, the model is re-estimated and used to produce forecasts for the pending volume on each price level for all future 5-min intervals during the remainder of the trading day. We show that our approach is able to outperform a naive prediction, where the current order book is used as a predictor for the remaining day. Moreover, it is illustrated that these results can be used to improve order execution strategies by reducing implied transaction costs.

The remainder of the paper is structured as follows: After the data description in Section 2, the Dynamic Semiparametric Factor Model (DSFM) is introduced in Section 3. Empirical results regarding the modelling and forecasting of liquidity supply are provided in Sections 4 and 5, respectively. Section 6 concludes.

2 Data

2.1 Trading at the ASX and Descriptive Statistics

The Australian Stock Exchange (ASX) is a continuous double auction electronic market, where 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. During continuous trading, 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. Note that there might be orders which are entered with an undisclosed or hidden volume if the total value of the order exceeds AUD 200,000. Since this applies only to a small fraction of the posted volumes, we can 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.

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 1: Number of market and limit orders for selected stocks at the ASX from July 8 to August 16, 2002

We select four companies traded at the ASX covering the period from July 8 to August 16, 2002 (30 trading days), namely Broken Hill Proprietary Limited (BHP), National Australia Bank Limited (NAB), MIM and Woolworths (WOW). The number of market

2 Data

and limit orders for the selected stocks is given in Table 1.

We observe 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. Furthermore, according to Table 2, for both traded as well as posted quantities it is shown that on average sell volumes are higher than buy volumes. Hence, confirming the result above, liquidity variations on the bid side are 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.

Quantity	BHP	NAB	MIM	WOW
Quoted quantity				
(i) buy	7,359	4,292	40,559	4,686
(ii) sell	12,247	4,593	54,162	5,579
Traded quantity				
(i) buy	5,456	2,761	36,033	3,232
(ii) sell	8,340	2,918	44,872	3,709

Table 2: Average quoted and traded quantities for selected stocks at the ASX from July 8 to August 16, 2002

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.

2.2 Notation and Data Preprocessing

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 $t = 1, \dots, T$. Particularly, 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. Hence, 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 $\tilde{Y}_{t,j}^b$ and $\tilde{Y}_{t,j}^a$ as the pending bid and ask volumes at bid and ask

2 Data

limit prices $\tilde{S}_{t,j}^b$ and $\tilde{S}_{t,j}^a$, respectively at time point t .

We define the best bid price at time t 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 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 $S_{t,j}^b = \check{S}_{t,j}^b/\tilde{S}_{t,101}^b$ and $S_{t,j}^a = \check{S}_{t,j}^a/\tilde{S}_{t,1}^a$, respectively.

Note that modelling order book curves in terms of a grid of (relative or absolute) price *deviations* from the prevailing best ask and bid quotes implies that we do not model price *levels* and thus, also take out the bid-ask spread. Capturing not only dynamics of order book curves but also of corresponding price *levels* would make the analysis significantly more complicate since (common) stochastic trends in price levels would have to be taken into account. See Hautsch and Huang (2009) for a corresponding model for quotes and depth. However, since the focus of the present paper is to capture dynamics in order book curves and to produce forecasts for pending volumes rather than predicting quotes themselves, it is sensible to refrain from price dynamics. Nevertheless, in situations where spread forecasts are required, the proposed model might be easily augmented by a corresponding (time series) model for spreads.

In order to account for intraday 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

$$Y_{t,j}^b = \frac{\tilde{Y}_{t,j}^b}{s_t^b} \quad (1)$$

$$Y_{t,j}^a = \frac{\tilde{Y}_{t,j}^a}{s_t^a}, \quad (2)$$

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

2 Data

flexible Fourier series approximation as proposed by Gallant (1981) and are given by

$$s_t^b = \delta^b \cdot \bar{t} + \sum_{m=1}^{M^b} \{ \delta_{c,m}^b \cos(\bar{t} \cdot 2\pi m) + \delta_{s,m}^b \sin(\bar{t} \cdot 2\pi m) \} \quad (3)$$

$$s_t^a = \delta^b \cdot \bar{t} + \sum_{m=1}^{M^a} \{ \delta_{c,m}^a \cos(\bar{t} \cdot 2\pi m) + \delta_{s,m}^a \sin(\bar{t} \cdot 2\pi m) \}. \quad (4)$$

Here δ^b , δ^a , $\delta_{c,m}^b$, $\delta_{c,m}^a$ and $\delta_{s,m}^b$ 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 intraday seasonal patterns for both sides of all limit order book markets are plotted in Figure 2.




Figure 2: Intraday seasonality factors for quantities offered at best bid prices (red) and for quantities supplied at best ask prices (blue).

For all stocks, we observe that the liquidity supply is increasing before 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). In contrast, for the less liquid stocks, the amount of posted volume is nearly monotonically increasing over the day.

3 The Dynamic Semiparametric Factor Model

Recall that the object of interest is the high-dimensional object of seasonally adjusted level-dependent order volume inventories $(Y_{t,j}^b, Y_{t,j}^a) \in \mathbb{R}^{202}$, observed on a 5-min frequency. Proposing a suitable statistical model results in the problem of finding an appropriate way of reducing the high dimension without losing too much information on the spatial and dynamic structure of the process. Moreover, applicability of the model requires computational tractability as well as numerical stability.

A common way to reduce the dimensionality of multivariate processes is to apply a factor decomposition. The underlying idea is that the high-dimensional process is ideally driven by only a few common factors which contain most underlying information. Factor models are often applied in the asset pricing literature to extract underlying common risk factors. In this spirit, a successful parametric factor model has been proposed, for instance, by Nelson and Siegel (1987) to model yield curves. In this framework, the shape of the curve is parametrically captured by Laguerre polynomials.

