Predicting Liquidity from Order Book Data

Knut Griese [*] and Alexander Kempf [†]

05.10.2004

Abstract

In this paper we investigate the functional form and temporal dynamics of the price impact function. Knowledge of the form and dynamics of the price impact function is important because it has serious implications for optimal liquidation strategies of large investors. Our empirical analysis shows that the functional form as well as the level of liquidity exhibits considerable variation over time. Furthermore, we find strong predictability in changes of the price impact function. Among competing models to predict future price impact functions a two parameter power function with the dynamics of the parameters modeled by a VAR(1) process shows the best forecast accuracy.

[*] Universität zu Köln
[†] Universität zu Köln

1 Introduction

Transaction costs in the trading process can seriously affect the performance of the trading strategies of institutional investors. The most significant trading cost institutional investors bear is price impact associated with large trades. Institutional investors therefore typically split their orders to avoid large price impact costs. Evidence for this behavior can be found for example in Chan and Lakonishok (1995) or Keim and Madhavan (1995).

In order to optimally split the order and to assess the corresponding transaction costs, investors must assess the specific characteristics of the liquidity of the market. Especially, the functional form and temporal dynamics of the price impact function have important implications for the optimal strategy. If, for example, the functional form of the price impact function is concave, incentives to split the order are weaker than in the case of a convex relationship.[1] As the decision problem requires the investor to evaluate the trade off between current and future price impact, it is not sufficient to know the current price impact function. In addition, the investor has to form expectations about the future price impact function. If, for example, an investor expects liquidity to increase in the future, he will ceteris paribus postpone some of his trades to avoid trading costs.

In this paper we analyze the functional form and temporal dynamics of the whole price impact function in order to address two questions. Firstly, how does the shape of the price impact function look like? Secondly, how pronounced are temporal variations of the price impact function and to what extent can changes in liquidity be predicted from the past? To answer these

[1] The theoretical literature on optimal order splitting strategies restricts price impact to be linear in trade size and to be deterministic, see for example Huberman and Stanzl (2003), Bertsimas and Lo (1998) and Almgren and Chriss (1998). Huberman and Stanzl (2004) argue that price impact should be linear in trade size to rule out (quasi-) arbitrage opportunities.

questions, we model and predict the price impact function employing order book data from the XETRA automated exchange. In the XETRA limit order market, the whole price impact function is visible to market participants. The data are therefore especially well suited. They reflect the relevant trading opportunities of investors and allow to analyze temporal variations of the price impact function.

We contribute in two ways to the existing literature. Firstly, we provide evidence of the functional form of the price impact function for different points of time. In this respect we differ from the large number of empirical studies which analyze the functional form of the price impact by employing transaction data.[2] These studies are not able to characterize the price impact function at a point of time and, as we argue, may be misleading in the context of a trader assessing the potential price impact of his order. We also differ from studies who employ order book data to analyze the functional form. These studies only investigate the shape of the average order book.[3]

Secondly, we propose different approaches to predict the whole price impact function. The dynamics of the price impact function has not been addressed in the literature so far. In addition we compare the accuracy of different forecasting models for the price impact function.

The main results of the empirical analysis can be summarized as follows. There is considerable time variation of the price impact function. This holds true for the level of liquidity as well as for the functional form. The functional form of the price impact function shows strong variation, with the shape being concave in about 50% of the cases and convex in about 50%. In addition we find strong predictability in liquidity. This holds true for the level as well as the shape of the price impact function. When evaluating the out of sample

[2]See for example Hasbrouck (1991), Hausman, Lo, and MacKinlay (1992) or Kempf and Korn (1999).

[3]See for example Maslov and Mills (2001) or Weber and Rosenow (2003).

forecast accuracy of different models, a linear price impact function with the parameter modeled as an AR(p) process works well in predicting the future price impact function. However a power function with the dynamics of the parameters modeled by a VAR(1)-process is shown to be significantly more accurate in predicting future price impact functions.

The paper is organized as follows: In section 2 we survey the empirical literature on nonlinearity of the price impact function and the literature on the dynamics of liquidity measures over time. In section 3 we derive optimal trading strategies for nonlinear price impact functions. In section 4 we describe the data. In section 5 we explain the empirical methodology and present the results. Section 6 concludes.

2 Literature Review

In this section we review the empirical literature which investigates the functional form of the price impact function as well as the empirical literature which investigates the dynamics of liquidity measures.

Empirical evidence on the functional form

There are numerous empirical studies which analyze the functional form of the price impact function. A common feature is that these studies almost solely rely on transactions data. Typically these studies employing transactions data detect a significantly concave relationship between order size and price impact, see among others Hasbrouck (1991), Madhavan and Smidt (1991), Hausman, Lo, and MacKinlay (1992), Kempf and Korn (1999), Knez and Ready (1996), Keim and Madhavan (1996), Algert (1990), Chen, Stanzl, and Watanabe (2002) or Spierdijk, Nijman, and Soest (2004). Only a small minority of studies find no significant deviation from linearity, see for example Engle and Lange (2001), Breen, Hodrick, and Korajczyk (2002) or Sadka

(2003). The studies which analyze the price impact from order book data get completely different results. Maslov and Mills (2001), Weber and Rosenow (2003) and Coppejans, Domowitz, and Madhavan (2003) detect a convex shape of the order book employing limit order book data from the ISLAND ECN respectively from the Swedish stock exchange. Interestingly Weber and Rosenow (2003) find a concave function when they rely on transactions data. Biais, Hillion, and Spatt (1995) detect no systematic deviation from linearity using order book data from the Paris Bourse.

We argue that the results of studies based on transactions data may be misleading in the context of a trader assessing the potential price impact of his order. The reason is that these studies are based on realized transactions and not on trading opportunities. The concentration on realized transactions may lead to systematically biased estimates of potential price impact. If liquidity is time varying and predictable for traders, they respond to variations in liquidity and place orders when the price impact is small. So it is not clear if the results from transactions data are of relevance for an optimizing trader, rather they may already reflect optimized behavior.[4] A further shortcoming of liquidity assessment based on transactions data is that temporal variations in liquidity can not be captured adequately, as the price impact function has to be estimated from the price impact of transactions of different size at different points of time.

The existing studies analyzing order book data aggregate the data over time in order to analyze the functional form. They do therefore not analyze the functional form at different points of time or the variability and dynamics of

[4] Such optimizing behavior is analyzed for example in Admati and Pfleiderer (1988). In Admati and Pfleiderer (1988) liquidity traders can choose the timing of their transactions in order to minimize the expected cost of their transactions. They prefer to trade more when their trading has little effect on prices. Coppejans, Domowitz, and Madhavan (2003) and Gomber, Schweickert, and Theissen (2004) provide empirical evidence that traders time their market orders, i.e. they trade more when liquidity is high.

the functional form.

Empirical evidence on liquidity dynamics

Recent empirical studies show that liquidity varies substantially over time and that liquidity changes seem to be at least partially predictable. These studies employ various liquidity measures and investigate quite different research questions.

Huberman and Halka (2001) study the commonality in unexpected changes of several liquidity proxies across different groups of stocks. These daily liquidity proxies, as spread or depth at the best quotes, show up to be highly autocorrelated. In order to decompose the changes in liquidity into expected and unexpected components Huberman and Halka (2001) estimate time series models for the liquidity proxies (averaged across different groups of stocks), notably AR(p) processes.

Amihud (2002) studies the effect of expected market illiquidity on expected stock excess returns in a time series context. His liquidity proxy is the daily ratio of absolute stock return to dollar trading volume (averaged across stocks) and is interpreted as a measure of price impact. Expected illiquidity is obtained by estimating an AR(1)-model for the liquidity measure. The liquidity measure is highly autocorrelated.

There are a few studies which use limit order book data to study the dynamics of liquidity measures. These data allow especially for an intraday analysis of the variation of liquidity.

Coppejans, Domowitz, and Madhavan (2003) study the dynamics of liquidity and its relation to returns and volatility employing limit order book data for Swedish stock index futures. They examine the time variation in depth of the limit order book at a 5-minute-frequency. They document considerable variation in observed depth as well as mean reversion in liquidity.

Beltran-Lopez, Giot, and Grammig (2002) extensively analyze time varying means and variances of a trading cost measure constructed from limit order book data from the XETRA automated exchange. They measure liquidity as hypothetical bid-ask returns which are closely related to the Irvine/Benston/Kandel(2000) cost of round trip measure and takes into account the price impact of a market order of given volume to get executed immediately. They model the dynamics of their liquidity measure by an AR(p) process with time varying volatility to account for conditional heteroscedasticity. In their analysis they also consider the effect of time of day patterns and international stock markets on the dynamics of the liquidity measure.

Kumar (2003) focuses on the forecast of a price impact cost measure employing a limit order book dataset from the Indian stock exchange. The impact costs are calculated for a given market order volume as the percentage change between the weighted average execution price and the quote midpoint. The dynamics of impact cost (sampled at hourly frequency) is captured by an ARMA(1,1)-model. He conducts out of sample tests in which the time series model of impact cost is shown to perform better compared to the naive forecast.

These studies give important insights into the time varying nature of liquidity and the possibility to forecast transaction costs. However, these studies leave open important aspects of liquidity dynamics and give an investor limited guidance when he is concerned with optimally breaking up his order. Only the measures in Beltran-Lopez, Giot, and Grammig (2002) and Kumar (2003) depend on the size of an order and assess liquidity from the relevant trading opportunities for investors, the limit order book. The major shortcoming of all of the employed liquidity measures is that they are valid only for a specific order size. For a trader concerned with the splitting of his order to minimize price impact, the order size is a decision variable. For this decision he needs to know the hypothetical prices for all possible order sizes today and a prediction of the prices for all possible order sizes tomorrow.[5] We therefore differ from the previous literature in that we do not only model the time series properties of a price impact measure for a fixed order size but model the functional form and the dynamics of the whole price impact function.[6]

3 Optimal liquidation strategies

In this chapter we show that the optimal strategy of an investor crucially depends on the shape of the price impact function. We consider an investor who wants to acquire a block of Q shares over two trading dates.[7] With his purchases the investor affects the transaction prices. The effect of trade size q_t in period t on the marginal price p_t is captured by the price impact function. We specify the price impact function as a power function with general exponent λ_2:

$$p_t = \widetilde{p_t} + \lambda_1 q_t^{\lambda_2} \tag{1}$$

$\widetilde{p_t}$ is the price which would prevail in the absence of any market impact (reservation price). For $\lambda_2 = 1$ the price impact function (1) is linear in trade size, for $\lambda_2 > 1$ it is convex and for $\lambda_2 < 1$ it is concave.

