
Department of Applied Mathematics
Faculty of EEMCS



University of Twente
The Netherlands

P.O. Box 217
7500 AE Enschede
The Netherlands

Phone: +31-53-4893400

Fax: +31-53-4893114

Email: memo@math.utwente.nl
www.math.utwente.nl/publications

Memorandum No. 1726

**Price dynamics and trading volume:
A semiparametric approach**

L. SPIERDIJK, T.E. NIJMAN¹

A.H.O. VAN SOEST²

June, 2004

ISSN 0169-2690

¹Dept. of Finance and CentER, Tilburg University, P.O. Box 90153, 5000 LE Tilburg

²RAND Corporation (USA) and Dept. of Econometrics & OR and CentER, Tilburg University, P.O. Box 90153, 5000 LE Tilburg

Price Dynamics and Trading Volume: A Semiparametric Approach

Laura Spierdijk*
Theo E. Nijman†
Arthur H.O. van Soest‡

1st June 2004

*Part of this paper has been written when Laura Spierdijk was affiliated to Tilburg University, the Netherlands.

†Theo Nijman is affiliated to the Department of Finance and CentER, Tilburg University, P.O. Box 90153, 5000 LE Tilburg, The Netherlands.

‡Arthur van Soest is affiliated to RAND Corporation (USA) and the Department of Econometrics & OR and CentER, Tilburg University, P.O. Box 90153, 5000 LE Tilburg, The Netherlands.

Abstract

In this paper we investigate the relation between price impact and trading volume for a sample of stocks listed on the New York Stock Exchange. The parametric VAR-models that have been used in the literature impose strong proportionality and symmetry restrictions on the price impact of trades, although market microstructure theory provides many reasons why these restrictions would not hold. We analyze a more flexible semiparametric partially linear specification and establish significant evidence for a nonlinear, asymmetric, increasing, and concave relation between trading volume and both immediate and persistent price impact. Moreover, we compare the price-impact functions obtained in the partially linear model to the ones generated by the parametric models and show that there are considerable differences. We test the parametric specifications against the partially linear model and show that the parametric models are rejected in favor of the semiparametric model. We also test the partially linear model against a more flexible fully nonparametric specification and show that this test does not reject the partially linear model.

Keywords: semiparametric modeling, price impact of trades, infrequently traded stocks, market microstructure

AMS Subject Classification: 91B84, 62P20, 62G08, 62G10

JEL classification: C41, C51, C52, G14

Introduction

An important issue for institutional investors and other traders who have to deal with block trades is how large trades affect market prices. Since, in efficient markets, security prices move in response to the release of new information, transactions cause traders and market makers to update their beliefs and prices to be revised. Market impact reflects the change in the security price that is caused by a trade. The relation between trading volume and prices determines to what extent particular trading strategies such as order splitting affect the costs of trading.

An extensive literature is available on the price impact of trades. Hasbrouck (1991a, 1991b) shows that the persistent impact of a trade on the midprice is larger when the spread is wide and is more significant for firms with smaller market capitalization. Kavajecz and Odders-White (2001) analyze how the price impact of trades depends on the information in the limit order book. Dufour and Engle (2000), Zebedee (2001), and Spierdijk (2002) show that, for liquid stocks, the price impact of a trade is larger and converges to its full information value faster when subsequent trades are close together in time, i.e. when the trading intensity is high. Spierdijk, Nijman, and Van Soest (2002) show that the latter effect is even stronger for infrequently traded stocks. Additionally, Spierdijk et al. (2002) establish the phenomenon of ‘overshooting’: after a trade, prices temporarily exceed the full information price, before they mean revert to this level. Glosten and Harris (1988), Madhavan and Smidt (1991), and De Jong, Nijman, and Röell (1996) allow prices to depend linearly on trading volume and measure the impact of trades on transaction prices.

Although the majority of models for prices are linear in trading volume, several empirical studies investigate the existence of nonlinearities. Hasbrouck (1991a, 1991b) investigates nonlinearities in the impact of trades on midprices using a VAR-model. He establishes an increasing and concave relation between price impact and order flow for several stocks traded on the New York Stock Exchange (NYSE). Hausman, Lo, and McKinlay (1992) use a Box-Cox transformation of trading volume as explanatory variable in an ordered probit-analysis of discrete price changes. They apply the model to several stocks listed on the NYSE and show that the impact of trades on midprices is increasing in trading volume, in a nonlinear fashion that differs from stock to stock. Kempf and Korn (1999) establish a nonlinear, increas-

ing, and concave relation between trading volume and prices of German futures, using neural networks. De Jong, Nijman, and Röell (1995) use data on French stocks traded on the Paris Bourse and SEAQ International and find that transaction prices are affected by trading volume in a nonlinear way.

This paper extends the existing literature in several ways. We investigate the relation between price impact and trading volume for a sample of infrequently traded stocks listed on the New York Stock Exchange. We show that the commonly used parametric VAR-models as introduced by Hasbrouck (1991a, 1991b) impose strong proportionality and symmetry restrictions on the price impact of trades, although market microstructure theory provides many reasons why these restrictions would not hold. We analyze the less restrictive semiparametric partially linear model of Engle, Granger, Rice, and Weiss (1986) and Robinson (1988a, 1988b) and establish significant evidence for a nonlinear, asymmetric, increasing, and concave relation between trading volume and both immediate and persistent price impact. Moreover, we compare the relation between price impact and order size obtained in the partially linear model to the price-order flow relation generated by some commonly used parametric VAR-models and show that there are considerable differences. In contrast to the partially linear model, the parametric models do not capture the nonlinearities in the price-order flow relation. We use the approach of Whang and Andrews (1993) to test the model specification and reject the parametric specifications in favor of the partially linear model. We also test the partially linear model against a more flexible fully nonparametric specification and show that this test does not reject the partially linear model.

The setup of this paper is as follows. Section I reviews several theoretical models of market microstructure that predict a nonlinear or asymmetric relation between prices and order flow. Section II introduces the data that are used in this paper. Section III discusses some properties of parametric VAR-models. In Section IV a partial linear specification is used to model the relation between trading volume and midprices. Section V focuses on the immediate and persistent impact of a trade on midprices and the relation to trading volume. Section VI investigates the temporary effects that trades have on prices. Finally, Section VII summarizes and concludes.

I Explanations for nonlinearities and asymmetries in the price-order flow relation

Various models predict a nonlinear price-order flow relation; for example models of reputation (Seppi (1990)), stealth trading (Barclay and Warner (1993)), counter party search (Keim and Madhavan (1996)), and bullish-bearish information (Dridi and Germain (2000)).

Seppi (1990) distinguishes market orders and block trades. While market orders are submitted anonymously to the market, dealers know the identity of the institution that initiates a block trade. This allows the dealer and the institution to enter into additional commitments apart from agreeing on the price and the quantity. The particular commitment examined in Seppi (1990) is one of ‘no bagging’, which prohibits subsequent trading by the institution. When an institution is uninformed, postponing a trade until the previous trade has been completely executed is not much of a concession. However, when an institution does possess private information, they want to benefit from that and will therefore be less willing to wait with trading. Therefore, when the dealer and the institution agree on ‘no bagging’, the institution releases the signal that it is most likely uninformed. Such a commitment affects the information content of a trade, and, consequently, it influences the price impact of the trade. This implies that block trades and market orders have different impact on prices. Since block trades are usually larger than market orders, this results in a nonlinear price-order flow relation.

A different explanation for a nonlinear price-order flow relation is given by Barclay and Warner (1993). Using a sample of firms listed on the NYSE, the authors show that most of a stock’s cumulative price change takes place on medium-size trades, which supports the ‘stealth-trading’ hypothesis that privately informed traders concentrate their trades in medium sizes. Since medium-size trades are associated to informed trading, larger trades add relatively little additional information. This results in a concave price-order flow relation. More evidence for the stealth-trading hypothesis is found in Chakravarty (2001).

Keim and Madhavan (1996) model the phenomenon of an upstairs market, where large (block) trades are processed through a search-brokerage mechanism. That is, first an intermediary or broker identifies counter parties to trade, after which the order is sent to the downstairs market for final ex-

ecution. By contrast, smaller trades are directly routed to the downstairs market, where market makers, floor traders, and limit orders provide liquidity on demand. The authors show that spreading the order among more traders – this is what happens at the upstairs market – lowers the liquidity costs. Since the number of counter parties found by the block broker increases with trading volume, the temporary price impact of a block trade is a nonlinear function of order size.

Dridi and Germain (2000) proceed in a different way and model a financial market where informed traders receive a signal that perfectly reveals the sign of the difference between the liquidation value of the asset and its true value, but not the exact value of this difference. This type of information is called bullish or bearish. By endowing informed traders with a buy or sell signal only, the authors deviate from the assumptions made in the model of, for instance, Kyle (1985). Dridi and Germain (2000) show that the assumption of bullish and bearish information has a large impact on prices. They find that the optimal trading strategies for the informed traders in equilibrium are not linear and that, consequently, the price impact of trades is a nonlinear function of trading volume.

Another part of the literature is devoted to the explanation of asymmetries in the price-order flow relation. Although the empirical analysis of Kempf and Korn (1999) does not lead to any evidence that buys have more persistent impact on prices than sells, Karpoff (1988), Madhavan and Smidt (1991), and Chan and Lakonishok (1993) find that buy orders are more informative than sell orders and thus have larger persistent impact on prices. Chan and Lakonishok (1993) provide an institutional explanation for this phenomenon. They put forward that there may be several liquidity-motivated reasons why institutional investors decide to sell a stock. Therefore, selling a stock does not necessarily have to convey negative information. However, buying a stock is likely to convey favorable firm-specific news. This institutional explanation has been formalized by Saar (2001).

A very different issue is put forward by Huberman and Stanzl (2001). They show that, when the price impact of trades is time stationary, the 'no quasi-arbitrage' requirement is only satisfied when the permanent price impact of trades depends linearly on trading volume. The intuition is that, when trade size affects prices in a nonlinear way, certain self-financing trading strategies based on buying large amounts of stocks first and selling small amounts

later (or vice versa) lead to quasi-arbitrage. That is, these strategies yield infinite profits with infinite Sharpe-ratios because of the nonlinearity. Using the same arguments, the no-arbitrage condition rules out asymmetric impact of buys and sells on prices. The empirical evidence for a nonlinear relation between trading volume and price impact found in Hasbrouck (1991a, 1991b), Hausman, Lo, and McKinlay (1992), and Kempf and Korn (1999) would indicate the existence of arbitrage possibilities in the framework of Huberman and Stanzl (2001). If transaction costs would outweigh the gains derived from quasi-arbitrage strategies, this could justify a nonlinear relation between trading volume and price impact. However, Huberman and Stanzl (2001) establish quasi-arbitrage possibilities even in the presence of transaction costs. Empirical evidence that the price impact depends nonlinearly on trade size implies that transaction costs are larger than assumed in Huberman and Stanzl (2001) or that other assumptions made in the latter paper do not reflect reality.