Since limit order book curves inherently reflect traders' price expectations and the supply and demand in the market (see, e.g. Glosten (1994) for a theoretical framework), the assumption of an underlying factor structure is economically very reasonable. However, as there is no obvious parametric form for ask and bid curves and we want to avoid imposing assumptions on functional form, we prefer to capture the curve's spatial structure in a nonparametric way. A natural and powerful class of models for these kind of problems is the class of Dynamic Semiparametric Factor Models (DSFMs) proposed by Fengler et al. (2007), Brüggemann et al. (2008), Park et al. (2009) and Cao et al. (2009). The DSFM model successfully combines the advantages of a nonparametric approach for cross-sectionally (spatially) fitting a curve and that of a parametric time series model for modelling persistent multivariate dynamics. In the following section we will discuss the DSFM model and its implementation for limit order book dynamics.

Assume that the observable J -dimensional random vector, $Y_{t,j}$, can be modelled based on the following orthogonal L -factor model,

$$Y_{t,j} = m_{0,j} + Z_{t,1}m_{1,j} + \cdots + Z_{t,L}m_{L,j} + \varepsilon_{t,j}, \quad (5)$$

where $m(\cdot) = (m_0, m_1, \dots, m_L)^\top$ denotes the time-invariant factors, a tuple of functions with the property $m_l : \mathbb{R}^d \rightarrow \mathbb{R}$, $l = 0, \dots, L$, $Z_t = (1_T, Z_{t,1}, \dots, Z_{t,L})^\top$ denotes the time series of factor loadings, and $\varepsilon_{t,j}$ represents a white noise error term. The time index is denoted by $t = 1, \dots, T$, whereas the cross-sectional index is $j = 1, \dots, J$. Note that this factor model is rather restrictive, because it does not take explanatory variable into account.

The DSFM is a generalization of the factor model given in (5) and allows the factors m_l

3 The Dynamic Semiparametric Factor Model

to depend upon explanatory variables, $X_{t,j}$. Its analytical form is given by

$$Y_{t,j} = \sum_{l=0}^L Z_{t,l} m_l(X_{t,j}) + \varepsilon_{t,j} = Z_t^\top m(X_{t,j}) + \varepsilon_{t,j}, \quad (6)$$

where the processes $X_{t,j}$, $\varepsilon_{t,j}$ and Z_t are assumed to be independent. Moreover, the number of underlying factors L should not exceed the dimension of the object, J . The main idea of the DSFM is that L is significantly smaller than J resulting in a severe dimension reduction of the process.

As suggested by Park et al. (2009), the estimation of the factors m_l is performed using a series estimator. For $K \geq 1$, appropriate functions $\psi_k : [0, 1]^d \rightarrow \mathbb{R}$, $k = 1, \dots, K$, which are normalized such that $\int \psi_k^2(x) dx = 1$ holds, are selected. Park et al. (2009) select tensor B-spline basis functions for ψ_k , whereas Fengler et al. (2007) use a kernel smoothing approach. In the present study, we follow the former strategy and employ tensor B-spline basis functions.

After selecting the functions ψ_k , the factors $m(\cdot) = (m_0, m_1, \dots, m_L)^\top$ are approximated by $A\psi$, where $A = (a_{l,k}) \in \mathbb{R}^{(L+1)K}$ is a coefficient matrix, and $\psi(\cdot) = (\psi_1, \dots, \psi_K)^\top$ denotes a vector of selected functions. Here, K denotes the number of knots used for the tensor B-spline functions and is interpretable as a bandwidth parameter. Thus, the first part in the right-hand side of (6), which incorporates all factors and factor loadings, can be rewritten as

$$Z_t^\top m(X_{t,j}) = \sum_{l=0}^L Z_{t,l} m_l(X_{t,j}) = \sum_{l=0}^L Z_{t,l} \sum_{k=1}^K a_{l,k} \psi_k(X_{t,j}) = Z_t^\top A\psi(X_{t,j}). \quad (7)$$

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

$$(\hat{Z}_t, \hat{A}) = \arg \min_{Z_t, A} S(A, Z_t) \quad (8)$$

$$= \arg \min_{Z_t, A} \sum_{t=1}^T \sum_{j=1}^J \left\{ Y_{t,j} - Z_t^\top A\psi(X_{t,j}) \right\}^2. \quad (9)$$

To find a solution of the minimization problem stated in (9), a Newton-Raphson algorithm is used. As shown by Park et al. (2009) this algorithm is shown to converge to a solution at a geometric rate under some weak conditions on the initial choice $\{vec(A)^{(0)}, Z_t^{(0)}\}$. Moreover, Park et al. (2009) prove that the difference between the estimated loadings \hat{Z}_t and the true loadings Z_t are asymptotically negligible. Consequently, it is justified to use in a second step multivariate time series specifications in order to model the dynamics of the factor loadings and thus that of the analyzed high-dimensional object.

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

$$EV(L) = 1 - RV(L) = 1 - \frac{\sum_{t=1}^T \sum_{j=1}^J \{Y_{t,j} - \sum_{l=0}^L \hat{Z}_{t,l} \hat{m}_l(X_{t,j})\}^2}{\sum_{t=1}^T \sum_{j=1}^J \{Y_{t,j} - \bar{Y}\}^2}. \quad (10)$$

Moreover, the knots used in the tensor B-spline functions should be specified in advance. We choose linearly spaced knots, with a starting point determined by the minimal value of the explanatory variable (corrected by -5%), and the end point corresponding to the maximal value (corrected by 5%). Sensitivity analysis shows that the results are quite stable regarding the choice of grid points.

Because of the use of tensor B-spline functions for the demand and supply curves, which are monotonous in the price levels, our estimated first factor \hat{m}_1 and the estimated quantities $\hat{Y}_{t,j}$ are adjusted for extreme price levels. Correspondingly, for the bid side we keep constant the first (lowest) ten level values, and analogously, for the ask side we fix the last (highest) ten level values.

The model's goodness-of-fit is evaluated using the root mean squared error (RMSE) criterion,

$$RMSE = \sqrt{\frac{1}{TJ} \sum_{t=1}^T \sum_{j=1}^J \{Y_{t,j} - \sum_{l=0}^L \hat{Z}_{t,l} \hat{m}_l(X_{t,j})\}^2}. \quad (11)$$

4 Modelling Limit Order Book Dynamics

We consider two possibilities to implement the DSFM approach:

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

First, we model the order book curves in dependence of the relative price levels solely. In the following step, we include additional explanatory variables, particularly, the past trading volume, past (realized) volatility as well as past log returns.