[5] If the price impact function was linear all the time, then the measure of Kumar (2003) e.g. would be a sufficient statistic to characterize the price impact function.

[6] The measures in Kumar (2003) and Beltran-Lopez, Giot, and Grammig (2002) consider the average price impact costs for a given order size. For an optimizing trader the marginal cost of the last unit is relevant for his decision. For this reason we model the price impact function in terms of marginal prices and not in terms of average costs.

[7] In order to develop the main consequences of different price impact functions, we restrict the analysis in this section to two periods. For the T period problem, we only obtain numerical results for concave and convex price impact. The second motivation for the restriction to two trading dates is that for $T > 2$ manipulation possibilities as discussed in Huberman and Stanzl (2004) arise.

Since the market order of size q_t is executed against limit orders with different prices along the price impact function (1), the total cost of acquiring q_t is given by:[8]

$$\int_0^{q_t} (\widetilde{p}_t + \lambda_1 x^{\lambda_2})dx = (\widetilde{p}_t + \frac{\lambda_1}{\lambda_2 + 1} q_t^{\lambda_2}) q_t \tag{2}$$

Since investors are able to split orders over time, we have to specify the dynamics of the reservation price. We assume that trade size has a permanent effect on the reservation price. A fraction α of the price impact is temporary and vanishes after one period so that $(1 - \alpha)$ of the price impact is reflected in the reservation price of the next period:

$$\widetilde{p}_t = \widetilde{p}_{t-1} + (1 - \alpha)(p_{t-1} - \widetilde{p}_{t-1}) + \varepsilon_t \tag{3}$$

ε_t is news incorporated into the price. ε_t is assumed to be white noise. The trader does know ε_t when he submits his market order in t.

The investor's objective is to minimize the expected cost C_1 of acquiring Q:

$$Min_{\{q_t\}} C_1 \tag{4}$$

with

$$C_1 = E_1[\sum_{t=1}^{2} \int_0^{q_t} (\widetilde{p}_t + \lambda_1 x^{\lambda_2})dx] \tag{5}$$

[8]In this respect we differ to Huberman and Stanzl (2003), Bertsimas and Lo (1998) or Almgren and Chriss (1998) who assume the execution price of the whole order of size q_t is the marginal price. The assumption of a unique execution price is more appropriate for a dealer market. As will be seen later, our assumption has consequences for optimal order splitting.

subject to (3) and the constraint

$$\sum_{t=1}^{2} q_t = Q \tag{6}$$

The problem is to find the optimal trading strategy $\{q_t\}$ that minimizes expected total cost. The first order condition is given by:[9]

$$\frac{\delta C_1}{\delta q_1} = \lambda_1 q_1^{\lambda_2} + (1-\alpha)\lambda_1 q_1^{\lambda_2 - 1}(\lambda_2 Q - (\lambda_2 + 1)q_1) - \lambda_1 (Q - q_1)^{\lambda_2} \doteq 0 \tag{7}$$

Closed form solutions for q_1 can be obtained only for some parameter values.[10] To analyze the consequences of different shapes of the function on the optimal strategy, we present closed form solutions for linear price impact, concave price impact with $\lambda_2 = 1/2$ and convex price impact with $\lambda_2 = 2$.[11]

Linear price impact

If the price impact is linear according to

$$p_t = \widetilde{p}_t + \lambda_1 \, q_t \tag{8}$$

the naive strategy of evenly splitting the order turns out to be the optimal strategy for $\alpha > 0$:[12]

$$q_1^* = \frac{1}{2} Q \tag{9}$$

[9] For a derivation see the appendix.

[10] These special cases are $\lambda_2 = 1/4, 1/3, 1/2, 2/3, 3/4, 1, 4/3, 3/2, 2, 3, 4,$ see the appendix.

[11] $\lambda_2 = 1/2$ is the estimate Maslov and Mills (2001) get, when they use transactions data and $\lambda_2 = 2$ when they estimate the price impact from order book data.

[12] For a derivation see the Appendix.

The minimum expected costs V_1 are an explicit function of the total traded quantity Q, the reservation price $\widetilde{p}_1$, the price impact parameter λ_1 and the fraction of the temporary price impact α:

$$V_1 = \widetilde{p}_1 Q + \frac{1}{2}\lambda_1 Q^2 - \frac{\alpha}{4}\lambda_1 Q^2 \tag{10}$$

For the special case of solely permanent price impact ($\alpha = 0$), the splitting across trading dates is arbitrary.[13]

Concave price impact

For a square root price impact function

$$p_t = \widetilde{p}_t + \lambda_1 \operatorname{sgn}(q_t)\sqrt{|q_t|} \tag{11}$$

the optimal trade size for $\alpha > 0$ is given by:

$$q_1^* = \frac{3 - 4\alpha + 3\alpha^2 + 2\sqrt{1 - 2\alpha + 2\alpha^2}}{5 + 9\alpha^2 - 6\alpha} Q \tag{12}$$

It can be shown that the investor trades now more in period 1, i.e. $q_1^* > \frac{1}{2}$ if $\alpha < 1$. q_1^* decreases with the fraction of temporary price impact, α, and approaches $q_1^* = \frac{1}{2}$ for $\alpha = 1$.

The minimum of expected costs is characterized by:

$$V_1 = \widetilde{p}_1 Q + (1-\alpha)\lambda_1\sqrt{q_1^*}Q + (\alpha - \frac{1}{3})\lambda_1 (q_1^*)^{\frac{3}{2}} + \frac{2}{3}\lambda_1 (Q - q_1^*)^{\frac{3}{2}} \tag{13}$$

For the special case $\alpha = 0$ the investor is indifferent between trading the total volume in period 1 or period 2:

[13] This result differs from Bertsimas and Lo (1998) and is a consequence of different execution prices as reflected in (2).

$$q_1^* \in \{0, Q\} \tag{14}$$

Convex price impact

For a power function with exponent 2

$$p_t = \widetilde{p}_t + \lambda_1 \operatorname{sgn}(q_t) q_t^2 \tag{15}$$

the optimal trade size is given by:

$$q_1^* = \frac{-2 + \alpha + \sqrt{1 - \alpha + \alpha^2}}{3(\alpha - 1)} Q \tag{16}$$

The investor trades more in the second period, $q_1^* < \frac{1}{2}$. q_1^* increaes with the fraction of temporary price impact, α, and approaches $q_1^* = \frac{1}{2}$ for $\alpha = 1$.

The minimum of expected costs is given by:

$$V_1 = \widetilde{p}_1 Q + \frac{1}{3}\lambda_1 q_1^{*3} + (1 - \alpha)\lambda_1 q_1^{*2}(Q - q_1^*) + \frac{1}{3}\lambda_1 (Q - q_1^*)^3 \tag{17}$$

In figure 1, the optimal fraction traded in the first date is plotted against the fraction of temporary price impact α. We distinguish between convex, linear and concave price impact functions.

Our example with three stylized price impact functions shows that the shape of the price impact function has a significant impact on the optimal strategy. For a concave price impact the investor trades more in the first period, linear price impact leads to evenly distributed trade sizes and for a convex price impact the investor trades more in the second period. Temporary price impact favors splitting the order more equally. If the price impact is solely temporary, the naive strategy is optimal irrespective of the specific functional form. Of

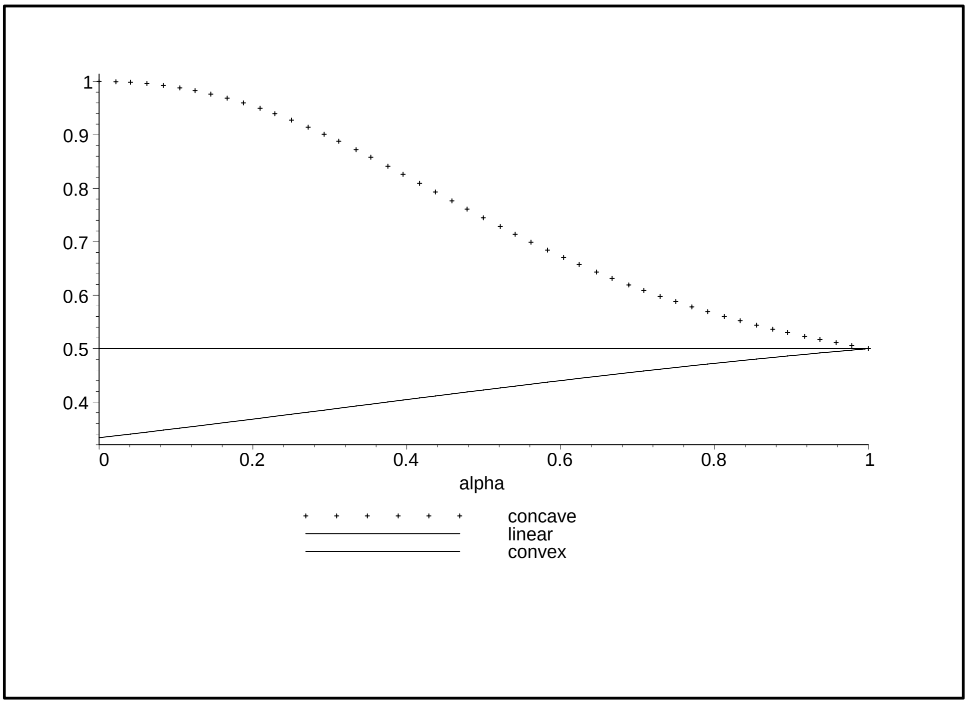

Figure 1: Optimal fraction q_1^*/Q for different price impact functions

course, a false assessment of the shape of the price impact function and the fraction of temporary price impact leads to increased trading costs.