II The data

We analyze a sample of stocks listed on the NYSE, taken from the *Trade and Quote* (TAQ) database. After ordering all NYSE stocks from least actively traded (decile one) to most actively traded (decile ten), we focus on stocks in liquidity decile 4. Since the semiparametric approach that we will follow requires a large number observations to achieve enough accuracy, we restrict the analysis to less frequently traded stocks taken from the fourth liquidity decile and do not consider stocks from the lower deciles. The results in Spierdijk et al. (2002) suggest that stocks in decile two and four have very similar behavior with respect to the price impact of trades, which further motivates the restriction to the fourth liquidity decile. We report results for a random subsample of the stocks in decile 4 and discuss in detail the results for the ‘representative stock’ of this decile (cf. Engle and Patton (2001)) which is the stock Commercial Intertech. The names of the five stocks considered in this paper are given in Table I.

On the NYSE the market starts at 9.30 AM with a call auction, while the remaining market is a continuous auction that ends each day at 4.00 PM. We remove all trades before 9.30 AM and after 4.00 PM. Moreover, we also delete trades that take place before the first quotes of the day are posted.

For each trade in every stock the following associated characteristics are recorded: transaction price p_t and unsigned trading volume $|y_t|$, where t indexes subsequent transactions (i.e. t indexes ‘transaction time’). To each trade we also associate a prevailing bid and ask quote, denoted by q_t^b and q_t^a . To obtain the prevailing quotes we use the ‘five-seconds rule’ by Lee and Ready (1991) which associates each trade to the quote posted at least five seconds before the trade, since quotes can be posted more quickly than trades are recorded. The five-second rule solves the problem of potential mismatching. On the basis of the prevailing quotes the prevailing midprice is obtained as $m_t = (q_t^b + q_t^a)/2$. The log return over the prevailing and subsequent midprice is expressed in basis points (bp) and denoted by $r_t = \log(m_{t+1}/m_t)$. The overnight returns are included in sample. We deal with dividend payments by deleting the first return in which the dividend payment is incorporated.

Since the transaction data provided by NYSE are not classified according to the nature of a trade (buy or sell), we use the Lee and Ready (1991) ‘midquote rule’ to classify a trade. With this rule, the prevailing midprice corresponding to a trade is used to decide whether a trade is a buy, a sell, or undecided. If the transaction price is lower (higher) than the midprice, it is viewed as a sell (buy). If the price is exactly at the midprice, its nature (buy or sell) remains undecided. To each trade we associate a trade indicator y_t^0 which indicates the nature of the trade: 1 (buy), -1 (sell), or 0 (undecided). In combination with unsigned trading volume $|y_t|$ that is associated to each trade, we construct signed trading volume y_t and signed log trading volume v_t .

It sometimes occurs that multiple trades take place at the same second. We follow Engle and Russell (1998) and treat multiple transactions at the same time as one single transaction and aggregate their trading volume and average their prices.

To get an idea of the sample properties of the data, we present an explorative data analysis. Table I shows some sample statistics (sample mean, median, and quantiles) of the trade characteristics of the stocks under consideration. Stocks in liquidity decile 4 are traded every 13-20 minutes (20-31 times a day). Moreover, the 5% sample quantile of trading volume varies from $-3,000$ to $-1,000$ shares and the 95% sample quantile lies between 1,000 and 2,500 shares.

Figure 1 displays the result of a simple (univariate) nonparametric kernel regression of midprice returns on trading volume, where we have kept trade size between its 5% and 95% sample quantile. A 95% point-wise confidence interval, based on the asymptotic distribution of the kernel estimator (as given in expression (B.11)), is also included. Figure 1 provides preliminary evidence for a nonlinear relation between returns and order flow.

III Properties of VAR-models for returns and trade size

By far the most popular model that is used in the literature to describe the dynamics between trading volume and prices is the VAR-model proposed by Hasbrouck (1991a, 1991b). Versions of the VAR-model have subsequently been used by Dufour and Engle (2000), Zebedee (2001), Spierdijk (2002), and Spierdijk et al. (2002), among many others. In this section we will show that the parametric VAR-models impose strong proportionality and symmetry restrictions on the price impact of trades, which are not grounded in theory as we have seen in Section I. In Section IV we will show how the validity of the restrictions can be tested and how they can be avoided.

Consider a simple, bivariate VAR-model for returns $(r_t)_t$ and signed trading volume $(y_t)_t$ defined in terms of the lag-operator as

$$\begin{pmatrix} a(L) & b(L) \\ c(L) & d(L) \end{pmatrix} \begin{pmatrix} r_t \\ y_t \end{pmatrix} = \begin{pmatrix} a_0 \\ b_0 \end{pmatrix} + \begin{pmatrix} \eta_{t,1} \\ \eta_{t,2} \end{pmatrix}, \quad (1)$$

with $(\eta_{t,1})_t$ and $(\eta_{t,2})_t$ mean-zero disturbances that are jointly and serially uncorrelated. We measure the price impact of trades by means of the cumulative impulse response function, cf. Hasbrouck (1991a, 1991b). Given a certain history of returns and trading volume up to time τ_t , the cumulative impulse response function at time τ_{t+k} corresponding to an unexpected buy of M shares at time τ_t is defined as

$$\alpha_{t+k|t}(M) = \mathbb{E}_{t-1}(r_t + \dots + r_{t+k} \mid \eta_{t,2} = M) - \mathbb{E}_{t-1}(r_t + \dots + r_{t+k}). \quad (2)$$

Hence, the cumulative impulse response function represents the expected price impact of an unexpected trade, relative to the expected price impact conditional on the history only. See, for instance, Koop, Pesaran, and Potter

(1996). Kraus and Stoll (1972) and Hasbrouck (1991a, 1991b) point out that the persistent price impact of an unexpected trade is naturally interpreted as the information content of the trade. The persistent impact is obtained for $k \rightarrow \infty$ in expression (2). In Appendix A we show that the price-impact functions corresponding to trades of different volumes are proportional to trading volume; i.e.

$$\frac{\alpha_{t+k|t}(M_1)}{\alpha_{t+k|t}(M_2)} = \frac{M_1}{M_2} \quad (3)$$

for any t, k . This implies that there is a linear relation between volume and prices in the bivariate VAR-model. We also show that the price impact of buys and sells obtained in the bivariate VAR-model is symmetric. This means that the magnitude of the impact of unexpected buys and sells of size is the same; i.e.

$$\alpha_{t+k|t}(M) = -\alpha_{t+k|t}(-M). \quad (4)$$

When we replace signed trading volume y_t by signed log trading volume v_t in the bivariate VAR-model, we get the log-linear VAR-model as proposed by Spierdijk (2002). The impulse response functions generated by the log-linear VAR-model are also symmetric. Moreover, the impulse response functions in the log-linear VAR-model are proportional to log trade size ; i.e. for any t, k

$$\frac{\alpha_{t+k|t}(M_1)}{\alpha_{t+k|t}(M_2)} = \frac{\log M_1}{\log M_2}. \quad (5)$$

Thus, the linear and log-linear VAR-models impose on two strong restrictions on price-impact functions: (log-) proportionality and symmetry. With respect to (log-) proportionality, this property determines the relation between price impact and trading volume and implies a (log) linearity of the price-order flow relation. However, in Section I we showed that there are various theoretical models predicting a complicated, nonlinear price-order flow relation, suggesting that it may be too restrictive to use a simple parametric specification for this relation. Regarding the symmetric impact of buys and sells, the literature suggests that this assumption may be too restrictive as we explained in Section I.

One way to make the bivariate VAR-model more flexible, is to include additional variables in its specification. The bivariate VAR-model has been extended in several ways. The extended linear VAR-model proposed by Hasbrouck (1991a) is a three-dimensional VAR-model for returns, trade sign,

and signed trading volume. Since small trades already have considerable price impact, the inclusion of trade sign in addition to signed trading volume helps to make the estimated price impact of these trades more accurate. The quadratic VAR-model, see Hasbrouck (1991a, 1991b), is a four-dimensional VAR-model for returns, trade sign, signed trading volume, and signed squared trading volume and is used by Hasbrouck (1991a, 1991b) to pick up any nonlinear effects in the price-order flow relation. Although the extended linear and the quadratic VAR-model do not have the proportionality property, the parametric structure of these models determines a priori the relation between price impact and trading volume. Moreover, the impulse response functions generated by the extended linear and the quadratic VAR-models are symmetric in trading volume.

IV A semilinear model for returns and trading volume

In Section I we provided some explanations for a nonlinear relation between price impact and trading volume. Moreover, we explained in Section III that some commonly used parametric VAR-models impose strong assumptions on the way volume affects prices, such as proportionality and symmetry. In this section we will use a more flexible semiparametric specification to allow for more complicated price-order flow relations.

The model that we use in this section is based upon the partially linear model introduced by Engle, Granger, Rice, and Weiss (1986) and Robinson (1988a, 1988b). The partially linear model is a semiparametric model, since the conditional mean of the dependent variable consists of both a parametric and a nonparametric part. In our setting, trading volume is allowed to affect returns in a nonlinear way, while lagged returns linearly affect current returns. Moreover, since we are not only be interested in immediate price effects but also in long-term price impact, we have to endogenize trading volume. We will use a partially linear model for trading volume as well, in which past returns and past trading volume affect current trading volume in a possibly nonlinear fashion. Furthermore, we will test the partially linear model against a wide range of alternative specifications, such as fully parametric and fully nonparametric models. For more details on estimation and testing of the partially linear model, we refer to Appendix B.

For notational convenience we write

$$z_{t-1} = (v_{t-1}, \dots, v_{t-m})' \text{ and } x_{t-1} = (r_{t-1}, \dots, r_{t-m})'. \quad (6)$$

For returns, the partially linear model is specified as

$$r_t = \beta_1' x_{t-1} + f_1(v_t, z_{t-1}) + \varepsilon_{t,1}, \quad \mathbb{E}(\varepsilon_{t,1} \mid x_{t-1}, v_t, z_{t-1}) = 0. \quad (7)$$

where β_1 is an $(m \times 1)$ vector of parameters and $f_1(\cdot)$ an unknown function. To model trading volume, we specify

$$v_t = \beta_2' x_{t-1} + f_2(z_{t-1}) + \varepsilon_{t,2}, \quad \mathbb{E}(\varepsilon_{t,2} \mid x_{t-1}, z_{t-1}) = 0. \quad (8)$$

where β_2 is an $(m \times 1)$ vector of parameters and $f_2(\cdot)$ an unknown function. Lagged returns appear linearly in the parametric part of equations (7) and (8). Contemporaneous and lagged trading volume, however, are in the nonparametric part of these specifications.