4.1 Modelling Order Book Curves in Dependence of Relative Price Levels

To model the curve dynamics in dependence of the relative price levels, i.e., the relative price deviations from the best bid price and best ask price, $S_{t,j}^b$ and $S_{t,j}^a$, respectively, we impose $K = 20$ knots for the B-spline functions in case of the separated approach and $K = 40$ knots in case of the combined approach. Using more knots does not result in significant improvements of the explained variance or in the corresponding RMSE, as defined in 10 and 11.

As shown in Tables 3 and 4, up to approximately 95% of the explained variation in order curves can be explained using $L = 2$ factors, whereas the marginal contribution of a potentially third factor is only very small. Consequently, a two-factor DSFM specification is 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 3: Explained variance (EV) of estimated order book variations depending on relative prices based on different number of factors L using both DSFM approaches.

Comparing the performance of the two alternative DSFM specifications, it turns out that 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. Figure 3 compares root mean squared errors for different absolute price levels j , $\check{S}_{t,j}^b$ and $\check{S}_{t,j}^a$, respectively. We observe that at almost every price level the DSFM-Separated approach outperforms the DSFM-Combined approach. Therefore, the remainder of the analysis will rely on the DSFM-Separated approach with two factors.

Figure 4 depicts the nonparametrically estimates of the first and second factor \hat{m}_1 and \hat{m}_2 in dependence of the relative price grids. The first factor obviously 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. In contrast, the second factor seems to capture order curve fluctuations around the overall slope and thus can be interpreted as a 'curvature' factor in the spirit of Nelson and Siegel (1987). The shape of this factor

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: 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 3: 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).

reveals that the curve's curvature is particularly distinct for levels close to the best quotes and for levels very deep in the book where the curve seems to spread out. The shapes of the estimated factors are remarkably similar for all stocks except for MIM. For the latter stock, the shapes of both factors are quite similar and significantly deviate from those reported for the other stocks. This finding is explained by the peculiarities of MIM for which the relative tick size is larger than for the other stocks. This implies that liquidity is concentrated on relatively few price levels around the best ask and bid quotes whereas the book flattens out for higher levels. This pattern is clearly revealed by the corresponding factors shown in Figure 4.




Figure 4: Estimated first and second factor of the limit order book depending on relative price levels using the DSFM-Separated approach with two factors.

Time series plots of the corresponding factor loadings \hat{Z}_t^b and \hat{Z}_t^a are shown in Figure 5. We observe that the loadings strongly vary over time reflecting time variations in the shape of the book. The series reveal clustering structures indicating a relatively high persistence in the processes. This result is not very surprising given the fact that order book inventories do not change too severely during short time horizons. Observing order book volumes on even higher frequencies than 5 minutes further increases this persistence, ultimately driving the processes toward unit root processes. Naturally, this behavior is particularly distinct for less frequently traded stocks and less severe for highly active stocks (cf. Hautsch and Huang (2009) for corresponding results for more liquid assets).

The high persistence is confirmed by autocorrelation functions of \hat{Z}_t^b and \hat{Z}_t^a (not shown in the paper) and corresponding unit root and stationarity tests. According to the

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 5: 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 6: 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.)

Schmidt-Phillips test (see Schmidt and Phillips (1992)) shown in Table 5, for all processes the null hypothesis of a unit root can be rejected at the 5% significance level. Conversely, testing the null hypothesis of stationarity using the KPSS test (see Kwiatkowski et al. (1992)) implies no rejections for the majority of the processes. Nevertheless, as shown in Table 6, in five cases we have to reject stationarity. Finally, to test for possible cointegration between the factor loadings, we perform Johansen's (1991) trace test (not shown in the paper) but do not find significant evidence for common stochastic trends underlying the order book.

As a graphical illustration for the goodness-of-fit of the model, Figure 6 depicts the estimated vs. the actually observed limit order book curves for all stocks on an arbitrarily selected day, namely July 8, 2002, at 11:00 and 13:00. The figure is quite representative for the fit over the whole sample and shows that the model fits the observed curves very well. This is particularly true for price levels close to the best ask and bid quotes. Slight deviations are observed for price levels deeply in the book. However, the latter case is less relevant for most applications in practice.

Given the diagnostics above, we conclude that it is sensible to treat the factor loadings as non-cointegrated, stationary processes suggesting a VAR specification as a natural




Figure 5: Estimated first and second factor loadings of the limit order book depending on relative price levels using the DSFM-Separated approach with two factors.




Figure 6: Estimated (dashed) and observed (solid) limit order book curves on July 8, 2002, 11:00 (upper panel) and 13:00 (lower panel).

choice of model to capture order book dynamics. Hence, we specify a VAR(p) model as

$$\begin{bmatrix} \hat{Z}_{1,t}^b \\ \hat{Z}_{2,t}^b \\ \hat{Z}_{1,t}^a \\ \hat{Z}_{2,t}^a \end{bmatrix} = c + B_1 \begin{bmatrix} \hat{Z}_{1,t-1}^b \\ \hat{Z}_{2,t-1}^b \\ \hat{Z}_{1,t-1}^a \\ \hat{Z}_{2,t-1}^a \end{bmatrix} + B_2 \begin{bmatrix} \hat{Z}_{1,t-2}^b \\ \hat{Z}_{2,t-2}^b \\ \hat{Z}_{1,t-2}^a \\ \hat{Z}_{2,t-2}^a \end{bmatrix} + \dots + B_p \begin{bmatrix} \hat{Z}_{1,t-p}^b \\ \hat{Z}_{2,t-p}^b \\ \hat{Z}_{1,t-p}^a \\ \hat{Z}_{2,t-p}^a \end{bmatrix} + \begin{bmatrix} \varepsilon_{1,t}^b \\ \varepsilon_{2,t}^b \\ \varepsilon_{1,t}^a \\ \varepsilon_{2,t}^a \end{bmatrix}. \quad (12)$$

Here b and a denote the bid and ask side, respectively, c denotes a vector with constants, and ε_t represents white noise error terms. The matrices B_1, B_2, \dots, B_p denote the corresponding parameter matrices. We determine the order p according to the BIC. In all cases, a maximum lag order of $p = 4$ is sufficient. In particular, the following VAR(p) models are selected: BHP and MIM - VAR(4), NAB - VAR(2), WOW - VAR(3).