4 Data

In our empirical analysis we use data from the automated auction system Xetra which is maintained by the German Stock Exchange.

Our dataset contains complete information about XETRA order book events, i.e. entries, cancellations, revisions, expirations and executions of market and limit orders for the three blue chip stocks DaimlerChrysler, Deutsche Telekom and SAP.[14] XETRA is the dominating system in the provision of liquidity for German blue chip stocks. A floor system competing with XETRA operates at the Frankfurt stock exchange. In 1999 the relative share of trading volume for XETRA was about 90% for the three stocks in our sample.[15]

The sample covers the 65 trading days from August 2, 1999 to October 29,1999. During this period in XETRA the whole content of the order book was visible for the market participants. This implies that the potential price impact of a market order is exactly known to the trader before submitting his market order.[16] Market orders at the XETRA system hit the order book until complete fill. [17] XETRA trading hours lasted from 8.30 a.m. to 5.00 p.m. In the middle of the sample period (September 20,1999) trading hours were shifted to 9.00 a.m. to 5.30 p.m.

Trading on XETRA is based on a continuous double auction mechanism. A computerized order book keeps track of all incoming market and limit orders. Orders are automatically matched based on clearly defined rules of price and

[14]We are grateful to Joachim Grammig and Helena Beltran-Lopez for providing us with the data and the programs to reconstruct the order book. Beltran-Lopez, Giot, and Grammig (2003) provide a very detailled descritption of the XETRA trading mechanism and the data.

[15]In 1999 there was no separate block trading system implemented for the XETRA trading system. In March 2001 a crossing system for the trading of large blocks called XETRA XXL was introduced.

[16]Since October 2000 hidden orders (often referred to as iceberg orders) were allowed in XETRA. Hidden orders are not visible in the displayed order book. This change in market transparency has the implication that a trader submitting a market order cannot asses the price impact of his order exactly ex ante due to the fact that his order may be executed against a hidden order. In this case the trader has to form expectations about the hidden part of the order book.

[17]This differs from e.g. the Paris Bourse where market orders are not necessarily immediately and fully executed. At the Paris Bourse market orders are executed at the best price. If the entire quantity of a market order is not filled at the best price in the order book, the remaining shares are transformed into a limit order at the transaction price. In October 2000 XETRA introduced so called market to limit orders which are executed against the best prices and then transformed into a limit order.

time priority. [18] Based on these rules of the XETRA trading protocol and the event histories a real time reconstruction of the order book sequences is performed. Starting from an initial state of the order book, each change in the order book implied by entry, partial or full fill, cancellation and expiration of market and limit orders is tracked. From the resulting real-time sequences of order books, snapshots at ten minute intervals during the continuous trading hours were taken.

5 Empirical analysis

The former section illustrated that the investor needs to have an estimate about the price impact function for each point of time. In this section we model the price impact function and allow for time varying price impact functions. Our basic concern is to get a good prediction concerning the future price impact function.

A good model of order book dynamics is able to reproduce the different possible shapes of the price impact function and to predict future liquidity accurately. In the following chapter different ways to model and predict the price impact function will be presented. We evaluate the models by their forecast accuracy and conduct test to discriminate between the competing models.

Our approach to predict the price impact function resembles a new approach for the prediction of yield curves. Diebold and Li (2003) predict the yield curve by modeling the three factors of Nelson and Siegel (1987) as autoregressive processes. The prediction of the parameters is utilized for a prediction of the yield curve. The approach of Diebold and Li (2003) works especially well when applied to out of sample predictions of the yield curve.

[18] In addition there are call auctions at the open, at mid-day and at the close.

The approach of modeling the form and dynamics of the price impact function consists of two steps. In the first step the order book is approximated by a parsimonious parametrized function at each point of time. In section 5.1 we consider a linear function, in section 5.2 a more flexible nonlinear function according to (1). In the second step we model the time series behavior of the parameters which characterize the price impact function with the aim to get a prediction of the price impact function. In the case of the linear model we estimate autoregressive models for the slope parameter, in the case of the nonlinear model we estimate vector autoregressive models for the two model parameters.

The separate analysis of the more restrictive linear function is justified by the fact that a priori it is not clear if the more flexible nonlinear function has superior forecast ability concerning the price impact function. It is well known that more parsimonious models often show superior forecast ability. In addition the linear model can serve as a benchmark model for the nonlinear model.

In each section we evaluate the performance of the models by comparing the predicted price impact functions to the actual price impact functions employing several measures of forecast accuracy. The forecast accuracy measures in section 5.1 and 5.2 are defined for various order sizes and indicate no strict dominance of either model. Therefore in section 5.3 we propose a measure which is independent of order size and employ a statistical test to directly compare the predictive ability of the linear model to the nonlinear model.

It is well known that liquidity measures exhibit time of the day patterns. Therefore we also estimated different models to incorporate seasonality into the linear and the nonlinear model. The basic results concerning the size and the significance of the parameter estimates of the autoregressive structure were not affected by the incorporation of time of the day effects. More importantly when evaluated by their predictive power none of the extended models showed superior forecast accuracy, often the forecast accuracy was even worse. We therefore abstain from presenting the results.

5.1 The linear model

In this section we will examine a linear price impact function which serves as a benchmark for the analysis of a more flexible functional form of the price impact function in 5.2. In 5.1.1 we will first outline the methodology in modeling and forecasting the price impact function. In 5.1.2 we present the results.

5.1.1 Methodology

Fitting the price impact function

We take 3315 ten minute snapshots of the order book to fit a linear price impact function according to:[19]

$$p_{t,i} = \widetilde{p}_t + \lambda_t q_i + v_{t,i} \tag{18}$$

$p_{t,i}$ are the hypothetical marginal transaction prices for different trade sizes q_i which are to be calculated from the state of the order book at time t. The midquote before the trade takes place serves as a proxy for the reservation price $\widetilde{p}_t$. v_t is the disturbance which is assumed to be $N(0, \sigma^2_{v_t})$. As the order book has different slopes at different points of time, we allow for time variation of the liquidity parameter λ.

[19]Equation (18) could be extended to account for a spread at zero quantity (see Glosten (1994), Proposition 3 for the existence of a spread at zero quantity in a limit order market). We do not consider this because the zero quantity spread is negligible in practice.

We estimate λ_t from (18) each period by OLS, treating the hypothetical price impacts $p_{t,i} - \widetilde{p}_t$ for z ex ante specified equidistant order sizes $q_1, ..., q_z$ as observations.

In implementing the approach we have to define the order sizes $q_1, ..., q_z$ in advance. They are chosen to be the same for all dates. We choose $q_z =$ 100000 shares for Daimler Chrysler and Deutsche Telekom and $q_z = 10000$ shares for SAP, which amounts to roughly 3% of daily turnover in these shares. Alternatively the maximum order size q_z could be chosen for example according to the average depth of the order book or a decile of the distribution of market order sizes. We calculate the hypothetical price impacts for $z = 10$ equidistant order sizes. For our choice of $q(z)$, the order book has sufficient depth at each point of time.

Dynamics of the parameter

As our aim is to get a prediction of the whole future price impact function, a model for the dynamics of the function is needed. In addition to the linear specification of the price impact function we specify a time series model for the slope parameter λ. To capture the dynamics, we estimate AR(p) processes for the liquidity parameter estimates λ: [20]

$$\lambda_t = c + \sum_{j=1}^{p} \gamma_j \lambda_{t-j} + \varepsilon_t \tag{19}$$

With $c = (1 - \gamma_1 - ... - \gamma_p)\lambda_0$, where λ_0 is the long term mean of the process. $\varepsilon_{j,t}$ is the noise term. We estimate (19) from the 3315 observations by least squares.

The one-step-prediction can be calculated as: [21]

[20] The order of the $AR(p)$-process is selected by the common model selection criteria Akaike Information Criterion and Schwarz Criterion.

[21] Hamilton (1994), p.260. The s-step-prediction could be calculated iteratively, see

$$\widehat{\lambda}_{t+1|t} = \widehat{c} + \widehat{\gamma}_1 \lambda_t + ... + \widehat{\gamma}_p \lambda_{t-p+1} \tag{20}$$

Evaluating the accuracy of the price impact curve forecast

We split the sample period in two parts and estimate the $AR(1)$-model in the first half. Then we use the estimated parameters $\widehat{c}$ and $\widehat{\gamma}$ to make out of sample forecasts in the second period. [22] To assess the forecast accuracy, we calculate the common statistics mean average error (ME), root mean squared error (RMSE) and Theils U (U) for the predicted price impacts at different order sizes:

$$ME_i = \frac{1}{T} \sum_{t=1}^{T} (p_{t,i} - \widehat{p}_{t,i}) \tag{21}$$

$$RMSE_i = \sqrt{\frac{1}{T} \sum_{t=1}^{T} (p_{t,i} - \widehat{p}_{t,i})^2} \tag{22}$$

$$U_i = \sqrt{\frac{\frac{1}{T} \sum_{t=1}^{T} (p_{t,i} - \widehat{p}_{t,i})^2}{\frac{1}{T} \sum_{t=1}^{T} p_{t,i}^2}} \tag{23}$$

Where $\widehat{p}_{t,i}$ is given by: [23]

Hamilton (1994), p.81.

[22]It would be inappropriate to run in sample time series regression of the liquidity parameters and compare the fitted and the actual value of the price impact curve due to the fact that the investor does not know the parameter values of the time series model. We therefore estimate the time series model for the liquidity parameters in the first half of the sample period and use the estimated coefficients to make out of sample predictions in the second half of the sample period. Instead of splitting the sample period in an estimation period and and evaluation period, we could run a rolling regression.