We compare the partially linear model to some commonly used parametric models that were introduced in Section III: the extended linear, the log-linear, and the quadratic VAR-model.

Estimation results

We consider the partially linear model for returns on the midprice as defined in equation (7). Following Hasbrouck (1991a, 1991b), we impose a low order on the recursive structure in the model ($m = 5$) which we estimate as explained in Appendix B. Point estimates and standard errors based on the procedure proposed by White (1980) for all five stocks are given in Table II. To get an idea of the impact of trade size on returns, we consider the function $f_1(\cdot)$ for the representative stock Commercial Intertech. Figure 2 shows a plot of the estimate $\hat{f}_{n,1}$ of $f_1(\cdot)$ as a function of its first argument; i.e. of

$$v_t \longrightarrow \hat{f}_{n,1}(v_t, \dots, v_{t-m}), \quad (9)$$

with the values of v_{t-1}, \dots, v_{t-m} set equal to the sample mean of signed trading volume¹. Clearly, $\hat{f}_{n,1}(\cdot)$ is an increasing and concave² function of its first argument.

We start with equation (7) and do several tests proposed by Whang and Andrews (1993) for which the corresponding p -values are reported in Table III. For an explanation of the testing procedure of Whang and Andrews (1993), see Appendix B. We start with testing the semilinear model against a fully nonparametric specification. For all stocks apart from F&M National, the

partially linear model is not rejected at a 5% level. For F&M National the semiparametric model is rejected at a 5% level, but the p -value of 0.0446 indicates that there is only limited evidence against the partially linear model. Furthermore, we test the log-linear, extended linear, and quadratic specifications – which are special cases of equation (7) as pointed out in Section III – against the partially linear model for all stocks under consideration. We reject at every reasonable significance level the null hypothesis that the parametric models are true. Subsequently, we test for autocorrelation in the disturbances (of the form (B.22)). For all five stocks the null of no autocorrelation in the disturbances of the return equation is not rejected at a 5% level. Finally, we use the test procedure of Whang and Andrews (1993) to detect conditional heteroskedasticity (of the form given in expression (B.25)). The null hypothesis of homoskedastic disturbances in the return equation (7) is rejected at a 5% significance level for Franklin Convey only. For this stock we assume that $\text{Var}(\varepsilon_{t,1} \mid v_t, \dots, v_{t-m}) = g(v_t, \dots, v_{t-m})$, since the Whang and Andrews (1993) test shows that there is no significant evidence for heteroskedasticity with respect to lagged returns. We estimate the function $g(\cdot)$ by means of a kernel regression of $\varepsilon_{t,1}^2$ on v_t, \dots, v_{t-m} .

To get a first impression of the relation between lagged trading volume and contemporaneous trading volume for the representative stocks Commercial Intertech, Figure 3 displays the results of a simple univariate kernel regression of lagged log trading volume on contemporaneous log trading volume (where we have kept trading volume between the 5% and 95% sample quantile). A 95% point-wise confidence interval is also given. Figure 3 provides preliminary evidence for a nonlinear relation between log trading volume and lagged log trading volume. Subsequently, we consider the model for trading volume as defined in expression (8), which we truncate at lag $m = 5$. The estimation results for the model given in equation (8) are reported in Table II. The function

$$v_{t-1} \longrightarrow \hat{f}_{n,2}(v_{t-1}, \dots, v_{t-m}), \quad (10)$$

with the values of v_{t-2}, \dots, v_{t-m} set equal to the sample mean of signed trading volume yields a plot that is very similar to that in Figure 3 and is therefore not displayed. The Whang and Andrews (1993) test results are given in Table III and show that, for all five stocks under consideration, there is no significant evidence against the partially linear structure assumed in (8). However, the null hypothesis of homoskedastic disturbances

$(\varepsilon_{t,2})_t$ in equation (8) is rejected at a 5% level for C&D Technology and Xtra Company. For these stocks, we proceed as before. We assume that $\text{Var}(\varepsilon_{t,2} \mid v_{t-1}, \dots, v_{t-m}) = h(v_{t-1}, \dots, v_{t-m})$ (again the results of the Whang and Andrews (1993) test show that there is no significant evidence for heteroskedasticity with respect to lagged returns) and we estimate $h(\cdot)$ in the usual way.

To estimate the price-impact functions, we need random values of the disturbances $(\varepsilon_{t,1})_t$ and $(\varepsilon_{t,2})_t$ as we will point out in Section V. We assume that the disturbances are independent and identically distributed. We do not impose any parametric assumptions on the distribution of the disturbances and will draw from the empirical distribution of the residuals.

V Immediate and persistent price impact and trade size

In this section we investigate the expected long-term and short-term price impact and the relation to trading volume.

To estimate the expected price impact of a trade, we use the impulse response function as defined in expression (2). Since the partially linear model is nonlinear in trading volume, we need the distribution of the disturbances $(\varepsilon_t)_t = (\varepsilon_{t,1}, \varepsilon_{t,2})'_t$ in equations (7) and (8) to estimate the impulse response function by means of simulation. We follow the approach of Hasbrouck (1991b) and assume that – given the test results obtained in Section IV – the disturbances are independent and identically distributed. Under this assumption it is possible to estimate the impulse response function by means of simulation. We jointly simulate paths of returns and trading volumes following an unexpected trade of M shares at time τ_0 , using the specification given in equations (7) and (8). For each path of returns, say $(r_t)_t$ for $t = 0, \dots, k$, we compute the corresponding cumulative midprice returns. Finally, we average the midprice changes over all $N = 10,000$ simulations to obtain the expected price impact. Appendix C explains in more detail how we simulate paths of returns and trading volumes and how we estimate impulse response functions in the partially linear model.

The market-impact curve

To gain more insight in the relation between price impact and trading volume, we consider the market-impact curve. This curve is defined as the

expected price impact of a trade (at a certain moment in time) as a function of unexpected trading volume, where we keep signed trading volume between the 5% and 95% sample quantiles of trade size. Note the difference between the impulse response function and the market-impact curve. The impulse response function is the expected price impact as a function of time for fixed initial trading volume, while the market-impact curve reflects the expected price impact as a function of the size of the initial trading volume, at a fixed moment in time.

We use the method of Appendix C to derive the market-impact curve from the impulse response functions at different moments in time: directly after the trade (immediate price impact) and in the long-run when the new efficient price has been reached (persistent price impact). As explained by Kraus and Stoll (1972) and Hasbrouck (1991a, 1991b), the permanent impact of a trade on prices reflects the change in the perception of the market maker due to the information contained in the initial trade. The difference between the persistent and the immediate price impact reflects temporary price movements caused by liquidity effects or lagged adjustment of prices to information. The market-impact curves for Commercial Intertech are shown in Figure 4.

The market-impact curves in Figure 4, together with the results in Tables IV and V, show that there is an increasing and concave relation between (immediate and persistent) price impact and trading volume. Tables IV and V are based on the market-impact curves and report the immediate and persistent price impact of buys and sells of different sizes for all five stocks under consideration (see the rows indicated by the abbreviation ‘PL’ that stands for ‘partially linear’) and show a similar increasing and concave relation between immediate and persistent price impact and order flow for the other stocks of the sample. This result is in line with the findings of Hasbrouck (1991a, 1991b), Hausman et al. (1992), and Kempf and Korn (1999). Since we establish an increasing and concave relation between the information content of a trade and the size of the trade, this provides evidence for the stealth-trading hypothesis of Barclay and Warner(1993) as discussed in Section I. The concave shape of the market-impact curves shows that trades of medium size contain relatively much information, suggesting that these trades have been initiated by informed traders. This is exactly the stealth-hypothesis as formulated by Barclay and Warner (1993).

The concavity of the market-impact curves shows that the immediate and persistent price impact are not proportional to trade size. Moreover, the price impact is not proportional to log trade size either. When we compare the market-impact curves generated by the partially linear model to those obtained in some commonly used parametric models, we find – in line with the test results obtained in Section IV – considerable differences. Figure 5 shows the market-impact curve for the representative stock Commercial Intertech in four different models: the partially linear model, the log-linear, the extended linear, and the quadratic VAR-model (see Section III). The quadratic VAR-model is used by Hasbrouck (1991a, 1991b) to pick up any nonlinear effects in the price-order flow relation. However, Figure 5 shows that the market-impact curve corresponding to the quadratic model differs considerably from the partially linear model. This suggests that the quadratic terms do not pick up well the nonlinearities that are reflected in the market-impact curve of the semiparametric model. The market-impact curve of the log-linear model also differs from the partially linear model. Tables IV and V report the immediate and persistent price impact of trades of different sizes in the log-linear VAR-model (indicated in Tables IV and V by the abbreviation ‘LL’), the extended linear VAR-model (abbreviated as ‘EL’) and in the quadratic VAR-model (abbreviated as ‘Q’). From these tables we see that the fully parametric models suffer from misspecification with respect to the relation between trading volume and returns. The largest differences are found for the representative stock Commercial Intertech. The difference between the partially linear and the log-linear model can amount as much as 34% (2,000 shares) for the immediate impact and 22% for the persistent impact (1,600 shares)³. The differences between the partially linear model and the extended linear VAR-model are even larger, 46% (immediate impact, 800 shares) and 35% (persistent impact, 1,600 shares). Finally, the largest differences with respect to the quadratic VAR-model are 43% (immediate impact, 800 shares) and 30% (persistent impact, 800 shares). Another important issue is the possible difference in persistent price impact between large buy and sell transactions. As explained in Section III, the parametric VAR-models yield symmetric impulse response functions, while the partially linear model allows for asymmetric impact of buys and sells on prices. We test the hypothesis of symmetry (explicitly given in expression (B.20)) using the approach of Whang and Andrews (1993). The null hypoth-

esis of symmetry is not rejected at a 5% level for Franklin Covey and Xtra Company; see Table III for the test results. We find significant evidence for asymmetric effects of buys and sells on prices for the remaining stocks C&D Technology, F&M National, and Commercial Intertech. However, Table V does not provide much evidence for the Chan and Lakonishok (1993) hypothesis (see Section I) that buys have larger persistent impact on prices than sells, since this is only the case for Commercial Intertech.

VI Temporary price effects and trade size

To gain more insight in the price adjustment process and the temporary price effects, this section investigates the expected price impact of a trade as a function of time.

Using the method explained in Appendix C we compute impulse response functions corresponding to trades of different volume and focus on the impact of trading volume on both immediate and persistent price impact over time. The impulse response functions corresponding to unexpected buys of 2,000, 1,200, 800, and 400 shares of Commercial Intertech are given in Figure 6. The impulse response functions for unexpected sells of the same size are shown in Figure 7.