For sake of brevity we refrain from showing all parameter estimates here, but just report the estimates of matrix B_1 for BHP, NAB, MIM and WOW which contains most relevant information (5% significance is denoted by an asterix (*)):

$$\begin{bmatrix} 0.91^* & 0.37^* & -0.04 & -0.27^* \\ 0.01 & 0.72^* & 0.00 & 0.02 \\ 0.06^* & 0.11 & 0.75^* & 0.02 \\ 0.00 & 0.03 & 0.02^* & 0.77^* \end{bmatrix}, \begin{bmatrix} 0.71^* & 0.17 & -0.03 & -0.19 \\ 0.04^* & 0.77^* & 0.00 & 0.07^* \\ 0.03 & 0.10 & 0.73^* & 0.18 \\ -0.02^* & -0.02 & 0.03^* & 0.71^* \end{bmatrix},$$

$$\begin{bmatrix} 0.88^* & 0.76^* & 0.01 & 0.26^* \\ 0.00 & 0.87^* & -0.01 & 0.01 \\ -0.04 & 0.89^* & -0.98^* & 0.16^* \\ 0.00 & -0.01 & 0.03^* & 0.83^* \end{bmatrix} \text{ and } \begin{bmatrix} 0.74^* & 0.03 & 0.08^* & 0.32^* \\ 0.04^* & 0.81^* & -0.03^* & -0.04 \\ 0.06 & 0.07 & 0.88^* & 0.16 \\ -0.03^* & 0.01 & 0.02^* & 0.82^* \end{bmatrix}.$$

As indicated by the diagonal elements, all processes reveal relatively strong own-process dynamics. Interestingly, most off-diagonal elements are comparably close to zero. This is particularly true for the more liquid assets (NAB and BHP), where spill-over effects are virtually zero. Hence, for these stocks, the dynamics associated with the slope and the curvature is widely unrelated. We also observe that there are no inter-dependencies between processes on the ask and bid side of the market indicating that time variations in the liquidity schedule on the buy side is not affected by those on the sell side and vice versa. However, in case of the less frequently traded stocks (MIM and WOW), these results are less distinct and we find stronger evidence for cross-dependencies between the factors and both sides of the market. These results provide some hints for the fact that for less liquid stocks interactions between the buy and sell side seem to be more pronounced. This might be due to the overall lower level of liquidity supply making strategic behavior regarding market imbalances more effective than in the case of comparably deep books of more liquid equities.

Having selected the VAR models using information criteria we will keep these specifications for the remainder of the analysis.

4.2 Including Further Explanatory Variables

In the previous analysis, order book curves depending on relative price levels have been modelled solely based on their own process history. In this section, we aim analyzing whether the model's fit can be further improved by the inclusion of additional regressors. As explanatory variables we select three variables for which we expect to observe the strongest impact on liquidity supply, namely the 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.

The buy and sell trading volumes at time t are given 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 $\tilde{Q}_t^s = \sum_{r=1}^{R_t^s} \tilde{Q}_r^s$, where R_t^b and R_t^s denote the number of buy and sell orders over the interval $(t-1, t]$, respectively. Correspondingly, log returns r_t and volatility V_t are computed as

$$r_t = \log \frac{\tilde{S}_t^*}{\tilde{S}_{t-1}^*} \quad (13)$$

$$V_t = r_t^2, \quad (14)$$

where \tilde{S}_t^* and \tilde{S}_{t-1}^* denote the mid-quotes observed at t and $t-1$, respectively. Note that the trading volumes as well as the volatility are seasonally adjusted following the procedure explained above. Moreover, the used nonparametric procedure requires the variables to be standardized between -1 and 1 . This standardization is performed based on the minimum and maximum observations of the corresponding variables. Finally, as commonly known, nonparametric regression becomes computationally cumbersome for a high number of regressors. To keep our approach computationally tractable and to avoid problems due to the curse of dimensionality, we include the regressors only individually (together with the relative price distances). This ultimately yields a three-dimensional problem.

Figures 7 and 8 show the estimated first factors for the bid and the ask side in dependence of the past 5-min sell and buy trading volumes, respectively. As expected, we observe that the past liquidity demand influences the order book curve. A high trading volume implies that a non-trivial part of the pending volume in the book is removed. In this context, recall that we do not model the spread but capture the curves on a grid of relative prices in relation to the current best quote. These relative price distances themselves do not strongly vary through 5-min intervals. Thus, most of the observed variation of the factor's shape is induced by the fact that either quoted price levels close to the best quotes have been completely absorbed and the remaining volume is correspondingly 'shifted down' in relation to the new best quote or, alternatively, only a part of the pending volume on the best quotes is removed changing the distribution of the pending volumes across the (relative) price levels.

As expected, the curve flattens in the area of high volumes. Strikingly, we also observe a decaying pattern if the volume sizes decline. Actually, in all pictures, 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 (though not zero) trading volumes. This pattern might be technically explained by the standardization procedure based on extreme values or by the usual boundary problems of non-parametric regression. On the other hand, note that due the curse-of-dimensionality problem we cannot simultaneously control for other variables. For instance, very small market-side-specific trading volumes can indicate the occurrence of market imbalances or, alternatively, might be associated with wide spreads. Both scenarios could force investors to post rather limit orders than market orders which might explain the decaying shape of the figures after having observed small trading volumes.




Figure 7: 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.

To evaluate whether the inclusion of past trading volume further increases the model's goodness-of-fit, Table 7 displays the corresponding RMSEs. Comparing these results with that reported above for the basis model shows that the included regressors yield higher estimation errors. Hence, obviously the inclusion of additional regressors ultimately generates more noise overcompensating a possibly higher explanatory power. Similar results are also found for the past log returns and past volatility serving as regressors. As shown by Tables 8, the inclusion of log returns yields smaller estimation errors than the inclusion of volatility. However, the overall performance is lower than in the cases above. Because of this reason, we refrain from showing corresponding graphs of the estimated factors.