[23]To focus on the accuracy of the forecast of the price impact function, we do not try to forecast $\widetilde{p}_t$ and assume it to be known at date $t-1$.

$$\widehat{p}_{t,i} = \widetilde{p}_t + \widehat{\lambda}_{t|t-1} q_i \tag{24}$$

5.1.2 Results

To get an impression of the suitability of the linear model we calculate the average price impact across the 65 trading days for different order sizes. From this average price impact function we estimate a slope parameter λ.

	$\widehat{\lambda}$	R^2	Std. error of est.
DaimlerChrysler	$1.0740 \cdot 10^{-5}$	0.959	0.074
SAP	$6.3759 \cdot 10^{-4}$	0.947	0.506
Deutsche Telekom	$6.1623 \cdot 10^{-6}$	0.991	0.018

Table 1: Liquidity of the aggregated price impact curve

Table 1 shows that the linear price impact function fits the aggregated data quite well with the adjusted R^2 ranging from 0.959 to 0.991. As we will see one is ill-advised to take the average price impact parameter as an estimate for the price impact parameter at a point of time. There is a lot of time variation in liquidity measured by the price impact parameter λ. Table 2 presents some descriptive statistics of the variation of the price impact parameter over the 3315 ten minute intervals.

The time variation in liquidity is substantial. The ratio of the maximum

	Mean $\widehat{\lambda}$	Median $\widehat{\lambda}$	Minimum $\widehat{\lambda}$	Maximum $\widehat{\lambda}$	$std(\widehat{\lambda})$
DaimlerChrysler	$1.0740 \cdot 10^{-5}$	$8.361 \cdot 10^{-6}$	$1.3182 \cdot 10^{-6}$	$8.7816 \cdot 10^{-5}$	$0.8174 \cdot 10^{-5}$
SAP	$6.3759 \cdot 10^{-4}$	$4.4506 \cdot 10^{-4}$	$3.1338 \cdot 10^{-5}$	$8.49 \cdot 10^{-3}$	$7.2372 \cdot 10^{-4}$
Deutsche Telekom	$6.1623 \cdot 10^{-6}$	$5.2013 \cdot 10^{-6}$	$5.6360 \cdot 10^{-7}$	$4.7275 \cdot 10^{-5}$	$3.9750 \cdot 10^{-6}$

Table 2: Variability of the price impact parameter estimate $\widehat{\lambda}$

	$\widehat{c}$	$\widehat{\gamma}$	R^2
DaimlerChrysler	$1.6381 \cdot 10^{-6}$	0.8208	0.7474
SAP	$1.1639 \cdot 10^{-4}$	0.7854	0.6976
Deutsche Telekom	$1.3403 \cdot 10^{-6}$	0.7557	0.6534

Table 3: Time series model for the price impact parameter estimates $\widehat{\lambda}$, AR(1) specification.

estimate and the minimum estimate is 67 for DaimlerChrysler, 271 for SAP and 84 for Deutsche Telekom. The standard deviation relative to the mean estimate is 76 % for DaimlerChrysler, 113 % for SAP and 65 % for Deutsche Telekom, emphasizing the importance of the time variation in liquidity.

For the investor it is important, if future liquidity measured by λ is purely random or if the changes in liquidity are predictable and how accurate the predictions are. Table 3 presents the parameter estimates of the time series model (19) for the three different stocks. [24] The γ estimated over the whole sample period lies in the range between 0.75 and 0.82. In each case γ is significantly different from zero and significantly different from one. The slope coefficient is therefore strongly autocorrelated. In addition the fact that γ differs from one implies that changes in liquidity are predictable. This predictability gives rise to the possibility of minimizing price impact costs by timing their market orders. If λ is above the long term mean, the investor will expect liquidity to increase (λ to decrease) and will postpone some of his trades.

To give an interpretation of the estimates γ we calculate halflifes, i.e. the time it takes after liquidity shocks have vanished by half. The halflife can be calculated as follows:[25]

[24] In each case the AR(1)-specification performs best according to the information criteria

[25] For a derivation see the Appendix.

	Std. dev. of prediction error	R^2 out of sample
DaimlerChrysler	$0.3380 \cdot 10^{-5}$	0.6897
SAP	$2.4441 \cdot 10^{-4}$	0.4716
Deutsche Telekom	$2.7683 \cdot 10^{-6}$	0.6249

Table 4: Forecast accuracy of the AR(1) model

$$\widehat{T^{*}-t} = \frac{\ln(0,5)}{\ln(\widehat{\gamma})} \tag{25}$$

The estimated parameter values imply halflifes of liquidity shocks from 25 to 35 minutes.

The investor is concerned how well the predicted parameters fit the realized parameter estimates. The high in sample R^2 in table 3 already indicates that the AR-model can explain a significant part of the variation of the slope parameter λ. To assess the accuracy of the $AR(1)$-model, we evaluate out of sample forecasts. We split the sample period in two parts and estimate the $AR(1)$-model in the first half. Then we use the estimated parameters $\widehat{c}$ and $\widehat{\gamma}$ to make out of sample forecasts in the second period. [26] As one can see from Table 4, the forecast accuracy of the $AR(1)$-model is high. Between 47% and 69% of the variation of the liquidity parameter λ can be explained by the time series models. The unexplained variation of the liquidity parameter indicated by the standard deviation of the prediction error in table 4 can be explained as a measure of liquidity risk.

So far we only evaluated the forecast accuracy concerning the liquidity parameter $\widehat{\lambda}$. But the aim is to get a good forecast for the whole price impact function. In fact we have a joint hypothesis about the functional form and the dynamics of the parameters of this functional form. An adequate test of fore-

[26] The parameter estimates of the time series model for the first 35 trading days can be found in table 15.

	ME	std	RMSE	Theils U
10000	0.0014	0.0839	0.0840	0.6736
20000	-0.0294	0.1496	0.1524	0.7275
30000	-0.0584	0.2076	0.2157	0.7194
40000	-0.0809	0.2513	0.2640	0.6598
50000	-0.0976	0.2934	0.3093	0.6006
60000	-0.0965	0.3332	0.3469	0.5054
70000	-0.0746	0.3777	0.3850	0.4295
80000	-0.0282	0.4337	0.4346	0.3843
90000	0.0634	0.5464	0.5501	0.3819
100000	0.1801	0.7572	0.7784	0.4335

Table 5: Accuracy of the linear price impact curve forecast, Daimler Chrysler, estimation period

cast accuracy lies in the comparison of the predicted price impact function and the real shape of the order book.

Table 5 and Table 6 evaluate the accuracy of the $AR(1)$-based price impact function forecast for Daimler Chrysler.[27] As one can see, the forecast accuracy measured by the Theils U is high. In general for larger order sizes, the forecast accuracy increases monotonically (except for the largest order size). The mean error is positive for small and for large order sizes, indicating some form of non-linearity. Therefore it seems promising to estimate a more flexible function than the linear model. In Figure (2) (see the Appendix) we plot the average price impact function for Daimler Chrysler. The plot indicates that the functional form of the average price impacts is slightly convex.

[27] The results for SAP and Deutsche Telekom can be found in the Appendix in tables 16 to 19.

	ME	std	RMSE	Theils U
10000	0.0088	0.0690	0.0696	0.5595
20000	-0.0148	0.1115	0.1125	0.5390
30000	-0.0356	0.1494	0.1536	0.5144
40000	-0.0507	0.1801	0.1871	0.4694
50000	-0.0569	0.2073	0.2150	0.4196
60000	-0.0553	0.2316	0.2381	0.3770
70000	-0.0450	0.2405	0.2447	0.3207
80000	-0.0192	0.2779	0.2786	0.2993
90000	0.0231	0.3745	0.3752	0.3341
100000	0.0862	0.4736	0.4813	0.3585

Table 6: Accuracy of the linear price impact curve forecast, Daimler Chrysler, evaluation period

5.2 The power function model

5.2.1 Methodology

In this section we fit a power function according to (1) to the price impact function which is more flexible in representing different shapes of the price impact function. This function has the property that it increases monotonically with order size (for $\lambda_1, \lambda_2 > 0$), concave, convex as well as linear forms are possible. $\lambda_2 = 1$ represents a linear function, $\lambda_2 < 1$ a concave function and $\lambda_2 > 1$ a convex function. As the order book looks different at different points of time, we allow for time varying liquidity parameters λ_1 and λ_2.

$$p_{t,i} = \widetilde{p}_t + \lambda_{1,t} q_i^{\lambda_{2,t}} + v_{t,i} \tag{26}$$

We estimate (26) by nonlinear least squares. The estimation is done in GAUSS employing the Levenberg-Marquardt variation of the Gauss-Newton method.

In the case of the power function the dynamics of the two parameters of the function which fits the order book have to be modeled. The dynamics of the parameters could be modeled univariatly. For each parameter an AR(p) process could be estimated and employed for prediction. However, this approach neglects joint dynamics of the parameters. To capture the joint dynamics of the parameters λ_1 and λ_2, we specify the following VAR(1)-system:

$$\begin{pmatrix} \log(\lambda_{1,t}) \\ \lambda_{2,t} \end{pmatrix} = \begin{pmatrix} c_1 \\ c_2 \end{pmatrix} + \begin{bmatrix} \gamma_{1,1} & \gamma_{1,2} \\ \gamma_{2,1} & \gamma_{2,2} \end{bmatrix} \begin{pmatrix} \log(\lambda_{1,t-1}) \\ \lambda_{2,t-1} \end{pmatrix} + \begin{pmatrix} \varepsilon_{1,t} \\ \varepsilon_{2,t} \end{pmatrix} \quad (27)$$

We transform the variable λ_1 using a logarithm (with base 10). This specification of the VAR is justified by the fact that the relationship between $\log \lambda_{1,t}$ and $\lambda_{2,t}$ is approximately linear, as can be seen from Figure 3 (see the Appendix) where (contemporaneous) values of $\log \lambda_1$ and λ_2 are plotted. A model with untransformed variables has very poor predictive ability.