In Section V we established a nonlinear relation between trading volume and both immediate and persistent price impact. The impulse response functions for the representative stock Commercial Intertech shown in Figures 6 and 7 show once more that there is a complicated, nonlinear price-order flow relation, since they are not (log-) proportional to trading volume.

Moreover, we compare the impulse response functions obtained in the partially linear model to the ones generated by the log-linear, the extended linear, and the quadratic VAR-model discussed in Section III, see Figure 8. This plot show the impulse response functions corresponding to a relatively large trade (1,800 shares; corresponding to the 95% sample quantile of trade size) in each of the four models, for the representative stock Commercial Intertech. We see that not only the market-impact curves in the partially linear model differ from those obtained in the parametric VAR-specifications as we showed in Section V, also the price impact over time is different. For the other stocks we find similar results.

In Section V we concluded that the partially linear model is flexible enough

to capture complicated price-order flow relations and we now see that it can also deal with complicated price movements over time. Clearly, a fully nonparametric model is even more flexible than the semiparametric partially linear model. The fact that the partially linear model is flexible enough for the data at hand is illustrated by the comparison to the fully nonparametric model. The tests of Whang and Andrews (1993) show that the semilinear model cannot be rejected in favor of the fully nonparametric model at a 5% level for all stocks under consideration (see Table III). Furthermore, the impulse response functions generated by a fully nonparametric model are very close to those generated by the partially linear model, which is in line with the test results.

The overshooting effect

From the impulse response functions of the representative stock Commercial Intertech in Figures 6 and 7 we see that the price-impact function ‘overshoots’ or ‘mean reverts’. This means that, after the buy, prices temporarily exceed the full information level, before they mean revert to this level. For sells we find a similar effect. The overshooting effect has also been established for infrequently traded stocks in the parametric VAR-model for returns, bid-ask spread, and trade sign in Spierdijk et al. (2002). In the parametric VAR-model used in latter paper, the overshooting effect – as the maximum over the adjustment path of the midprice minus its long-run equilibrium level – is generally larger than in the partially linear model. This can be explained by the fact that the model in Spierdijk et al. (2002) also condition on the bid-ask spread and the trading intensity to which the overshooting effect is strongly related.

Spierdijk et al. (2002) provide several explanations for the phenomenon of overshooting, such as order imbalances in the limit-order book, inventory effects, asymmetric information and the monopoly position of the market maker. However, since Spierdijk et al. (2002) model trade sign and not trading volume, they do not determine the relation between the size of the trade and the degree of overshooting. When we related the degree of overshooting to trade size, we find that there is a positive, approximately concave relation between overshooting and trade size. The positive relation between the overshooting effect and the size of the trade is consistent with the possible explanations of the overshooting effect as given in Spierdijk et al. (2002). Large trades will lead to larger imbalances in the limit-order book, which are

likely to cause more overshooting. Similarly, large trades are associated with larger inventory imbalances which are also likely cause more mean reversion in prices. The positive relation between trading volume and overshooting allows us to gain more insight in the role of asymmetric information – another possible explanation for the overshooting effect as pointed out in Spierdijk et al. (2002). In Easley and O’Hara (1987) it is put forward that the risk of informed trading is higher for large trades. Consequently, they demonstrate that information affects the price-quantity relationship in a complicated way, since both the size and the sequence of trades determine this relation in a situation of asymmetric information and event uncertainty. For example, when a large block buy takes place, the market maker increases his price for the next small trade. He does this because the large buy changes his perception of the risk of informed trading, which he considers to be higher. When more small trades follow, the market maker adjusts his perception on the risk of informed trading downwards, which leads to a partial recovery of prices. In this case, informational effects lead to overshooting and the overshooting effect increases with the volume of the block trade. See also Kraus and Stoll (1972) and Dann, Mayers, and Raab (1977) who find overshooting caused by block trades in their empirical analysis. Although the above argument of asymmetric information would equally well be applicable for frequently traded stocks, our empirical results show that prices of these stocks do not overshoot after a large trade. The risk of informed trading is larger for infrequently traded stocks than for frequently traded stocks (cf. Easley et al. (1996)), hence large trades in infrequently traded stocks will generally have more impact on the perception of the market maker than large trades in frequently traded stocks. This suggests that for frequently traded stocks the adjustments in perception are too small to lead to visible overshooting.

VII Conclusions

In this paper we investigated the relation between price impact and trading volume. The parametric VAR-models that have been used in the literature starting with Hasbrouck (1991a, 1991b) impose strong proportionality and symmetry restrictions on the price impact of trades, although market microstructure theory provides many reasons why these restrictions would not hold. We analyzed a more flexible semiparametric partially linear model of

Engle, Granger, Rice, and Weiss (1986) and applied the model to a sample of infrequently traded stocks listed on the NYSE in the year 1999. We established significant evidence for a nonlinear, asymmetric, increasing, and concave relation between trading volume and both immediate and persistent price impact. Moreover, we compared the relation between price impact and order size obtained in the partially linear model to the price-order flow relation generated by some commonly used parametric VAR-models and showed that there are considerable differences. In contrast to the partially linear model, the parametric models do not capture the nonlinearities in the price-order flow relation. We used the approach of Whang and Andrews (1993) to test the semiparametric specification and showed that the parametric models are rejected in favor of the partially linear model. We also tested the partially linear model against a more flexible fully nonparametric specification, but this test did not reject the partially linear model for the stocks under consideration.

The nonlinear, increasing, and concave price-order flow relation for order splitting can be explained in several ways (see for instance Seppi (1990), Barclay and Warner (1993), Keim and Madhavan (1996), and Dridi and Germain (2000)), but it implies that order splitting on the basis of price impact leads to an increase in the costs of trading. This suggests that, for order splitting, other costs such as temporary market impact and opportunity costs should be taken into account as well. This is left as an important topic for further research.

Acknowledgements

The main part of this paper has been written when Laura Spierdijk was affiliated to Tilburg University, The Netherlands. She is grateful for the financial support from the Netherlands Organization for Scientific Research she received during that period.

The authors would like to thank the participants of the economics seminar at the University of York (UK), the members of the Financial Stability Group at the Dutch Central Bank, the participants of the Young Financial Researchers Day in Leuven (Belgium), and the participants of the conference ‘The Econometrics of the Microstructure of Financial Markets’ organized by CentER, Tilburg University (The Netherlands). The usual disclaimer

applies.

A Properties of parametric VAR-models for returns and trading volume

In this appendix we focus on some important properties of parametric VAR-models.

Consider a bivariate VAR-model for returns $(r_t)_t$ and signed trading volume $(y_t)_t$, which we specify in terms of the lag-polynomial L as

$$\begin{pmatrix} a(L) & b(L) \\ c(L) & d(L) \end{pmatrix} \begin{pmatrix} r_t \\ y_t \end{pmatrix} = \begin{pmatrix} a_0 \\ b_0 \end{pmatrix} + \begin{pmatrix} \eta_{t,1} \\ \eta_{t,2} \end{pmatrix}, \quad (\text{A.1})$$

where

$$\begin{aligned} \mathbb{E}\eta_{t,i} &= \mathbb{E}\eta_{t,i}\eta_{s,i} = 0 & [t \neq s; i = 1, 2]; \\ \mathbb{E}\eta_{t,1}\eta_{s,2} &= 0. \end{aligned}$$

We are interested in the properties of the linear VAR-model defined in equation (A.1) regarding the price impact of trades. We measure the price impact of trades by means of the cumulative impulse response function, cf. Hasbrouck (1991a, 1991b). Given a certain history of returns and trading volume up to time τ_t , the cumulative impulse response function at time τ_{t+k} corresponding to an unexpected buy of M shares at time τ_t is defined as

$$\alpha_{t+k|t}(M) = \mathbb{E}_{t-1}(r_t + \dots + r_{t+k} \mid \eta_{t,2} = M) - \mathbb{E}_{t-1}(r_t + \dots + r_{t+k}). \quad (\text{A.2})$$

Hence, the cumulative impulse response function represents the expected price impact of an unexpected trade, relative to the expected price impact conditional on the history only. See, for instance, Koop, Pesaran, and Potter (1996). Kraus and Stoll (1972) and Hasbrouck (1991a, 1991b) point out that the persistent price impact of an unexpected trade is naturally interpreted as the information content of the trade. The persistent impact is obtained for $k \rightarrow \infty$ in expression (A.2). When we rewrite the VAR-model in expression (A.1) as a vector moving average, we obtain

$$\begin{pmatrix} r_t \\ y_t \end{pmatrix} = \begin{pmatrix} a^*(L) & b^*(L) \\ c^*(L) & d^*(L) \end{pmatrix} \left[\begin{pmatrix} a_0 \\ b_0 \end{pmatrix} + \begin{pmatrix} \eta_{t,1} \\ \eta_{t,2} \end{pmatrix} \right]. \quad (\text{A.3})$$

Let β_j denote the coefficient of L^j in the polynomial $b^*(L)$. Equation (A.3) shows that, at time τ_{t+k} , the impulse response function corresponding to an

unexpected trade of size M initiated at time τ_t equals

$$\alpha_{t+k|t}(M) = M \sum_{j=0}^k \beta_j. \quad (\text{A.4})$$

The persistent price impact, obtained for $k \rightarrow \infty$, yields $\alpha(M) = b^*(1)M$ as long-term impulse response.

From expression (A.4) we can derive some important properties of the impulse response function. Firstly, the price-impact functions corresponding to trades of different volumes are proportional; i.e.

$$\frac{\alpha_{t+k|t}(M_1)}{\alpha_{t+k|t}(M_2)} = \frac{M_1}{M_2} \quad (\text{A.5})$$

for any t, k . This implies that there is a linear relation between volume and prices in the VAR-model of equation (A.1). Secondly, the price impact of buys and sells is symmetric. This means that the magnitude of the impact of unexpected buys and sells of size is the same; i.e.

$$\alpha_{t+k|t}(M) = -\alpha_{t+k|t}(-M). \quad (\text{A.6})$$

B The partially linear model

The model that is used to specify the possibly nonlinear relation between prices and trading volume is the partially linear or semilinear model, introduced by Engle, Granger, Rice, and Weiss (1986) and Robinson (1988a, 1988b) and will be discussed in this appendix.