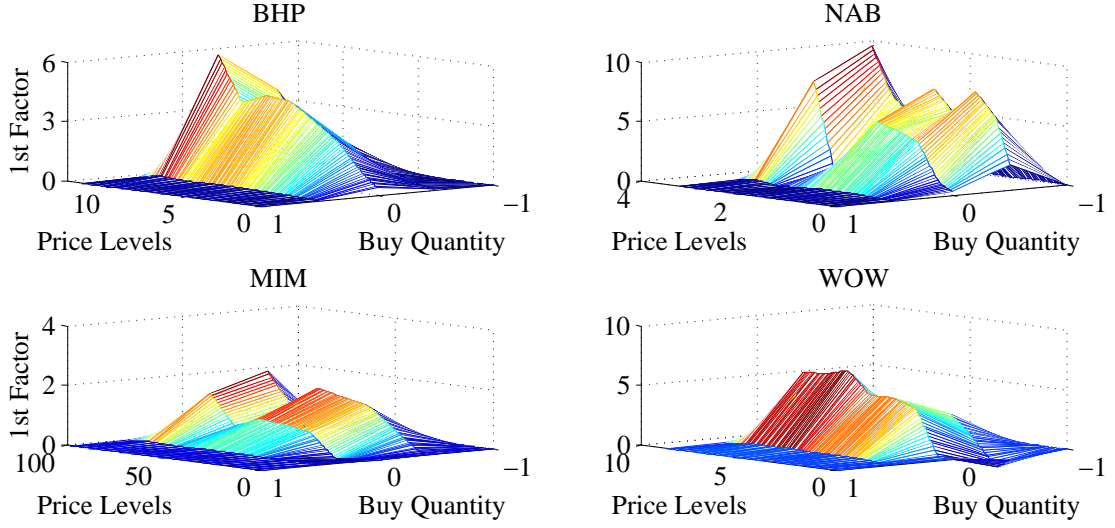


Figure 8: Estimated first factors of the ask side with respect to relative price levels and the past log traded buy volume using the DSFM-Separated approach with two factors.

A possible reason for the declining model performance in case of included regressors might be the lower dimensionality of the regressors in comparison with that of the limit order book. Note that the included regressors do not reveal any variation across the levels of the book. Consequently, the explanatory variables cannot improve the model's spatial fit but just its dynamic fit. Obviously, the latter is not sufficient to obtain an overall reduction of estimation errors.

Variable	BID				ASK			
	BHP	NAB	MIM	WOW	BHP	NAB	MIM	WOW
Q_t^s	10.37	8.17	5.41	6.31	7.38	8.30	5.72	9.18
Q_t^b	10.42	8.41	4.37	6.29	7.30	8.42	7.22	8.88

Table 7: Root mean squared errors (RMSEs) implied by estimated order book variations depending on relative prices and log de-seasonalized buy and sell trading quantities, Q_t^s and Q_t^b .

Variable	BID				ASK			
	BHP	NAB	MIM	WOW	BHP	NAB	MIM	WOW
r_t	21.93	23.09	39.47	175.40	18.00	22.13	45.54	236.08
RV_t	95.74	87.12	258.37	-	78.62	63.63	192.87	-

Table 8: Root mean squared errors (RMSEs) implied by estimated order book variations depending on relative prices and log returns, r_t , as well as the volatility RV_t . We denote cases with numerical instability by (-).

5 Forecasting Liquidity Supply

5.1 Setup

The aim of this section is analyze the model’s forecasting performance in a realistic setting mimicking the situation in financial applications. We consider an investor observing the limit order book at 5-minute snapshots together with the history over the past 10 trading days. It is assumed that during a trading day an investor updates limit order book every 5 minutes and requires producing forecasts for all (5 minutes) intervals of the remainder of the day. Such information might be useful in order to optimally balance order execution during the course of a day. Since we do not exceed beyond the end of the trading day (in order to avoid overnight effects), the forecasting horizon h subsequently declines if we approach market closure. Hence, starting at 10:30, we produce multi-step forecasts for all remaining $h = 66$ intervals during the day. Correspondingly, at 15:50, we are left with a horizon of $h = 1$.

Consequently, the model is re-estimated every five minutes exploiting past information over a fixed window of 10 trading days (including the recent observation). Due to the length of the estimation period, we do not produce forecasts for the first two weeks of our sample but focus on the period between July 22 and August 16, 2002, thereby covering the period of 20 trading days. In accordance with our in-sample results reported in the previous section, we choose the DSFM-Separated approach based on two factors without additional regressors as underlying specification.

A natural benchmark to evaluate our model is the naive forecast. In this context, we assume that the investor has no appropriate prediction model but just uses the current liquidity supply as a forecast for the remainder of the day. More formally, we suppose that our investor can use the following two approaches in order to forecast liquidity supply $\hat{Y}_{t'+h,j}$ at a given time point t' from July 22 at 10:25 until August 16, 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 :

- (i) DSFM approach: Firstly, the factors and factor loadings are estimated using

the DSFM-Separated approach with two factors, $K = 20$ knots used for the B-spline basis functions, and with past 690 observed (de-seasonalized) limit order book curves. More precisely, at time point t' , relative price levels $S_{t'-691:t',j}^b$ and $S_{t'-691:t',j}^a$ and de-seasonalized observed bid and ask sides $Y_{t'-691:t',j}^b$ and $Y_{t'-691:t',j}^a$ enter the estimation procedures. This yields estimates for the bid (ask) side, 66 times per day for each stock, in total 1320 times over 20 days.

Secondly, the specified VAR(p) models for the individual stocks according to Chapter 4, are used to forecast the factor loadings over the forecasting period $\hat{Z}_{t'+h}$. Then, the forecasted factor loadings together with the estimated time-invariant factors \hat{m}_l are used to predict the liquidity supply.