In short notation (27) is given by:

$$\lambda_t = c + \sum_{j=1}^{p} \Gamma_j \lambda_{t-j} + \varepsilon_t \quad (28)$$

$c = (I - \Gamma_1 - ... - \Gamma_p)\lambda_0$, where λ_0 is the vector of long term means. The noise terms $\varepsilon_{j,t}$ are assumed to be only contemporary correlated. Equation (27) can be estimated for each single equation by least squares.[28]

The one-step-prediction can be calculated as: [29]

[28] See for example Hamilton (1994), p. 294.

[29] Note that we predict $\log(\lambda_{1,t})$, so we have to retransform it to calculate the predicted price impact function.

$$\widehat{\lambda}_{t+1|t} = \lambda_0 + \widehat{\Gamma}_1(\lambda_t - \lambda_0) + ... + \widehat{\Gamma}_p(\lambda_{t-p+1} - \lambda_0) \qquad (29)$$

5.2.2 Results

As in the case of the linear model, we first estimate the parameters of the price impact function from the aggregated price impact curve. Table 7 presents the results for the parameter estimates $\widehat{\lambda}_1$ and $\widehat{\lambda}_2$. For Daimler Chrysler and SAP, the parameter estimate $\widehat{\lambda}_2$ is significantly greater than one. The adjusted R^2 is higher compared to Table 1, indicating a fit superior to the linear model.

	$\widehat{\lambda}_1$	$\widehat{\lambda}_2$	$\overline{R}^2$	Std. error of est.
DaimlerChrysler	$3.056 \cdot 10^{-7}$	1.3166 (***)	0.988	0.041
SAP	$2.684 \cdot 10^{-5}$	1.3542 (***)	0.979	0.316
Deutsche Telekom	$3.037 \cdot 10^{-6}$	1.0630	0.992	0.017

Table 7: Liquidity of the aggregated price impact curve, power function

Again it has to be considered, that the price impact function is not constant over time, and it is not convex at each point of time. We estimate the price impact function for the 3315 ten minute snapshots. Table 8 presents descriptive statistics on the estimated exponent of the power function.

In each case the median lies around one, which means that in about half of the periods the price impact is convex, in the other half concave. The high

	mean $\widehat{\lambda}_2$	median $\widehat{\lambda}_2$	Min. $\widehat{\lambda}_2$	Max. $\widehat{\lambda}_2$	$std(\widehat{\lambda}_2)$
DaimlerChrysler	1.2523	1.0460	0.0803	8.5285	0.7648
SAP	1.1762	0.9319	0.0155	16.5754	0.9500
Deutsche Telekom	1.0533	0.9713	0.0417	5.0856	0.4648

Table 8: Variability of the price impact parameter estimate $\widehat{\lambda}_2$

standard deviation of λ_2 implies that the deviations from linearity are substantial. To get an idea if the variation of the functional form implied by the parameter λ_2 is systematic or purely random, Table 9 gives the unconditional probability of a convex relationship as well as the conditional probability of observing a convex relationship given that the last price impact function was convex.

	$P(\widehat{\lambda_2} > 1)$	$P(\widehat{\lambda_2} > 1 \mid \widehat{\lambda}_{2,-1} > 1)$
DaimlerChrysler	0.550	0.788
SAP	0.451	0.727
Deutsche Telekom	0.466	0.712

Table 9: Probability of a convex price impact function

In each case this conditional probability is more than 70 percent which strongly deviates from the unconditional probability, indicating predictability for the shape parameter.

Table 10 presents the results of the VAR(1)-model. In each case all parameters for the lagged variables $\log(\lambda_1)$ and λ_2 are significant. [30]

Again we split the sample in two parts to assess the accuracy of the VAR-model to forecast future price impact parameters. The out of sample R^2 of the time series model for second period is presented in Table 11.[31] The out of sample R^2 is in each case higher than for the univariate time series model (we do not report the results for the univariate models here).

Since the forecast accuracy indicated by the out-of-sample R^2 of the VAR-model for the two liquidity parameters cannot be compared directly to the forecast accuracy of the linear model, we again compare the whole predicted

[30] Stability of the system requires that the roots of the polynomial $(1-\gamma_{11}L)(1-\gamma_{22}L) - (\gamma_{12}\gamma_{21}L^2)$ lie outside the unit circle. See e.g. Enders S. 297ff. As can be checked the stability condition is met for each stock.

[31] Table 21 in the Appendix reports the parameter estimates for the first 35 trading days.

		$\widehat{c}$	$\widehat{\gamma}_{i,1}$	$\widehat{\gamma}_{i,2}$	R^2
Daimler Chrysler	$\log\lambda_1$	-2.8476	-1.9549	-12.5400	0.4730
	λ_2	0.5703	0.5675	3.3882	0.4974
SAP	$\log\lambda_1$	-0.6843	-1.1337	-6.7207	0.3972
	λ_2	0.1962	0.4956	2.5272	0.4329
Deutsche Telekom	$\log\lambda_1$	-3.4590	-1.0140	-7.3060	0.2917
	λ_2	0.6800	0.3679	2.2868	0.3360

Table 10: Time series model for the price impact parameter estimates λ_1 and λ_2: VAR(1)

		R^2 out of sample
DaimlerChrysler	$\log\lambda_1$	0.4026
	λ_2	0.4311
SAP	$\log\lambda_1$	0.1991
	λ_2	0.2191
Deutsche Telekom	$\log\lambda_1$	0.2547
	λ_2	0.3008

Table 11: Forecast accuracy of the VAR(1) model

price impact function and the real shape of the order book. Tables 12 and 13 present the results for Daimler Chrysler. [32] When compared to tables 5 and 6, the power function model is superior when evaluated by the measures Theils U, root mean squared error and standard deviation of prediction error for most of the order sizes. This holds also true for the mean absolute error. In general the estimate of the price impact function in the power function case seems to be biased downwards as the mean errors are always positive.

[32]The results for SAP and Deutsche Telekom can be found in the Appendix in tables 20 to 24.

	ME	std	RMSE	Theils U
10000	0.0481	0.0646	0.0805	0.6461
20000	0.0557	0.0941	0.1093	0.5218
30000	0.0544	0.1210	0.1326	0.4425
40000	0.0482	0.1446	0.1525	0.3810
50000	0.0354	0.1805	0.1839	0.3572
60000	0.0267	0.2669	0.2682	0.3907
70000	0.0236	0.3625	0.3632	0.4053
80000	0.0277	0.4421	0.4429	0.3917
90000	0.0577	0.5281	0.5312	0.3688
100000	0.0907	0.6444	0.6507	0.3625

Table 12: Accuracy of the power function price impact curve forecast, Daimler Chrysler, estimation period

	ME	std	RMSE	Theils U
10000	0.0470	0.0606	0.0767	0.6169
20000	0.0509	0.0870	0.1008	0.4831
30000	0.0477	0.1112	0.1210	0.4053
40000	0.0411	0.1410	0.1468	0.3683
50000	0.0345	0.1749	0.1783	0.3479
60000	0.0270	0.2044	0.2062	0.3264
70000	0.0195	0.2243	0.2252	0.2952
80000	0.0189	0.2767	0.2774	0.2980
90000	0.0260	0.3708	0.3718	0.3310
100000	0.0450	0.4522	0.4544	0.3384

Table 13: Accuracy of the power function price impact curve forecast, Daimler Chrysler, evaluation period

5.3 Testing for the difference in forecast accuracy of the competing models

The evaluation of the forecast accuracy of the linear and the nonlinear model in the former sections indicated a superiority of the nonlinear model. However, the forecast accuracy measures in section 5.1 and 5.2 were defined for various order sizes and indicated no strict dominance of either model. Therefore we propose a measure which is independent of order size and employ a statistical test to directly compare the predictive ability of the linear model to the nonlinear model.

To test for difference in forecast accuracy we employ the Diebold and Mariano (2002) test. [33] The Diebold and Mariano (2002) t- statistics for the one step ahead forecast is calculated as:

$$DM = \sqrt{T-1}\frac{\overline{d}}{\widehat{std}(\overline{d})} \tag{30}$$

$$\overline{d} = \frac{1}{T-R}\sum_{t=R+1}^{T} d_t \tag{31}$$

$$d_t = L(\widehat{\nu}_{1,t}) - L(\widehat{\nu}_{2,t}) \tag{32}$$

R is the number of observations used to estimate the parameters. $\widehat{\nu}_{1,t}$ is the vector of price impact forecast errors for model 1 for different order sizes observed at time t associated with forecasts from time $t-1$. We specify the loss function as the sum of squared forecast errors:[34]

[33] This test is used for example in Goyal and Welch (2004).

[34] The specification of the loss function is somewhat heuristic since the economic loss associated with a forecast error of a particular size and sign need not have this specific form.

Daimler Chrysler	linear, AR	linear,n.f.	linear, h.a.
power function, VAR(1)	-4.0377 (***)	-4.8721 (***)	-8.3026 (***)
linear, AR	-	-4.5523 (***)	-6.8634 (***)
linear, n.f.	-	-	-6.2380 (***)
SAP	linear, AR	linear,n.f.	linear, h.a.
power function, VAR(1)	-1.3851	-2.4118 (***)	-11.1958 (***)
linear, AR	-	-2.9097 (***)	-15.0647 (***)
linear, n.f.	-	-	-12.9956 (***)
Deutsche Telekom	linear, AR	linear,n.f.	linear, h.a.
power function, VAR(1)	-2.6525 (***)	-3.2786 (***)	-6.6936 (***)
linear, AR	-	-2.7663 (***)	-6.2932 (***)
linear, n.f.	-	-	-5.4183 (***)

Table 14: Out-of-sample forecast accuracy comparisons

$$L(\widehat{\nu}_{1,t}) = \sum_{i=1}^{z} \widehat{\nu}_{1,t,i}^2 \tag{33}$$

For example, in the case of the linear model the forecast error $\widehat{\nu}_{1,t,i}$ is calculated as:

$$\widehat{\nu}_{1,t,i} = \sum_{i=1}^{z} (p_{t,i} - \widetilde{p}_t - \widehat{\lambda}_t q_i) \tag{34}$$

where the forecast $\widehat{\lambda}_t$ is either the naive forecast λ_{t-1}, the historical average $\overline{\lambda}$ or is calculated from the time series model for λ according to (29).