The partially linear model is a semiparametric model, since the conditional mean of the dependent variable y consists of both a parametric and a nonparametric part. The parametric part is a linear transformation of a vector of explanatory variables x of dimension k . The nonparametric part is formed by a transformation of another vector of explanatory variables, say z , that has dimension ℓ . Thus,

$$y_t = \beta' x_t + f(z_t) + \varepsilon_t, \quad \mathbb{E}(\varepsilon_t | x_t, z_t) = 0 \quad [t = 1, \dots, n], \quad (\text{B.1})$$

where β denotes a $(k \times 1)$ vector of parameters and $f(\cdot)$ an unknown function. Note that, to ensure identification, β should not contain an intercept. For the same reason x_t and z_t should not have any variables in common. In this paper the variable z_t is a vector (lagged) trading volumes, which may nonlinearly affect returns.

The coefficient β in the partially linear model is estimated by means of two kernel regressions. A kernel regression of y_t on z_t is used to estimate $\mathbb{E}(y_t | z_t)$ and a kernel regression of x_t on z_t is used to estimate $\mathbb{E}(x_t | z_t)$. Finally, β is estimated from the model

$$y_t - \mathbb{E}(y_t | z_t) = \beta'(x_t - \mathbb{E}(x_t | z_t)) + \varepsilon_t, \quad (\text{B.2})$$

which is implied by the initial model given in expression (B.1). To estimate (B.2) by means of OLS, $\mathbb{E}(y_t | z_t)$ and $\mathbb{E}(x_t | z_t)$ are replaced by the corresponding kernel estimates $\mathbb{E}_n(y_t | z_t)$ and $\mathbb{E}_n(x_t | z_t)$, respectively. Despite the nonparametric estimation stage that precedes the computation of the OLS-estimator, the estimation errors do not affect the asymptotic distribution of the OLS-estimator $\hat{\beta}_n$ and, under appropriate regularity conditions, the asymptotic distribution of $\hat{\beta}_n$ is

$$\sqrt{n}(\hat{\beta}_n - \beta) \xrightarrow{d} \mathcal{N}(0, A^{-1}BA^{-1}), \quad (\text{B.3})$$

where

$$A = \mathbb{E}[(x_t - \mathbb{E}(x_t | z_t))(x_t - \mathbb{E}(x_t | z_t))']; \quad (\text{B.4})$$

$$B = \mathbb{E}[(x_t - \mathbb{E}(x_t | z_t))(x_t - \mathbb{E}(x_t | z_t))' \varepsilon_t^2]. \quad (\text{B.5})$$

See for instance Lee (1996).

The function $f(z)$ can be estimated by means of a kernel regression of $y_t - \hat{\beta}_n' x_t$ on z_t . The kernel estimator $f_n(z)$ of $f(z)$ is given by

$$f_n(z) = \frac{g_n(z)}{h_n(z)}, \quad (\text{B.6})$$

where

$$g_n(z) = \frac{1}{n\gamma_n^\ell} \sum_{t=1}^n K\left(\frac{z_t - z}{\gamma_n}\right) (y_t - \hat{\beta}_n' x_t); \quad (\text{B.7})$$

$$h_n(z) = \frac{1}{n\gamma_n^\ell} \sum_{t=1}^n K\left(\frac{z_t - z}{\gamma_n}\right). \quad (\text{B.8})$$

Here $h_n(\cdot)$ is an estimate of the density $h(\cdot)$ of z_t . The function $K(\cdot)$ is a kernel function which is bounded, symmetric around zero and which integrates to one; e.g. the Gaussian density function (which is the kernel function used in this paper). Moreover, γ_n satisfies $n\gamma_n^\ell \rightarrow \infty$ and $\gamma_n \rightarrow 0$ for $n \rightarrow \infty$.

The kernel estimator $f_n(z)$ satisfies, for $n\gamma_n^{\ell+4} \rightarrow 0$,

$$\sqrt{n\gamma_n^\ell} (f_n(z) - f(z)) \xrightarrow{d} \mathcal{N}\left(0, h(z)^{-1} \text{Var}(\varepsilon | z) \int K^2(u) du\right). \quad (\text{B.9})$$

The latter asymptotic result is based on the assumption that the variable z is continuous. However, it is important to note that, in this paper, the regressor z represents the discrete variable trading volume. In this case a different asymptotic result applies, as proved in Bierens (1987). Suppose that $K(0) = 1$, $\gamma_n \rightarrow 0$ for $n \rightarrow \infty$, and that for every $\lambda > 0$

$$\sup_{|u| > \lambda/\gamma_n} \|K(u)\| \rightarrow 0. \quad (\text{B.10})$$

Then

$$\sqrt{n}(f_n(z) - f(z)) \xrightarrow{d} \mathcal{N}(0, \text{Var}(\varepsilon | z)h(z)^{-1}). \quad (\text{B.11})$$

Hence, in the discrete case, the problem of the curse of dimensionality is not encountered. For more details on kernel regression with discrete regressors we refer to Bierens (1987) and Delgado and Mora (1995).

As explained in Bierens (1987), the γ_n that minimizes the mean integrated squared error is of the form $cn^{-1/(\ell+4)}$, where the optimal value of c can be determined by means of leave-one-out cross-validation. Instead of using one single bandwidth parameter for all the regressors, we use the adaptive metric kernel estimation procedure proposed by Goutte and Larsen (2000). This approach allows for different bandwidths for each of the input variables. This is useful in our case, since the regressors in the kernel regression represent different lags of signed trading volume which are likely to be of different importance. We use the BFGS-algorithm of Broyden, Fletcher, Goldfarb, and Shanno (see Shanno (1970)) to carry out the numerical optimization procedure required for the V -fold cross-validation.

Whang and Andrews (1993) have developed several diagnostic tests for the validity of the assumptions underlying the partially linear model. The test statistic that they propose is based upon

$$\bar{r}_n(\beta, \pi) = \frac{1}{n} \sum_{t=1}^n r_t(\beta, \pi), \quad (\text{B.12})$$

where π represents a nonparametric regression function and r_t is chosen in such a way that, under the null hypothesis of correct specification, $\mathbb{E}(r_t) = 0$ for $t = 1, \dots, n$. Hence, we would expect that the sum in expression (B.12), with β and π replaced by $\hat{\beta}_n \xrightarrow{p} \beta$ and $\hat{\pi} \xrightarrow{p} \pi$ respectively, is close to zero when the model is correctly specified. Whang and Andrews (1993) show that, under sufficient conditions,

$$\sqrt{n}\bar{r}_n(\hat{\beta}_n, \hat{\pi}) \xrightarrow{d} \mathcal{N}(0, \Psi), \quad (\text{B.13})$$

under the null hypothesis, for some nonsingular covariance matrix Ψ . Their test statistic has the form

$$T = n\bar{r}'_n \hat{\Psi}^{-1} \bar{r}_n, \quad (\text{B.14})$$

where $\hat{\Psi}$ is a consistent estimator of the matrix Ψ of size $(q \times q)$, with q the dimension of r_t . They also show that the test statistic in (B.14) has, under the null hypothesis of correct specification, a χ^2_q limit distribution. For more details on the precise form of the test statistic and the underlying conditions we refer to Whang and Andrews (1993).

The null hypothesis that the fully parametric model is ‘true’ is formulated as

$$H_0 : \mathbb{P}(f(z_t) = \gamma' z_t) = 1, \quad (\text{B.15})$$

for some value of γ . The test statistic in (B.14) is based on

$$\begin{aligned} r_t = & [y_t - \mathbb{E}(y_t | z_t) - \bar{\beta}'(x_t - \mathbb{E}(x_t | z_t))] \\ & \times [x_t - \mathbb{E}(x_t | z_t)] / \mathbb{E}(\varepsilon_t^2 | z_t). \end{aligned} \quad (\text{B.16})$$

For the computation of (B.12), $\bar{\beta}$ is replaced by the OLS-estimate of β in the fully parametric model. Under the null hypothesis, $\bar{\beta} = \beta_0$ (where β_0 indicates the value of β under the null hypothesis that the fully parametric model is true) and $\mathbb{E}(r_t) = 0$ Whang and Andrews (1993) show that the test statistic in (B.14) is asymptotically χ^2_k distributed under the null hypothesis, with k the dimension of x_t . The test is consistent against alternatives for which $\hat{\beta}_{OLS} \xrightarrow{p} \beta_0$.

We also consider testing the partially linear model against the fully non-parametric model given by

$$y_t = f^*(x_t, z_t) + \varepsilon_t^*, \quad \mathbb{E}(\varepsilon_t^* | x_t, z_t) = 0 \quad [t = 1, \dots, n]. \quad (\text{B.17})$$

The null hypothesis that the semiparametric model is ‘true’ is formulated as

$$H_0 : \mathbb{P}(f^*(x_t, z_t) = \beta' x_t + f(z_t)) = 1, \quad (\text{B.18})$$

for some β . In this case the test statistic in (B.14) is asymptotically χ^2_1 distributed. In a similar fashion, the fully parametric model can be tested against the fully nonparametric model. To test this hypothesis the sample is split up in two independent subsamples. The sample is split up to avoid degeneracy of the limiting distribution of the test statistic and the subsamples

are used to estimate the model under the null hypothesis and the alternative hypothesis. The test statistic (B.14) is based on

$$r_t = [y_t - \mathbb{E}(y_t | z_t) - \beta'(x_t - \mathbb{E}(x_t | z_t))]^2 - [y_t^* - \mathbb{E}(y_t^* | x_t^*, z_t^*)]^2, \quad (\text{B.19})$$

where the variables with a star (“*”) refer to the second subsample.

Following the approach of Whang and Andrews (1993), we construct a test for symmetry of the function $f(\cdot)$. The null hypothesis that this function is symmetric in z is stated as

$$H_0 : \mathbb{P}(f(z_t) = -f(-z_t)) = 1. \quad (\text{B.20})$$

The test statistic in (B.14) is based on

$$r_t = [y_t - \hat{\beta}'_n x_t - f_n(z_t)]^2 - [y_t^* - \hat{\beta}'_n x_t^* + f_n(-z_t^*)]^2, \quad (\text{B.21})$$

and is asymptotically χ_1^2 distributed under the null hypothesis. We will later use this to test whether the impact of buys and sells on prices is symmetric. As will be explained later, we are interested in the distribution of the disturbances $(\varepsilon_t)_t$. Therefore, we need some statistical tests for autocorrelation and heteroskedasticity. The test procedure of Whang and Andrews (1993) can also be used for this. To test whether the disturbances $(\varepsilon_t)_t$ are autocorrelated, we consider a simple MA(1) error structure

$$\varepsilon_t = \rho u_{t-1} + u_t \quad [|\rho| < 1], \quad (\text{B.22})$$

where $(u_t)_t$ is a sequence of iid and zero mean variables such that u_t is independent of (x_t, z_t) . The null hypothesis of no autocorrelation is expressed as

$$H_0 : \rho = 0. \quad (\text{B.23})$$

The test statistic is based on

$$r_t = [y_t - \mathbb{E}(y_t | z_t) - \beta'(x_t - \mathbb{E}(x_t | z_t))] \times [y_{t-1} - \mathbb{E}(y_{t-1} | z_{t-1}) - \beta'(x_{t-1} - \mathbb{E}(x_{t-1} | z_{t-1}))] \quad (\text{B.24})$$

and is asymptotically χ_1^2 distributed under the null hypothesis of no autocorrelation and is similar to the test statistic proposed by Pagan and Hall (1983) for the same null hypothesis in the fully linear model. Note the assumption of MA(1) disturbances is not required. The same test statistic can

be used for autocorrelation other than the MA(1)-type; e.g. of the AR(1)-form. The test has power against any alternative for which the first-order autocorrelation is not zero.