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

The predictions are evaluated using the root mean squared prediction error (RMSPE), given by

$$RMSPE = \sqrt{\frac{1}{hJ} \sum_{h \geq 1} \sum_{j=1}^J \{Y_{h,j} - \sum_{l=0}^L \hat{Z}_{h,l} \hat{m}_l(X_{t,j})\}^2}, \quad (15)$$

where h denotes the forecasting horizon. Note that we do not predict future quotes and thus do not forecast future relative price grids. Under the assumption that quotes themselves follow random walk processes and the spread remains constant, future quotes are predicted using the current one. Consequently, the predicted future relative price grid remains constant.

A graphical illustration of the forecasted limit order book curves and the actually observed ones for each stock on July 22, 2002, at 11:00 and 15:00 is shown in Figure 9.

5.2 Forecasting Results

Figure 10 shows the RMSPEs for each required forecasting horizon h during a trading day implied by the DSFM as well as the naive model. The following results can be summarized: First, overall the DSFM forecasts outperform the naive ones. Nevertheless, the naive forecast is a serious competitor which is hard to beat. This result is not surprising given the high persistence in liquidity supply. Second, the model's forecasting performance is obviously higher on the bid side than on the ask side. This result might be explained by the fact that during the sample period we observe a downward market inducing higher activities on the bid side than on the ask side. This is confirmed by the descriptive statistics shown above. Third, the DSFM outperforms the naive model




Figure 9: Forecasted (dashed) and realized (solid) limit order book curves on July 22, 2002, at 11:00 (upper panel) and 15:00 (lower panel). The naive forecast is depicted by the black solid line.

particularly over horizons up to 1 to 2 hours. For longer horizons, the picture is less clear.

Analyzing average RMSPEs (averaged over all forecasting horizons and both sides of the market) as reported by Table 9 indicate that the overall prediction performance of the DSFM approach is significantly higher than that of the benchmark.

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 9: Average root mean squared prediction errors (RMPSEs) implied by the DSFM approach and the naive model.

5.3 Financial and Economic Applications

The results in the previous section show that the DSFM approach successfully predicts the liquidity supply over various forecasting horizons during a day. In this subsection, we apply these results in two practical examples. The first one is devoted to an order




Figure 10: Root mean squared prediction errors (RMSPEs) implied by the DSFM-Separated approach for the bid side (red) as well as the ask side (blue) and by the naive approach (black) for all intraday forecasting horizons (in hours).

execution strategy, whereas the second one deals with forecasts of demand and supply elasticities.

EXAMPLE 1. (Trading Strategy)

Suppose an institutional investor decides to buy (sell) a certain number of shares over the course of a trading day, starting from 10:30 until 15:55. In order to have comparable results for the individual stocks we choose stock-specific volume sizes and distinguish between three different cases:

- (a) BHP - 100,000 shares; NAB and WOW - 10,000 shares; MIM - 500,000 shares
- (b) BHP - 200,000 shares; NAB and WOW - 20,000 shares; MIM - 1,000,000 shares
- (c) BHP - 300,000 shares; NAB and WOW - 30,000 shares; MIM - 1,500,000 shares.

We assume that after market opening (10:30), the investor has to decide between two execution strategies:

- (i) Splitting the buy (sell) order proportionally over the trading day (i.e., every 5 minutes).
- (ii) Placing one buy (sell) order at a time where the predicted transaction costs are minimal.

To implement these strategies, we consider 20 forecasting days covering the period from July 22 to August 16, 2002. In contrast to the forecasting exercise documented above, we assume that the investor makes his trading decision at 10:30 but does not monitor the market anymore during the day. Consequently, her forecasting horizon covers $h = 66$ periods at each trading day. Using (and keeping constant) the ask and bid quotes prevailing at 10:25, we predict the transaction costs (i.e., the effective costs to buy or to sell the order) using the DSFM for all 5min periods during the day. Then, the order is placed at the period, where expected transaction costs are smallest. Alternatively, according to the proportional trading strategy, the quantities are split in 66 equal market buy (sell) orders per day traded every 5 minutes.

Table 10 gives the stock-specific order sizes expressed as percentages of the average depth prevailing at the best bid and ask quotes. In almost all cases, the order size significantly exceeds the average posted first level depth. Hence, on average, complete all-in-one execution of a buy (sell) order implies that the order has to 'walk up (down)' the book. This increases transaction costs compared to a splitting strategy where the split orders are ideally executed against the first level depth. The following analysis will show whether the suggested model is successful in predicting optimal execution times where the book is sufficiently deep in order to make an all-in-one execution profitable.

Case	Sell				Buy			
	BHP	NAB	MIM	WOW	BHP	NAB	MIM	WOW
(a)	338.34	204.42	47.20	100.99	311.19	179.08	44.95	99.80
(b)	676.69	408.85	94.40	201.99	622.39	358.17	89.90	199.60
(c)	1015.03	613.27	141.60	302.98	933.58	537.25	134.85	299.39

Table 10: Order sizes expressed as percentages of the average depth posted at the best bid and ask quotes over the period from July 22 to August 16, 2002.

Tables 11 and 12 summarize the average daily transaction costs, expressed in AUD, for selling and buying different number of shares using both strategies. Note that the realized transaction costs are naturally driven by the actual depth as well as the underlying prevailing quotes. While the first component is predicted by our model, the latter is unpredictable. Since these unpredictable quote movements affect both trading strategies and we can expect that over the analyzed period positive and negative movements will cancel out, our findings should not be systematically affected by this component which is not controlled by our model.

The realized transaction costs implied by both execution strategies show that the DSFM-drive strategy performs equally well and in several cases even outperforms the splitting strategy. This is remarkable given the fact that the DSFM strategy represents the extreme case where an investor completely executes the volume using just one market order. Typically, such a strategy is quite expensive and is avoided by traders. Hence, in

Strategy	BHP	NAB	MIM	WOW
Proportional				
(a)	916,247	334,711	576,123	121,288
(b)	1,832,438	669,402	1,152,243	242,567
(c)	2,748,587	1,004,100	1,728,364	363,846
DSFM approach				
(a)	914,864	338,573	589,500	121,885
(b)	1,827,103	677,107	1,157,000	244,020
(c)	2,736,822	1,015,343	1,718,012	365,921

Table 11: Average daily transaction costs in AUD for selling shares according to different cases, using the proportional splitting strategy and the DSFM approach from July 22 to August 16, 2002.