The DM-test allows for forecast errors which are non gaussian, have nonzero mean, are serially and contemporaneously correlated. In the estimation of the standard deviation we make a correction for serial correlation according to Newey and West (1987).

The null hypothesis is that the two forecasts have the same predictive ability with respect to the loss function L. Under the null hypothesis of equal forecast accuracy, the DM statistic is asymptotically normally distributed when testing non nested models. In our case the models are nested so the test statistic is not normally distributed. Therefore we use the corrected critical values from Table 3 in McCracken (1999).

In table 14 for each stock the forecast accuracy of the two models is compared. The results show that in each case the power function model shows the best predictive power relative to the linear model. This difference in forecast accuracy is significant except for the case of SAP. In the class of linear models we also consider alternative models of using the average liquidity level as predictor of future liquidity and the naive prediction of using the last observation of the liquidity parameter as our estimate.[35] As a result the AR(1) specification is superior to the naive prediction. In each case the mean specification is clearly beaten. This shows that one is ill- advised to take the average historical liquidity as predictor for future liquidity.

6 Conclusion

In this paper we investigated the form and temporal dynamics of the price impact function.

To motivate our empirical analysis we demonstrated theoretically the impact of different forms of the price impact function on optimal order splitting strategies. The optimal strategy looks completely different if the price impact function is concave or convex.

To study empirically how the functional form and the temporal dynamics of

[35] Underlying this approach is the idea that the current order book constitues the best predictor for the future order book. How well this model performs depends apparently on the length of the horizon of the prediction and the extent of unsystematic changes of the order book.

the price impact function look like, we employed order book data from the XETRA automated exchange. These data are especially well suited for our analysis as they reflect the relevant trading opportunities of investors and allow to analyze temporal variations of the whole price impact function. In our empirical analysis we documented that the functional form as well as the level of liquidity exhibits considerable variation over time. The functional form is in about 50% of the cases concave and in about 50% convex.

Most importantly, the strong variation in the price impact function is not purely random but we find strong predictability in liquidity. This holds true for the level as well as the shape of the price impact function. When evaluating the out of sample forecast accuracy of different models, a linear price impact function with the parameter modeled as an AR(p) process works well in predicting the future price impact function. However, a power function with the dynamics of the parameters modeled by a VAR(1)-process is shown to be significantly more accurate in predicting future price impact functions.

Our study leaves open several important questions for further research. From a theoretical point of view it is of interest to analyze the consequences of the complex dynamics of the price impact function explored in this paper for optimal trading strategies of large investors in more detail than in the stylized examples in this paper. Especially the incorporation of stochastically time varying price impact functions has not been addressed in the literature so far. This would allow to assess the possible gains from exploiting the predictability of liquidity by means of timing market orders.

Concerning the empirical analysis we predicted the order book solely based on its own history. A more complete analysis would consider the joint dynamics of liquidity supply (represented by the order book) and liquidity demand (represented by market orders). If incorporating this common dynamics into the model adds significantly to the forecast ability is left for further research.

References

Admati, Anat R., and Paul Pfleiderer, 1988, A theory of intraday patterns : Volume and price variability, *review of financial studies* 1, 3–40.

Algert, Peter M., 1990, Estimates of non-linearity in the response of stock prices to order imbalances, Working Paper.

Almgren, Robert, and Neil Chriss, 1998, Optimal liquidation, Working Paper.

Amihud, Yakov, 2002, Illiquidity and stock returns: cross-section and time-series effects, *Journal of Financial Markets* 5, 31–56.

Beltran-Lopez, Helena, Pierre Giot, and Joachim Grammig, 2002, Expected and unexpected cost of trading in the XETRA automated auction market, Working Paper.

———, 2003, Liquidity, volatility and trading activity in the xetra automated auction market, Working Paper.

Bertsimas, D., and A. Lo, 1998, Optimal control of execution costs, *Journal of financial markets* 1, 1–50.

Biais, B., P. Hillion, and C. Spatt, 1995, An empirical analysis of the order flow and order book in the paris bourse, *Journal of Finance* 50, 1655–1689.

Breen, William J., Laurie Simon Hodrick, and Robert A. Korajczyk, 2002, Predicting equity liquidity, *Management Science* 48, 470–483.

Chan, Louis K., and Josef Lakonishok, 1995, The behavoir of stock prices around institutional trades, *Journal of Finance* 50, 1147–1174.

Chen, Zhiwu, Werner Stanzl, and Masahiro Watanabe, 2002, Price impact costs and the limit of arbitrage, Working Paper.

Coppejans, Mark, Ian Domowitz, and Ananth Madhavan, 2003, Resiliency in an automated auction, Working Paper.

Diebold, Francis X., and Canlin Li, 2003, Forecasting the term structure of government bond yields, Working Paper.

Diebold, Francis X., and Roberto S. Mariano, 2002, Comparing predictive accuracy, *Journal of business and economic statistics* 20, 134–144.

Engle, Robert F., and Joe Lange, 2001, Predicting VNET: A model of the dynamics of market depth, *Journal of Financial Markets* 4, 113–142.

Glosten, L.R., 1994, Is the electronic open limit order book inevitable?, *Journal of Finance* 49, 1127–1161.

Gomber, Peter, Uwe Schweickert, and Erik Theissen, 2004, Zooming in on liquidity, Working Paper.

Goyal, Amit, and Ivo Welch, 2004, A comprehensive look at the empirical performance of equity premium prediction, Working Paper.

Hamilton, James D., 1994, *Time series analysis* (Princeton University Press: New Jersey).

Hasbrouck, Joel, 1991, Measuring the information content of stock trades, *Journal of Finance* 46, 179–207.

Hausman, Jerry A., Andrew W. Lo, and A. Craig MacKinlay, 1992, An ordered probit analysis of transaction stock prices, *Journal of Financial Economics* 31, 319–379.

Huberman, Gur, and Dominika Halka, 2001, Systematic liquidity, *Journal of Financial Research* 24, 161–178.

Huberman, G., and W. Stanzl, 2003, Optimal liquidity trading, Working Paper.

——— , 2004, Quasi-arbitrage and price manipulation, *Econometrica* 72, 1247–1275.

Keim, Donald B., and Ananth Madhavan, 1995, Anatomy of the trading process : Empirical evidence on the behavior of institutional traders, *Journal of financial economics* 37, 371–398.

——— , 1996, The upstairs market for large -block transactions: Analysis and measurement of price effects, *Review of financial studies* 9, 1–36.

Kempf, Alexander, and Olaf Korn, 1999, Market depth and order size, *Journal of Financial Markets* 2, 29–48.

Knez, Peter J., and Mark J. Ready, 1996, Estimating the profits from trading strategies, *The review of financial studies* 9, 1121–1163.

Kumar, Santosh, 2003, Liquidity: Persistence and forecast in open electronic limit order book market, Working Paper.

Madhavan, Ananth, and Seymour Smidt, 1991, A bayesian model of intraday specialist pricing, *Journal of financial economics* 30, 99–134.

Maslov, Sergei, and Mark Mills, 2001, Price fluctuations from the order book perspective - empirical facts and a simple model, *Physica A* 299, 234–246.

McCracken, Michael W., 1999, Asymtotics for out of sample tests of causality, Working Paper.

Nelson, C.R., and A.F. Siegel, 1987, Parsimonious modeling of yield curves, *Journal of Business* 60, 473–489.

Newey, Whitney K., and Kenneth D. West, 1987, A simple positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix, *Econometrica* 55, 703–708.

Sadka, Ronnie, 2003, Momentum, liquidity risk, and limits to arbitrage, Working Paper.

Spierdijk, Laura, Theo E. Nijman, and Arthur H.O. Van Soest, 2004, Price dynamics and trading volume: A semiparametric approach, Working Paper.

Weber, Philipp, and Bernd Rosenow, 2003, Order book approach to price impact, Working Paper.

7 Appendix

Price Impact with general exponent

The price impact function is given by a power function with general exponent λ_2:

$$p_t = \widetilde{p}_t + \lambda_1 q_t^{\lambda_2} \tag{35}$$

The price update function is given by:

$$\widetilde{p}_t = \alpha \widetilde{p}_{t-1} + (1-\alpha) p_{t-1} + \varepsilon_t \tag{36}$$

In $t = 2$ the trading volume is determined by the trading restriction:

$$q_2 = Q_2 = Q - q_1 \tag{37}$$

The value function in $t = 2$ is given by:

$$V_2(\widetilde{p}_2, Q_2) = \int_0^{Q_2} (\widetilde{p}_2 + \lambda_1 x^{\lambda_2}) dx = (\widetilde{p}_2 + \frac{\lambda_1}{\lambda_2 + 1} Q_2^{\lambda_2}) Q_2 \tag{38}$$

The value function in $t = 1$ is given by:

$$V_1(\widetilde{p}_1, Q_1) = \min_{q_1} E_1[\int_0^{q_1} (\widetilde{p}_1 + \lambda_1 x^{\lambda_2}) dx + V_2(p_1, Q_2)] \tag{39}$$