Finally, we test for conditional heteroskedasticity in the disturbances $(\varepsilon_t)_t$ in equation (B.1). Under the null hypothesis, the disturbances ε_t are assumed to satisfy

$$\varepsilon_t = \sigma(w_t)\eta_t, \quad \sigma(w_t) = 1 + k(\gamma'w_t), \quad (\text{B.25})$$

where w_t is a p -dimensional vector of variables related to x_t and z_t , γ a $(p \times 1)$ vector of coefficients, $(\eta_t)_t$ a sequence of homoskedastic variables and $k(\cdot)$ a known function such that $k(0) = 0$. The null hypothesis of homoskedasticity is formulated as

$$H_0 : \gamma = 0. \quad (\text{B.26})$$

This test statistic is based on

$$r_t = w_t[(y_t - \mathbf{E}(y_t | z_t) - \beta'(x_t - \mathbf{E}(x_t | z_t)))^2 - \sigma^2] \quad (\text{B.27})$$

with $\text{Var}(\varepsilon_t) = \sigma^2$. Whang and Andrews (1993) show that the test statistic is – under appropriate conditions – asymptotically χ_p^2 distributed under the null hypothesis of homoskedasticity; i.e. its asymptotic distribution is independent of the functional form of $k(\cdot)$. It is noted that this test is the analogue for the partially linear model of the heteroskedasticity tests proposed by Breusch and Pagan (1979) and Koenker (1981) designed for the fully linear model.

C Estimation of the impulse response functions

In this appendix we explain how to estimate the impulse response functions in the partially linear model given by equations (7) and (8). Given a certain history of returns and trading volume up to time τ_t , the cumulative impulse response function at time τ_{t+k} corresponding to an unexpected buy of M shares at time τ_t is defined as

$$\begin{aligned} \alpha_{t+k|t}(M) &= \mathbf{E}_{t-1}(r_t + \dots + r_{t+k} | \varepsilon_{t,2} = \log(M)) \\ &\quad - \mathbf{E}_{t-1}(r_t + \dots + r_{t+k}). \end{aligned} \quad (\text{C.1})$$

The immediate price impact is obtained by taking $k = 0$ in expression (C.1) and the persistent impact is obtained for $k \rightarrow \infty$.

We simulate paths of returns and trading volumes in the following way:

- Initialize $v_{t+k} = \bar{v}_n, r_{t+k} = \bar{r}_n$ for $k < 0$;
- For $k = 0, \dots, K$:
 - For $k = 0$ set $e_{t+k,2} = \log(M)$. Randomly draw a disturbance $e_{t+k,1}$ and for $k > 0$ also a disturbance $e_{t+k,2}$ from the corresponding empirical distribution;
 - Compute, consecutively,

$$v_{t+k} = (r_{t+k-1}, \dots, r_{t+k-m})\hat{\beta}_{n,2} + f_{n,2}(v_{t+k-1}, \dots, v_{t+k-m}) + e_{t+k,2}. \quad (\text{C.2})$$

- Calculate

$$r_{t+k} = (r_{t+k-1}, \dots, r_{t+k-m})\hat{\beta}_{n,1} + f_{n,1}(v_{t+k}, \dots, v_{t+k-m}) + e_{t+k,1}. \quad (\text{C.3})$$

Subsequently, obtain

$$\alpha_{t+k|t}^M = r_0 + \dots + r_{t+k}. \quad (\text{C.4})$$

We repeat the above schedule $N = 10,000$ times and average the paths of price changes over the N simulations. This yields a sequence of estimates of expected price changes $\mathbb{E}_N(\alpha_{t+k|t}^M)$, for $k = 0, \dots, K$. In a similar way we estimate the expected price by conditioning on the history only (thus averaging out the initial unexpected trade), yielding $\mathbb{E}_N(\alpha_{t+k|t})$. The difference

$$\hat{\alpha}_{t+k|t}^N(M) = \mathbb{E}_N(\alpha_{t+k|t}^M - \alpha_{t+k|t}) \quad [k = 0, \dots, K] \quad (\text{C.5})$$

represents the estimated cumulative impulse response function.

ticker symbol	CHP	FC	FMN	TEC	XTR
company name	C&D Techn. Inc.	Franklin Covey Corp.	F&M National Corp.	Commercial Intertech Corp.	Xtra Corp.
#transactions	7,802	6,898	6,122	5,105	5,632
mean # trades a day	31	27	24	20	22
returns (bp)					
mean	0.3219	-1.1461	0.0013	0.0200	0.0440
median	0.0000	0.0000	0.0000	0.0000	0.0000
trade size (# shares)					
mean	-25	-134	33	-1	47
median	0	0	0	0	0
5%	-2,000	-3,000	-1,000	-2,000	-2,000
95%	2,000	2,000	1,000	1,800	2,500

Table I: Ticker symbols, company names, and some sample statistics

returns	CHP		FC		FMN		TEC		XTR	
	estim.	std.error								
$\beta_{1,1}$	0.0068	0.0176	-0.0416	0.0181	0.0224	0.0168	-0.0031	0.0215	0.0103	0.0196
$\beta_{1,2}$	0.0785	0.0122	-0.0092	0.0193	0.0522	0.0142	-0.0161	0.0153	0.0527	0.0145
$\beta_{1,3}$	0.0508	0.0125	0.0299	0.0136	0.0309	0.0126	-0.0009	0.0156	0.0388	0.0142
$\beta_{1,4}$	0.0420	0.0126	0.0161	0.0147	-0.0046	0.0340	0.0143	0.0158	0.0134	0.0143
$\beta_{1,5}$	0.0276	0.0100	0.0098	0.0116	-0.0024	0.0117	-0.0186	0.0129	0.0173	0.0119
volume										
$\beta_{1,1}$	0.0041	0.0024	-0.0037	0.0014	0.0033	0.0030	-0.0057	0.0014	0.0002	0.0039
$\beta_{2,1}$	0.0030	0.0025	-0.0022	0.0015	-0.0005	0.0030	-0.0017	0.0014	0.0013	0.0038
$\beta_{3,1}$	0.0008	0.0024	-0.0048	0.0014	-0.0052	0.0031	-0.0016	0.0014	0.0003	0.0038
$\beta_{4,1}$	0.0042	0.0024	-0.0003	0.0015	-0.0014	0.0030	-0.0013	0.0014	0.0015	0.0039
$\beta_{5,1}$	0.0015	0.0021	0.0005	0.0012	0.0005	0.0026	0.0015	0.0012	0.0086	0.0032

Table II: Estimation results for the partially linear model

This table reports the estimation results for the partially linear model for returns on the midprice and trading volume; cf. equations (7) and (8). The standard errors in the columns on the right-hand-side are computed from White (1980)'s heteroskedasticity-consistent covariance matrix.

	CHP	FC	FMN	TEC	XTR
returns					
semilinear vs. nonparametric	0.0636	0.3964	0.0446	0.0723	0.5304
linear vs. semilinear	0.0000	0.0000	0.0000	0.0000	0.0000
quadratic vs. semilinear	0.0000	0.0000	0.0000	0.0000	0.0000
log-linear vs. semilinear	0.0000	0.0000	0.0000	0.0000	0.0000
autocorrelation	0.8529	0.8988	0.9255	0.8969	0.9953
heteroskedasticity	0.3872	0.0288	0.3839	0.5191	0.2062
symmetric vs. nonsymmetric	0.0233	0.4108	0.0373	0.0046	0.6463
trading volume					
semilinear vs. nonparametric	0.0611	0.2579	0.0957	0.1703	0.0680
log-linear vs. semilinear	0.0000	0.0000	0.0000	0.0000	0.0000
autocorrelation	0.7072	0.1248	0.1603	0.6646	0.5563
heteroskedasticity	0.0494	0.1828	0.3982	0.1384	0.0122

Table III: The p -values of the Whang and Andrews (1993) tests

The p -values correspond to the Whang and Andrews (1993) tests (see Appendix B) described in the left-hand-side column and apply to the partially linear model for returns on the midprice and trading volume given in equations (7) and (8).

	CHP		FC		FMN		TEC		XTR	
	buy	sell	buy	sell	buy	sell	buy	sell	buy	sell
400 shares										
PL	9.2	-12.4	16.8	-14.2	10.9	-13.9	26.8	-20.3	7.8	-7.7
LL	9.3	-9.3	15.1	-15.1	11.6	-11.6	21.2	-21.7	7.1	-7.1
EL	7.6	-7.6	12.8	-12.8	9.8	-9.8	15.9	-15.9	5.4	-5.4
Q	7.2	-7.2	11.2	-11.2	9.5	-9.5	14.2	-14.2	5.3	-5.3
800 shares										
PL	10.6	-13.6	21.3	-19.3	13.5	-15.7	33.8	-29.3	9.1	-8.6
LL	10.3	-10.3	16.8	-16.8	13.0	-13.0	23.6	-23.6	8.0	-8.0
EL	8.1	-8.1	13.4	-13.4	11.4	-11.4	18.3	-18.3	6.1	-6.1
Q	8.0	-8.0	12.7	-12.7	12.4	-12.4	19.4	-19.4	6.1	-6.1
1,200 shares										
PL	11.4	-14.3	24.0	-22.6	15.1	-16.6	37.4	-34.4	10.2	-9.4
LL	11.0	-11.0	17.9	-17.9	13.7	-13.7	25.0	-25.0	8.4	-8.4
EL	8.6	-8.6	14.1	-14.1	13.0	-13.9	20.7	-20.7	6.8	-6.8
Q	8.8	-8.8	14.2	-14.2	15.2	-15.2	24.4	-24.4	6.9	-6.9
1,600 shares										
PL	11.9	-14.8	25.7	-25.0	16.1	-17.2	39.4	-37.5	11.1	-10.0
LL	11.4	-11.4	18.6	-18.6	14.3	-14.3	26.5	-26.5	8.8	-8.8
EL	9.1	-9.1	14.7	-14.7	14.6	-14.6	23.0	-23.0	7.5	-7.5
Q	9.5	-9.5	16.5	-16.5	17.9	-17.9	29.3	-29.3	7.8	-7.8
2,000 shares										
PL	12.3	-15.2	26.9	-26.8	16.7	-17.6	40.8	-39.3	11.7	-10.5
LL	11.7	-11.7	18.9	-18.9	14.7	-14.7	26.9	-26.9	9.1	-9.1
EL	9.7	-9.7	15.4	-15.4	16.1	-16.1	25.4	-25.4	8.2	-8.2
Q	10.3	-10.3	17.3	-17.3	20.4	-20.4	34.0	-34.0	8.6	-8.6

Table IV: Expected immediate price impact of buys and sells of different sizes

This table reports the expected immediate price impact (in bp) of buys and sells of different sizes (in shares) on midprices estimated by the partially linear model given by equations (7) and (8) (abbreviated as ‘PL’), the log-linear VAR-model (‘LL’), the extended linear VAR-model (‘EL’), and the quadratic VAR-model (‘Q’).