Strategy	BHP	NAB	MIM	WOW
Proportional				
(a)	922,907	335,115	591,138	121,646
(b)	1,845,952	670,215	1,182,274	243,292
(c)	2,769,216	1,005,345	1,773,410	364,943
DSFM approach				
(a)	913,852	333,925	593,000	121,390
(b)	1,828,862	677,352	1,191,188	243,160
(c)	2,747,783	1,014,982	1,785,581	367,138

Table 12: Average daily transaction costs in AUD for buying shares according to different cases, using the proportional splitting strategy and the DSFM approach from July 22 to August 16, 2002.

practice, one would try to further reduce transaction costs by executing the order not at one but at several optimally chosen (prediction based) time points. In this respect, the all-in-one strategy can be seen as a rather conservative "practitioners' test" of the model. However, our findings indicate that the model is successful in predicting times where the market is sufficiently deep in order to execute a big order. The fact that the model performs reasonably well is promising for more elaborate practical applications of the DSFM. Moreover, note that the reported results are valid under the assumption that there are no transaction fees. Actually, in practice, a proportional splitting strategy induces higher transaction costs as a complete execution via a market order. This component is not taken into account here and would even increase the performance of the DSFM-based execution strategy.

EXAMPLE 2. (Demand and Supply Elasticity)

A straightforward dimension-less measure for the order book slope is the curve's elasticity which we compute at best bid ($\tilde{S}_{t',101}^b$) and best ask prices ($\tilde{S}_{t',1}^a$) as

$$\hat{E}_{t'+h}^d = \frac{\hat{Y}_{t'+h,1}^b - \hat{Y}_{t'+h,101}^b}{\hat{Y}_{t'+h,101}^b} / \frac{\tilde{S}_{t',1}^b - \tilde{S}_{t',101}^b}{\tilde{S}_{t',101}^b}, \quad (16)$$

$$\hat{E}_{t'+h}^s = \frac{\hat{Y}_{t'+h,101}^a - \hat{Y}_{t'+h,1}^a}{\hat{Y}_{t'+h,1}^a} / \frac{\tilde{S}_{t',101}^a - \tilde{S}_{t',1}^a}{\tilde{S}_{t',1}^a}, \quad (17)$$

for the demand (bid) and supply (ask) side, respectively. The elasticity is a measure for the marginal trading costs reflecting the curve's curvature.

Suppose at 10:30 an investor aims predicting the demand and supply elasticity at best bid and best ask prices for all 5-min intervals during the trading day covering the forecast horizons $h = 1, \dots, 66$. As above, the forecasts are computed using the last 10 trading days. Since we are not forecasting the price process, the last observed ask and bid quotes are used for prediction. Figure 11 shows the 10:30 predictions of demand and supply elasticities at best bid and best ask prices during all trading days. We observe that marginal trading costs exhibit significant variations over time. The fact that predicted elasticities reveal temporarily trending patterns might be used for improving trading strategies.

Consider the case of NAB on July 24 and July 30, 2002. We observe that the demand elasticities (in absolute terms) are increasing on the first day, and decreasing on the second day. Practically, it would be better to sell shares late on July 24, and early on July 30, under the assumption that the price does not change significantly over both trading days. The supply elasticities show converse patterns across the days. Consequently, it would be advisable to buy shares early on July 24, and late on July 30, provided that the prices remain unchanged. While this section aims to illustrate possible applications of the DSFM approach, more detailed elaborations of dynamic strategies are beyond the scope of the paper.




Figure 11: Predicted demand and supply elasticities at best bid (red) and best ask prices (blue) from July 22 to August 2, 2002 (upper panels) and from August 5 to August 16, 2002 (lower panels)

6 Conclusions

In this paper, we propose a dynamic semiparametric factor model (DSFM) for limit order book curves. The main idea of the DSFM as proposed by Fengler et al. (2007), Brüggemann et al. (2008), Park et al. (2009) and Cao et al. (2009) is to capture the order curve's spatial structure across various price levels using a factor decomposition which is estimated nonparametrically. To capture the order book's time variations the corresponding factor loadings are modelled using a vector autoregressive (VAR) model. The framework is flexible though parsimonious and turns out to provide a powerful way to reduce the high dimension of the book and to extract the relevant underlying information.

The model is applied to four stocks traded at the Australian Stock Exchange (ASX). It is shown that two underlying factors can explain up to 95% of in-sample variations of ask and bid liquidity supply. While the first factor captures the overall order curve's slope, the second factor is associated with the curve's curvature. The extracted factor loadings reveal highly persistent though weakly stationary dynamics which are successfully captured by VAR specifications. Though it is shown that the order curves' shapes are driven by explanatory variables reflecting the recent liquidity demand, volatility as well as mid-quote returns, these variables do not improve the model's goodness-of-fit.

Employing the DSFM approach in an extensive and realistic out-of-sample forecasting exercise we show that the model successfully predicts the liquidity supply over various forecasting horizons during a trading day and outperforms a naive approach. Using the forecasting results in a trading strategy it is shown that order execution costs can be reduced if orders are optimally placed according to predictions of liquidity supply. In particular, it turns out that optimal order placement in periods of high liquidity results in smaller transaction costs than in the case of a proportional splitting over time.

In summary, the DSFM approach is suitable for modelling and forecasting the limit order book. Since it is computationally tractable it can serve as a valuable building block for automated trading models.