The expected costs are:

$$\begin{aligned} C_1 &= E_1[\int_0^{q_1} (\widetilde{p}_1 + \lambda_1 x^{\lambda_2}) dx + V_2(p_1, Q_2)] \\ &= E_1[(\widetilde{p}_1 + \frac{\lambda_1}{\lambda_2 + 1} q_1^{\lambda_2}) q_1 + (\alpha \widetilde{p}_1 + (1-\alpha) p_1 + \varepsilon_2 + \frac{\lambda_1}{\lambda_2 + 1} (Q - q_1)^{\lambda_2})(Q - q_1)] \end{aligned} \tag{40}$$

$$= E_1[(\widetilde{p}_1 + \frac{\lambda_1}{\lambda_2+1} q_1^{\lambda_2})q_1 + (\widetilde{p}_1 + (1-\alpha)\lambda_1 q_1^{\lambda_2} + \varepsilon_2 + \frac{\lambda_1}{\lambda_2+1}(Q-q_1)^{\lambda_2})(Q-q_1)]$$
$$= (\widetilde{p}_1 + \frac{\lambda_1}{\lambda_2+1} q_1^{\lambda_2})q_1 + (\widetilde{p}_1 + (1-\alpha)\lambda_1 q_1^{\lambda_2} + \frac{\lambda_1}{\lambda_2+1}(Q-q_1)^{\lambda_2})(Q-q_1)$$

The first order conditions is then given by:

$$\frac{\delta C_1}{\delta q_1} = \lambda_1 q_1^{\lambda_2} + (1-\alpha)\lambda_1 q_1^{\lambda_2-1}(\lambda_2 Q - (\lambda_2+1)q_1) - \lambda_1(Q-q_1)^{\lambda_2} \doteq 0 \quad (41)$$

Rearranging yields:

$$\frac{\delta C_1}{\delta q_1} = (1-(1-\alpha)(1+\lambda_2))\lambda_1 q_1^{\lambda_2} + (1-\alpha)\lambda_2 Q \lambda_1 q_1^{\lambda_2-1} - \lambda_1(Q-q_1)^{\lambda_2} \doteq 0 \quad (42)$$

A closed form solution for q_1 can be obtained only for some special cases. This can be shown if we rewrite (7) as

$$a_1 q_1^{\lambda_2} + a_2 q_1^{\lambda_2-1} + a_3(Q-q_1)^{\lambda_2} = 0 \quad (43)$$

where $\lambda_2 = \frac{m}{n}$, with m and n positive integers. Subtracting the third term, taking the exponent n and multiplying with $q_1^{max[n-m,0]}$ yields a polynomial of order $max(m,n)$. Closed form solutions can only be given for polynomials of order not higher than 4. Accordingly the special cases with closed form solutions are: $\lambda_2 = 1/4$, $\lambda_2 = 1/3$, $\lambda_2 = 1/2$, $\lambda_2 = 2/3$, $\lambda_2 = 3/4$ $\lambda_2 = 1$, $\lambda_2 = 4/3$,$\lambda_2 = 3/2$,$\lambda_2 = 2$, $\lambda_2 = 3$,$\lambda_2 = 4$.

Concave price impact with exponent 1/2

If $\lambda_2 = \frac{1}{2}$, the first order condition becomes:

$$\frac{\delta C_1}{\delta q_1} = \lambda q_1^{1/2} + \frac{1}{2}(1-\alpha)\lambda(q_1^{-1/2}Q - 3q_1^{1/2}) - \lambda(Q-q_1)^{1/2} \doteq 0$$

The minimum of expected costs is associated with:[36]

$$q_1^* = \frac{3 - 4\alpha + 3\alpha^2 + 2\sqrt{1 - 2\alpha + 2\alpha^2}}{5 + 9\alpha^2 - 6\alpha} Q \tag{44}$$

From (40) it follows that the value function V_1 takes the form:

$$V_1(\widetilde{p}_1, Q_1) \;=\; \widetilde{p}_1 Q + (\alpha - \frac{1}{3})\lambda(q_1^*)^{3/2} + \lambda(1-\alpha)(q_1^*)^{1/2}Q + \frac{2}{3}\lambda(Q - q_1^*)^{3/2}$$

Linear price impact

For $\lambda_2 = 1$, the first order condition becomes:

$$\frac{\delta C_1}{\delta q_1} = -\alpha\lambda_1(Q - 2q_1) \doteq 0 \tag{45}$$

The optimal order size is:

$$q_1^* = \frac{1}{2}Q \tag{46}$$

From (40) and (46) it follows that the value function V_1 is given by:

$$V_1 = \widetilde{p}_1 Q + \frac{1}{2}\lambda Q^2 - \frac{\alpha}{4}\lambda Q^2 \tag{47}$$

Convex price impact with exponent 2

[36]The second solution to (44) corresponds to the maximum of expected costs.

For $\lambda_2 = 2$, the first order condition becomes:

$$\frac{\delta C_1}{\delta q_1} = \lambda_1 q_1^2 + (1-\alpha)\lambda_1 q_1 (2Q - 3q_1) - \lambda_1 (Q - q_1)^2 \doteq 0 \tag{48}$$

The optimal order size is given by:[37]

$$q_1^* = \frac{-2 + \alpha + \sqrt{1 - \alpha + \alpha^2}}{3(\alpha - 1)} Q \tag{49}$$

The value function V_1 takes the form:

$$V_1 = \widetilde{p}_1 Q + \frac{1}{3}\lambda q_1^{*3} + (1-\alpha)\lambda q_1^{*2}(Q - q_1^*) + \frac{1}{3}\lambda (Q - q_1^*)^3 \tag{50}$$

Calculating the halflife

We look for T^*, so that

$$\frac{E(\lambda_{T^*}) - \lambda_0}{\lambda_t - \lambda_0} = 0.5 \tag{51}$$

The dynamics of the parameter is given by:

$$\lambda_{t+1} = c + \gamma \lambda_t + \varepsilon_{t+1} \tag{52}$$

The long term mean of the process is given by $\lambda_0 = c/(1-\gamma)$. We rewrite (52):

$$\lambda_{t+1} = \lambda_0 (1-\gamma) + \gamma \lambda_t + \varepsilon_{t+1} = \lambda_0 + \gamma(\lambda_t - \lambda_0) + \varepsilon_{t+1} \tag{53}$$

Taking expectations yields:

[37] The second solution to (44) corresponds to the maximum of expected costs.

$$E(\lambda_{t+1}) = \lambda_0 + \gamma(\lambda_t - \lambda_0) \tag{54}$$

λ_{t+2} can be accordingly rewritten as:

$$\lambda_{t+2} = \lambda_0 + \gamma^2(\lambda_t - \lambda_0) + \gamma\varepsilon_{t+1} + \varepsilon_{t+2} \tag{55}$$

The two step ahead forecast can be calculated as:

$$E(\lambda_{t+2}) = \lambda_0 + \gamma^2(\lambda_t - \lambda_0) \tag{56}$$

If we proceed in this fashion, the $T - t$ step ahead forecast can be calculated as:

$$E(\lambda_T) = \lambda_0 + \gamma^{T-t}(\lambda_t - \lambda_0) \tag{57}$$

The fraction of the deviation from the mean which is expected to vanish in the period $T - t$ is given by:

$$\frac{E(\lambda_T) - \lambda_0}{\lambda_t - \lambda_0} = \gamma^{T-t} \tag{58}$$

From (51) and (58) we get the halflife:

$$T^* - t = \frac{\ln(0,5)}{\ln(\gamma)} \tag{59}$$

As γ has to be estimated, we calculate HL as:

$$\widehat{T^* - t} = \frac{\ln(0,5)}{\ln(\widehat{\gamma})} \tag{60}$$

	$\widehat{c}$	$\widehat{\gamma}$	R^2
DaimlerChrysler	$1.5210 \cdot 10^{-6}$	0.8355	0.7686
SAP	$1.3070 \cdot 10^{-4}$	0.7943	0.7128
Deutsche Telekom	$1.2981 \cdot 10^{-6}$	0.7333	0.6796

Table 15: Time series model for the price impact parameter estimates $\widehat{\lambda}$, AR(1) specification, estimation period.

	ME	std	RMSE	Theils U
10000	0.0230	0.0451	0.0506	0.5656
20000	0.0158	0.0646	0.0665	0.4719
30000	0.0100	0.0797	0.0803	0.4177
40000	0.0034	0.0902	0.0903	0.3707
50000	-0.0020	0.1009	0.1009	0.3407
60000	-0.0070	0.1094	0.1096	0.3124
70000	-0.0093	0.1173	0.1177	0.2880
80000	-0.0056	0.1288	0.1289	0.2736
90000	0.0000	0.1456	0.1456	0.2705
100000	0.0064	0.1648	0.1650	0.2722

Table 16: Accuracy of the linear price impact curve forecast, Deutsche Telekom, estimation period

	ME	std	RMSE	Theils U
10000	0.0178	0.0534	0.0563	0.5848
20000	0.0062	0.0826	0.0829	0.5291
30000	-0.0026	0.1060	0.1061	0.4812
40000	-0.0096	0.1246	0.1250	0.4358
50000	-0.0107	0.1415	0.1419	0.3917
60000	-0.0091	0.1598	0.1601	0.3622
70000	0.0003	0.1896	0.1896	0.3532
80000	0.0196	0.2224	0.2233	0.3484
90000	0.0503	0.2771	0.2816	0.3676
100000	0.0991	0.3670	0.3802	0.4108

Table 17: Accuracy of the linear price impact curve forecast, Deutsche Telekom, evaluation period

	ME	std	RMSE	Theils U
1000	0.0601	0.7549	0.7573	0.7557
2000	-0.2127	1.3382	1.3550	0.8629
3000	-0.4836	1.8903	1.9512	0.9011
4000	-0.7067	2.2910	2.3976	0.8346
5000	-0.8146	2.6522	2.7745	0.6945
6000	-0.7983	3.0777	3.1796	0.6037
7000	-0.6488	3.5821	3.6404	0.5333
8000	-0.1800	4.2550	4.2588	0.4611
9000	0.6011	5.8372	5.8681	0.4790
10000	1.4065	6.9944	7.1345	0.4842