	CHP		FC		FMN		TEC		XTR	
	buy	sell	buy	sell	buy	sell	buy	sell	buy	sell
400 shares										
PL	22.2	-27.0	26.9	-26.6	18.0	-21.6	32.8	-26.9	14.6	-16.2
LL	24.9	-24.9	30.7	-30.7	20.4	-20.4	28.3	-28.3	13.7	-13.7
EL	24.1	-24.1	29.0	-29.0	18.2	-18.2	25.4	-25.4	12.0	-12.0
Q	23.8	-23.8	26.9	-26.9	17.7	-17.7	24.2	-24.2	11.8	-11.8
800 shares										
PL	24.6	-28.8	31.7	-31.2	21.0	-23.9	39.4	-35.2	16.4	-17.4
LL	27.8	-27.8	34.2	-34.2	22.8	-22.8	31.5	-31.5	15.3	-15.3
EL	24.5	-24.5	29.8	-29.8	19.2	-19.2	26.7	-26.7	12.7	-12.7
Q	24.4	-24.4	28.8	-28.8	20.5	-20.5	27.5	-27.5	12.8	-12.8
1,200 shares										
PL	25.8	-29.8	34.5	-34.0	22.7	-24.8	42.8	-39.9	17.8	-18.3
LL	29.5	-29.5	36.3	-36.3	24.2	-24.2	33.4	-33.4	16.2	-16.2
EL	24.9	-24.9	30.6	-30.6	20.2	-20.2	27.9	-27.9	13.3	-13.3
Q	25.0	-25.0	30.7	-30.7	23.1	-23.1	30.7	-30.7	13.8	-13.8
1,600 shares										
PL	26.6	-30.5	36.3	-36.0	23.7	-25.4	44.8	-42.7	18.9	-19.0
LL	30.7	-30.7	37.8	-37.8	25.2	-25.2	34.8	-34.8	16.9	-16.9
EL	25.2	-25.2	31.4	-31.4	21.1	-21.1	29.1	-29.1	14.0	-14.0
Q	25.6	-25.6	33.6	-33.6	25.5	-25.5	33.8	-33.8	14.8	-14.8
2,000 shares										
PL	27.2	-31.0	37.4	-37.4	24.3	-25.7	46.0	-44.4	19.8	-19.6
LL	31.7	-31.7	38.4	-38.4	25.9	-25.9	35.9	-35.9	17.4	-17.4
EL	25.6	-25.6	32.3	-32.3	22.1	-22.1	30.3	-30.3	14.7	-14.7
Q	26.2	-26.2	34.0	-34.0	27.6	-27.6	36.7	-36.7	15.8	-15.8

Table V: Expected persistent price impact of buys and sells of different sizes

This table reports the expected persistent price impact (in bp) of buys and sells of different sizes (in shares) on midprices estimated by the partially linear model given by equations (7) and (8) (abbreviated as ‘PL’), the log-linear VAR-model (‘LL’), the extended linear VAR-model (‘EL’), and the quadratic VAR-model (‘Q’).




Figure 1: Kernel regression: returns versus trading volume

This figure displays the relation between returns and trading volume for Commercial Intertech, based upon a kernel regression. The dashed lines indicate the boundaries of a 95% point-wise confidence interval based on the asymptotic distribution of the kernel estimator as given in expression (B.11) in Appendix B.




Figure 2: Relation between returns and trading volume

This figure shows the function $v_t \rightarrow \hat{f}_{n,1}(v_t, \dots, v_{t-m})$ in the model for returns and trading volume given by equations (7) and (8) for Commercial Intertech.

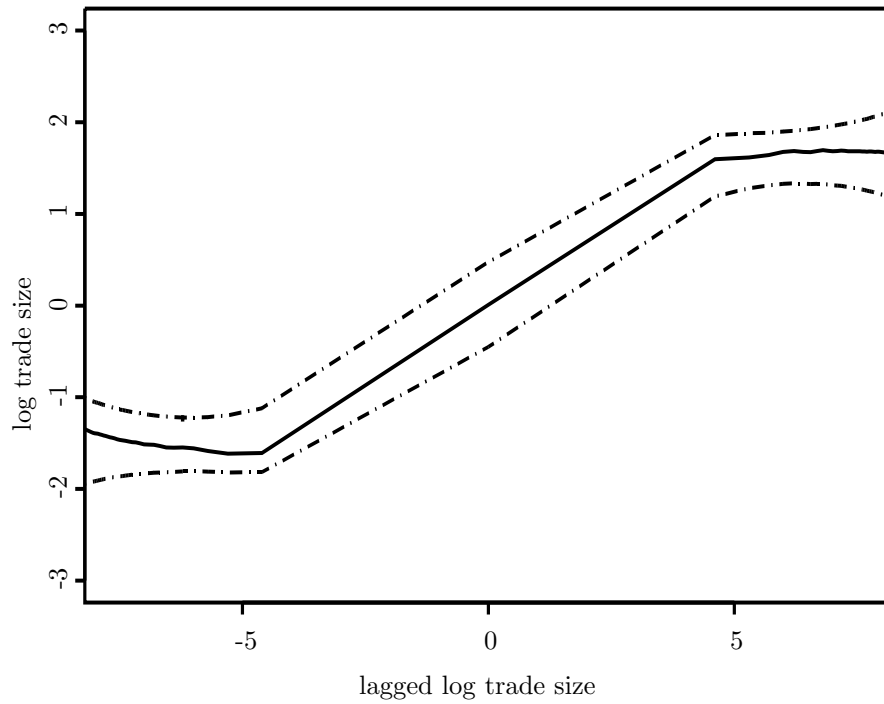


Figure 3: Kernel regression: log trading volume versus lagged log trading volume

This figure shows the relation between log trading volume and lagged log trading volume for Commercial Intertech, based upon a kernel regression. The dashed lines indicate the boundaries of a 95% point-wise confidence interval based on the asymptotic distribution of the kernel estimator as given in expression (B.11) of Appendix B.

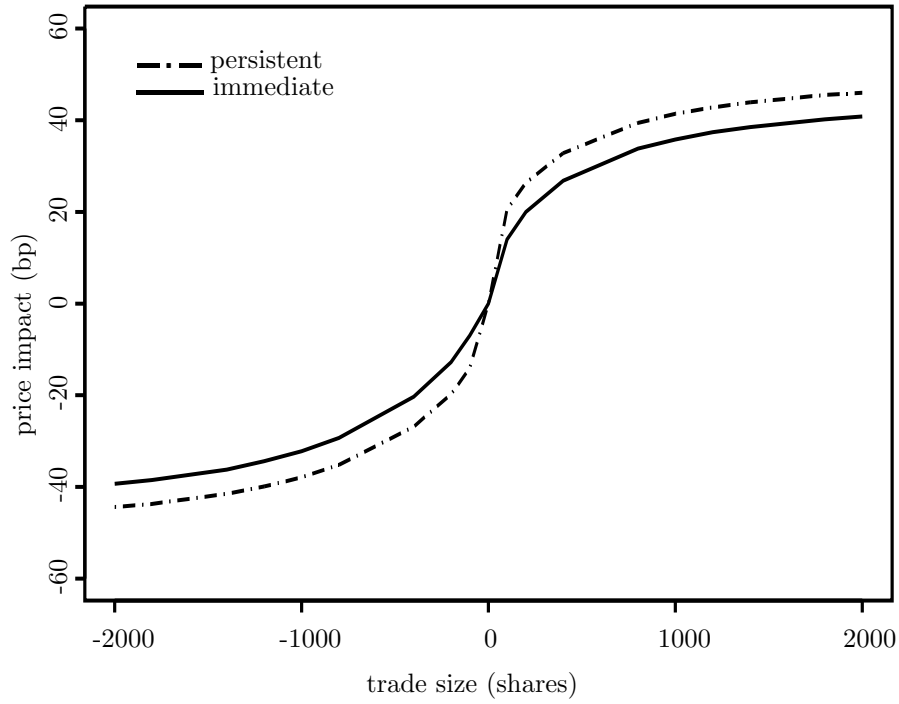


Figure 4: Market-impact curve: short-term and long-term

The market-impact curve shows the relation between signed trading volume (in shares) and expected price impact (in bp) in the partially linear model given by equations (7) and (8) for Commercial Intertech.

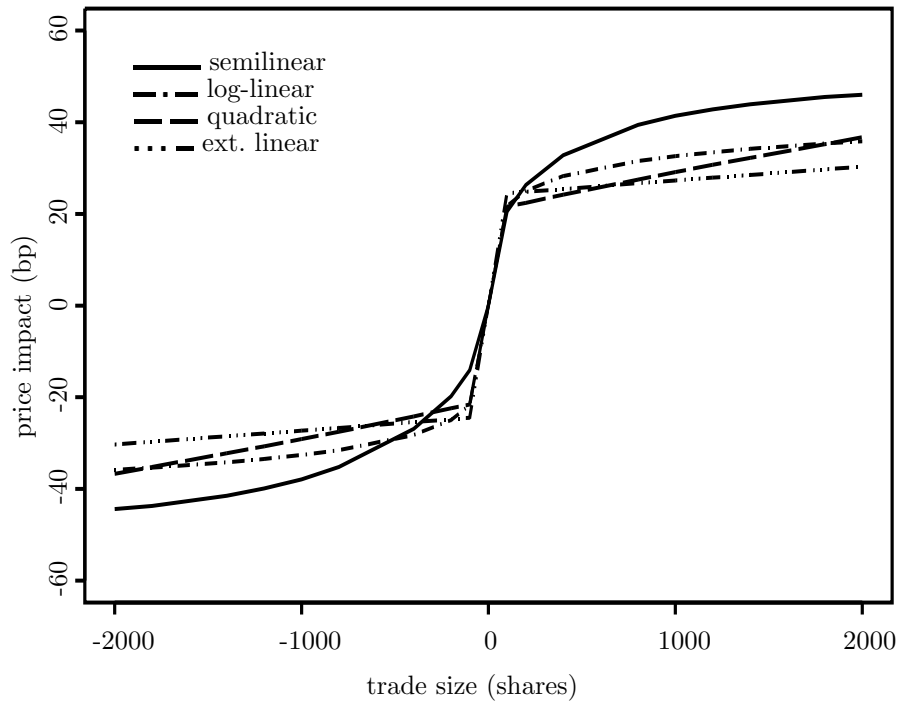


Figure 5: Long-term market-impact curve in four different models

The long-term market-impact curve shows the relation between signed trading volume (in shares) and expected persistent price impact (in bp) in the partially linear model, the linear VAR-model, the log-linear VAR-model, the extended linear VAR-model, and the quadratic VAR-model, for Commercial Intertech as defined in Section III.