References

- Ahn, H. J., Bae, K. H. and Chan, K. (2001). Limit orders, depth, and volatility: evidence from the stock exchange of hong kong, *Journal of Finance* **56**: 767–788.
- Baruch, S. (2005). Who benefits from an open limit-order book?, *Journal of Business* **78**: 1267–1306.
- Biais, B., Hillion, P. and Spatt, C. (1995). An empirical analysis of the limit order book and the order flow in the paris bourse, *Journal of Finance* **50**: 1655–1689.
- Bloomfield, R., O’Hara, M. and Saar, G. (2005). The ”make or take” decision in an electronic market: evidence on the evolution of liquidity., *Journal of Financial Economics* **75**: 165–200.
- Boehmer, E., Saar, G. and Yu, L. (2005). Lifting the veil: an analysis of pre-trade transparency at the nyse, *Journal of Finance* **60**: 783–815.
- Brownlees, C. T., Cipollini, F. and Giampiero, M. G. (2009). Intra-daily volume modeling and prediction for algorithmic trading, *Discussion paper*, Stern School of Business.
- Brüggemann, R., Härdle, W., Mungo, J. and Trenkler, C. (2008). Var modelling for dynamic semiparametric factors of volatility strings, *Journal of Business* **6**(3): 361–381.
- Cao, J., Härdle, W. and Mungo, J. (2009). 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.
- Chacko, G. C., Jurek, J. W. and Stafford, E. (2008). The price of immediacy, *Journal of Finance* **63**: 1253–1290.
- Comerton-Forde, C. and Tang, K. M. (2009). Anonymity, liquidity and fragmentation, *Journal of Financial Markets* **12**: 337–367.

References

- Engle, R. F. and Ferstenberg, R. (2007). Execution risk, *Journal of Portfolio Management* **33**: 34–45.
- Eom, K. S., Ok, J. and Park, J. H. (2007). Pre-trade transparency and market quality, *Journal of Financial Markets* **10**: 319–341.
- Fengler, M. R., Härdle, W. and Mammen, E. (2007). A dynamic semiparametric factor model for implied volatility string dynamics, *Journal of Financial Econometrics* **5**(2): 189–218.
- Gallant, A. R. (1981). On the bias of flexible functional forms and an essentially unbiased form, *Journal of Econometrics* **15**: 211–245.
- Garvey, R. and Wu, F. (2009). Intraday time and order execution quality dimensions, *Journal of Financial Markets* **12**: 203–228.
- Glosten, L. (1994). Is the electronic limit order book inevitable, *Journal of Finance* **49**(4): 1127–1161.
- Goyenko, R. Y., Holden, C. W. and Trzcinka, C. A. (2009). Do liquidity measures measure liquidity?, *Journal of Financial Economics* **92**: 153–181.
- Griffiths, M. D., Smith, B. F., Turnbull, D. A. S. and White, R. W. (2000). The costs and determinants of order aggressiveness, *Journal of Financial Economics* **56**: 65–88.
- Hall, A. D. and Hautsch, N. (2006). Order aggressiveness and order book dynamics, *Empirical Economics* **30**: 973–1005.
- Hall, A. D. and Hautsch, N. (2007). Modelling the buy and sell intensity in a limit order book market, *Journal of Financial Markets* **10**(3): 249–286.
- Hasbrouck, J. (2009). Trading costs and returns for u.s. equities: Estimating effective costs from daily data, *Journal of Finance* **64**: 1445–1477.
- Hasbrouck, J. and Saar, G. (2009). Technology and liquidity provision: The blurring of traditional definitions, *Journal of Financial Markets* **12**: 143–172.
- Hautsch, N. (2008). Capturing common components in high-frequency financial time series: A multivariate stochastic multiplicative error model, *Journal of Economic Dynamics & Control* **32**: 3978–4009.
- Hautsch, N. and Huang, R. (2009). The market impact of a limit order, *Discussion Paper forthcoming*, Collaborative Research Center 649 "Economic Risk", Humboldt-Universität zu Berlin.
- Hollifield, B., Miller, R. A. and Sandås, P. (2004). Empirical analysis of limit order markets, *Review of Economic Studies* **71**: 1027–1063.
- Johansen, S. (1991). Estimation and hypothesis testing of cointegration vectors in gaussian vector autoregressive models, *Econometrica* **59**(6): 1551–1580.

References

- Johnson, T. C. (2008). Volume, liquidity and liquidity risk, *Journal of Financial Economics* **87**: 388–417.
- Kwiatkowski, D., Phillips, P. C. B., Schmidt, P. and Shin, Y. (1992). 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.
- Liu, W.-M. (2009). Monitoring and limit order submission risks, *Journal of Financial Markets* **12**: 107–141.
- Nelson, C. R. and Siegel, A. F. (1987). Parsimonious modelling of yield curves, *Journal of Business* **60**: 473–489.
- Park, B., Mammen, E., Härdle, W. and Borak, S. (2009). Time series modelling with semiparametric factor dynamics, *Journal of the American Statistical Association* **104**(485): 284–298.
- Ranaldo, A. (2004). Order aggressiveness in limit order book markets, *Journal of Financial Markets* **7**: 53–74.
- Schmidt, P. and Phillips, P. C. B. (1992). Lm tests for a unit root in the presence of deterministic trends, *Oxford Bulletin of Economics and Statistics* **54**: 257–287.

CFS Working Paper Series:

No.	Author(s)	Title
2009/17	John F. Cogan Tobias Cwik John B. Taylor Volker Wieland	New Keynesian versus Old Keynesian Government Spending Multipliers
2009/16	Christopher D. Carroll	Precautionary Saving and the Marginal Propensity to Consume Out of Permanent Income
2009/15	Christopher D. Carroll Olivier Jeanne	A Tractable Model of Precautionary Reserves, Net Foreign Assets, or Sovereign Wealth Funds
2009/14	Christopher D. Carroll Patrick Toche	A Tractable Model of Buffer Stock Saving
2009/13	Jan Pieter Krahnen Günter Franke	Instabile Finanzmärkte
2009/12	Christopher D. Carroll Jiri Slacalek	The American Consumer: Reforming, Or Just Resting?
2009/11	Jan Pieter Krahnen Christian Wilde	CDOs and Systematic Risk: Why bond ratings are inadequate
2009/10	Peter Gomber Markus Gsell	Algorithmic Trading Engines Versus Human Traders – Do They Behave Different in Securities Markets?
2009/09	Christian Laux Christian Leuz	The Crisis of Fair Value Accounting: Making Sense of the Recent Debate
2009/08	Annamaria Lusardi Peter Tufano	Debt Literacy, Financial Experiences, and Overindebtedness

Copies of working papers can be downloaded at <http://www.ifk-cfs.de>