Table 18: Accuracy of the linear price impact curve forecast, SAP, estimation period

	ME	std	RMSE	Theils U
1000	0.1868	0.4439	0.4816	0.5875
2000	0.0518	0.6620	0.6640	0.5415
3000	-0.1076	0.8507	0.8574	0.5341
4000	-0.2640	1.0513	1.0840	0.5412
5000	-0.3845	1.2639	1.3211	0.5404
6000	-0.4720	1.4134	1.4902	0.5047
7000	-0.4726	1.6169	1.6846	0.4563
8000	-0.4213	2.0825	2.1247	0.4717
9000	-0.3469	2.5080	2.5319	0.4668
10000	-0.0700	3.8204	3.8211	0.5419

Table 19: Accuracy of the linear price impact curve forecast, SAP, evaluation period

	ME	std	RMSE	Theils U
10000	0.0240	0.0458	0.0517	0.5779
20000	0.0212	0.0626	0.0661	0.4688
30000	0.0193	0.0750	0.0775	0.4030
40000	0.0160	0.0854	0.0869	0.3571
50000	0.0134	0.0964	0.0974	0.3287
60000	0.0105	0.1074	0.1079	0.3075
70000	0.0096	0.1165	0.1169	0.2862
80000	0.0141	0.1251	0.1259	0.2672
90000	0.0200	0.1372	0.1386	0.2575
100000	0.0260	0.1504	0.1526	0.2519

Table 20: Accuracy of the power function price impact curve forecast, Deutsche Telekom, estimation period

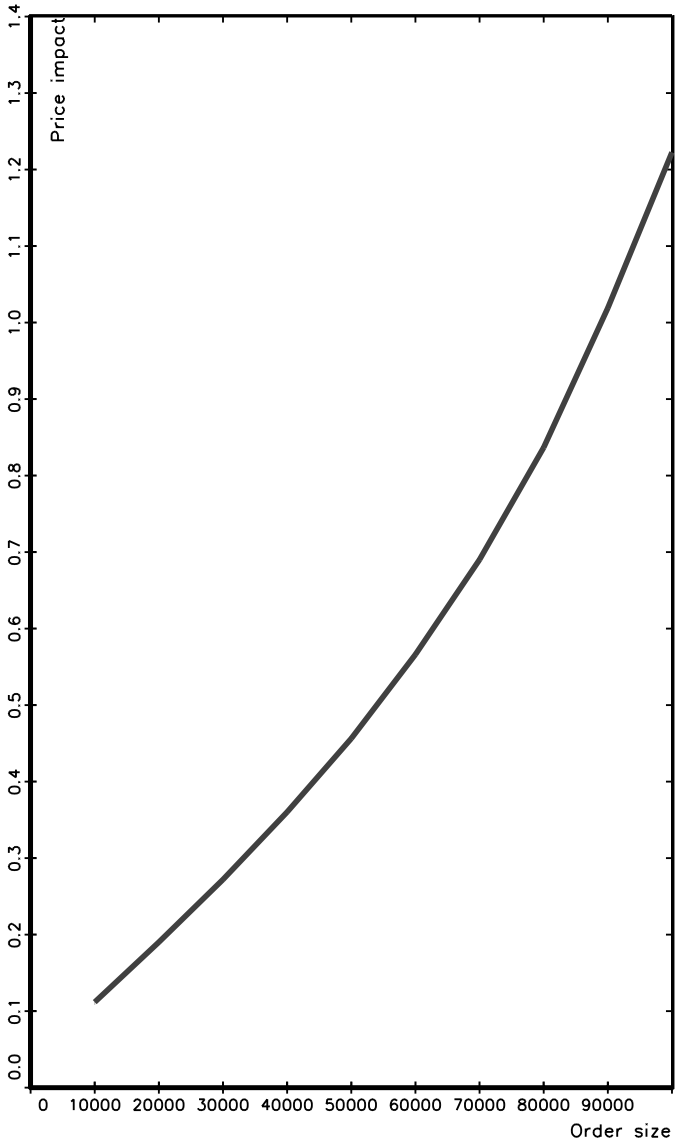

Figure 2: Average Price impact function for DaimlerChrysler

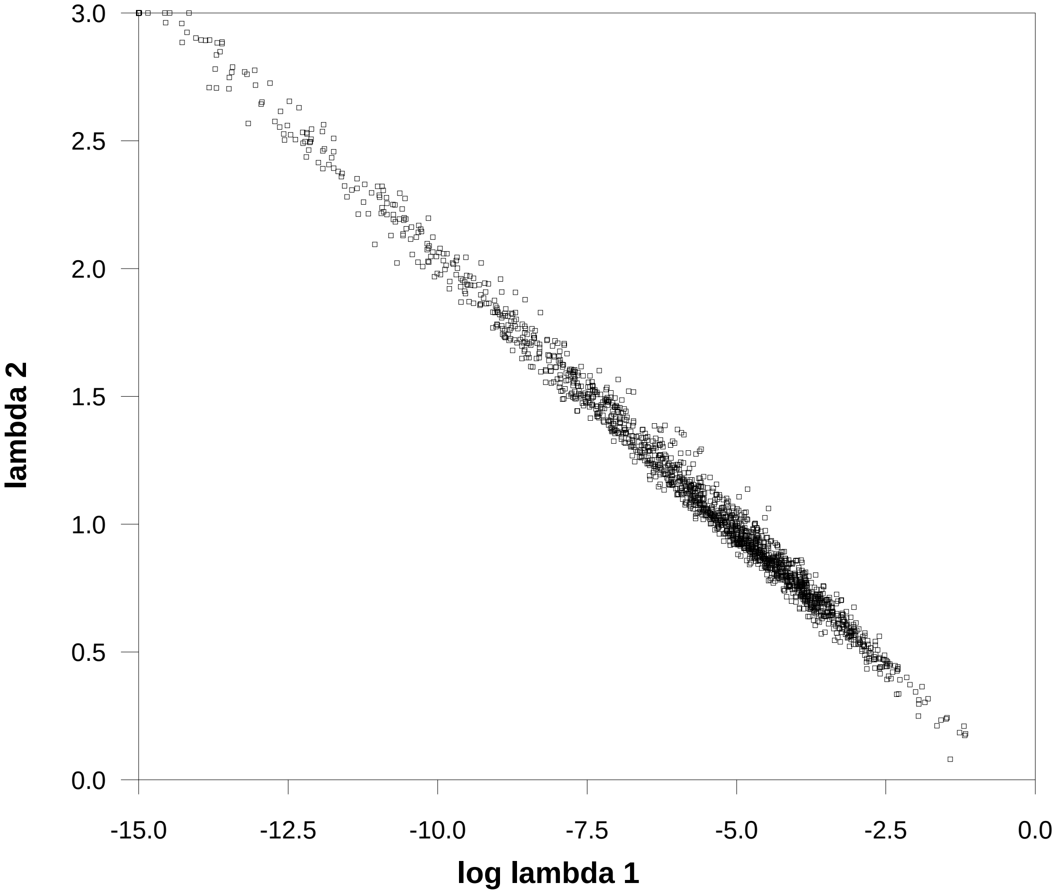

Figure 3: Daimler Chrysler

		$\widehat{c}$	$\widehat{\gamma}_{i,1}$	$\widehat{\gamma}_{i,2}$	R^2
DaimlerChrysler	$\log \lambda_1$	-3.0030	-2.3950	-14.7135	0.4935
	λ_2	0.6012	0.6576	3.8336	0.5172
SAP	$\log \lambda_1$	-0.3113	-1.3396	-7.7606	0.4653
	λ_2	0.0983	0.5514	2.8076	0.5057
Deutsche Telekom	$\log \lambda_1$	-3.0275	-0.6470	-5.7026	0.3016
	λ_2	0.5894	0.2921	1.9559	0.3416

Table 21: Time series model for the price impact parameter estimates λ_1 and λ_2: VAR(1), estimation period

	ME	std	RMSE	Theils U
10000	0.0314	0.0495	0.0586	0.6091
20000	0.0301	0.0737	0.0796	0.5085
30000	0.0277	0.0939	0.0979	0.4441
40000	0.0237	0.1113	0.1138	0.3968
50000	0.0225	0.1312	0.1331	0.3675
60000	0.0211	0.1529	0.1543	0.3491
70000	0.0247	0.1878	0.1894	0.3529
80000	0.0355	0.2212	0.2241	0.3497
90000	0.0552	0.2755	0.2810	0.3668
100000	0.0906	0.3544	0.3659	0.3953

Table 22: Accuracy of the power function price impact curve forecast, Deutsche Telekom, evaluation period

	ME	std	RMSE	Theils U
1000	0.4013	0.6215	0.7399	0.7383
2000	0.4590	0.8500	0.9661	0.6152
3000	0.4478	1.0615	1.1521	0.5321
4000	0.3993	1.3131	1.3725	0.4778
5000	0.3673	1.9461	1.9804	0.4957
6000	0.3459	2.5493	2.5727	0.4885
7000	0.3280	3.3075	3.3237	0.4869
8000	0.4803	4.3557	4.3821	0.4744
9000	0.7744	5.6003	5.6536	0.4615
10000	0.8968	5.8767	5.9448	0.4035

Table 23: Accuracy of the power function price impact curve forecast, SAP, estimation period

	ME	std	RMSE	Theils U
1000	0.2666	0.4580	0.5300	0.6466
2000	0.2601	0.5838	0.6392	0.5213
3000	0.2132	0.6627	0.6962	0.4336
4000	0.1468	0.7589	0.7730	0.3859
5000	0.0896	0.8938	0.8982	0.3674
6000	0.0346	1.0723	1.0729	0.3634
7000	0.0305	1.4541	1.4544	0.3939
8000	0.0362	2.0183	2.0187	0.4482
9000	0.0153	2.5969	2.5969	0.4788
10000	0.1378	3.9256	3.9280	0.5571

Table 24: Accuracy of the power function price impact curve forecast, SAP, evaluation period