Figure 6: Impulse response function for buys of different sizes

The expected price impact (in bp) of unexpected buys of 2,000, 1,200, 800, and 400 shares of Commercial Intertech in the partially linear model given by equations (7) and (8).

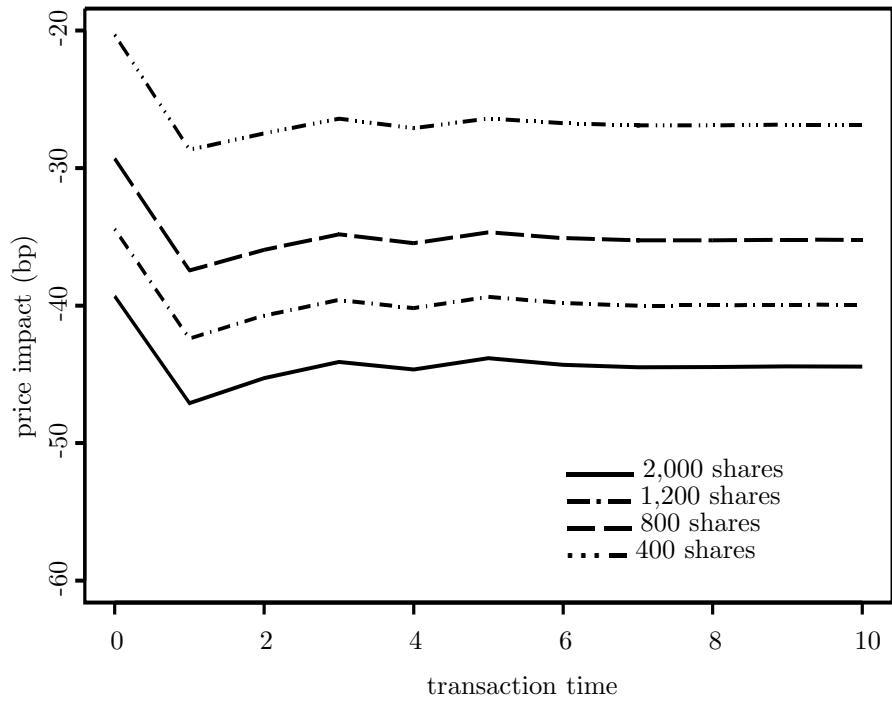


Figure 7: Impulse response function for sells of different sizes

The expected price impact (in bp) of unexpected sells of 2,000, 1,200, 800, and 400 shares of Commercial Intertech in the partially linear model given by equations (7) and (8).

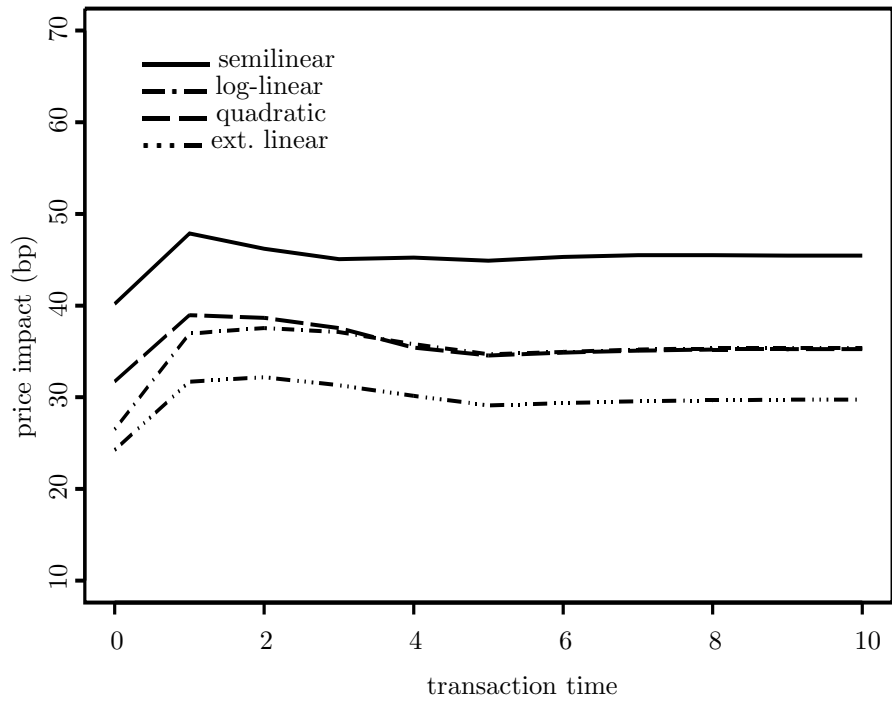


Figure 8: Impulse response functions in four different models

This figure shows the impulse response function corresponding to a buy of 1,800 shares of Commercial Intertech in the partially linear model, the log-linear VAR-model, the extended linear VAR-model, and the quadratic VAR-model as defined in Section III.

References

- [1] Admati, A. R., Pfleiderer, P., 1988, A Theory of Intraday Patterns: Volume and Price Variability, *The Review of Financial Studies*, 1, pp. 3-40.
- [2] Barclay, M.J., Warner, J.B., 1993, Stealth Trading and Volatility. Which Trades Move Prices?, *Journal of Financial Economics*, 34, pp. 281-305.
- [3] Bauwens, L., Giot, P., 2000, The Logarithmic ACD-model: An Application to the Bid-Ask Quote Process of three NYSE-Stocks, *Annales d'Economie et de Statistique*, 60, pp. 117-150.
- [4] Berndt, E.R., Hall, B.H., Hall, R.E., Hausman, J.A., 1974, Estimation and Inference in Nonlinear Structural Models, *Annals of Economic and Social Measurement*, pp. 653-665.
- [5] Bierens, H.J., 1987, Kernel Estimators of Regression Functions, in T. F. Bewley (ed.), *Advances in Econometrics: Fifth World Congress*, Cambridge University Press.
- [6] Breusch, T.S., Pagan, A.R., 1979, A Simple Test for Heteroskedasticity of Residual Security Returns, *Journal of Business and Economic Statistics*, 1, pp. 163-168.
- [7] Chakravarty, S., 2001, Stealth Trading: Which traders' trades move stock prices?, *Journal of Financial Economics*, 61, pp. 289-307.
- [8] Chan, L.K.C., Lakonishok, J., 1993, Institutional Trades and Intraday Stock Price Behavior, *Journal of Financial Economics*, 33, pp. 173-199.
- [9] Dann, L., Mayers, D., Raab, R.J., 1977, Trading Rules, Large Blocks, and the Speed of Price Adjustment, *Journal of Financial Economics*, 4, pp. 3-22.
- [10] De Jong, F., Nijman, Th.E., Röell, A., 1995, A Comparison of the Cost of Trading French Shares on the Paris Bourse and on SEAQ International, *European Economic Review*, 39, pp. 1277-1301.

- [11] De Jong, F., Nijman, Th.E., Röell, A., 1996, Price Effects of Trading and Components of the Bid-Ask Spread on the Paris Bourse, *Journal of Empirical Finance*, 3, pp. 193-213.
- [12] Delgado, M.A., Mora, J., 1995, Nonparametric and Semiparametric Estimation with Discrete Regressors, *Econometrica*, 63, pp. 1477-1484.
- [13] Diamond, D.W., Verrecchia, R.E., 1987, Constraints on Short-Selling and Asset Price Adjustments to Private Information, *Journal of Financial Economics*, 18, pp. 277-311.
- [14] Dridi, R., Germain, L., 2000, Bullish-Bearish Strategies of Trading: A Nonlinear Equilibrium, unpublished working paper, École Supérieure de Toulouse.
- [15] Drost F.C., Werker, B.J.M., 2001, Semiparametric Duration Models, Discussion Paper Tilburg University and CentER, nr. 2001-11, forthcoming *Journal of Business and Economic Statistics*.
- [16] Dufour, A., Engle, R.F., 2000, Time and the Impact of a Trade, *Journal of Finance*, 55, pp. 2467-2498.
- [17] Easley, D., O'Hara, M., 1987, Price, Size and the Information in Security Markets, *Journal of Financial Economics*, 16, pp. 69-90.
- [18] Easley D., O'Hara, M., 1992, Time and the Process of Security Price Adjustment, *Journal of Finance*, 47, pp. 577-605.
- [19] Engle, R. F., Granger, C.W.J., Rice, J., Weiss, A., 1986, Semiparametric Estimates of the Relation between Weather and Electricity Sales, *Journal of the American Statistical Association*, 81, pp. 310-320.
- [20] Engle, R.F., Hendry, D.F., Richard, J.F., 1983, Exogeneity, *Econometrica*, 51, pp. 277-304.
- [21] Engle, R.F., Lunde, 1998, Trades and Quotes: a Bivariate Point Process, unpublished working paper, University of California San Diego.

- [22] Engle, R. F., Patton, A.J., 2001, Impact of Trades in an Error-Correction Model of Quote Prices, forthcoming *Journal of Financial Markets*.
- [23] Engle, R.F., Russell, J.,1998, Autoregressive Conditional Duration: A New Model for Irregularly Spaced transaction data, *Econometrica*, 66, pp. 1127-1162.
- [24] Glosten, L., Harris, L., 1988, Estimating the Components of the Bid-Ask Spread, *Journal of Financial Economics*, 21, pp. 123-142.
- [25] Goutte, C., Larsen J., 2000, Adaptive Metric Kernel Regression, *Journal of VLSI Signal Processing Systems*, 26, pp. 155-167.
- [26] Hasbrouck, J., 1991a, Measuring the Information Content of Stock Trades, *Journal of Finance*, 66, pp. 179-207.
- [27] Hasbrouck, J., 1991b, The Summary Informativeness of Stock Trades: an Econometric Analysis, *Review of Financial Studies*, 4, pp. 571-595.
- [28] Hausman, J.A., Lo, A.W., McKinlay, A.C., 1992, An Ordered Probit Analysis of Transaction Stock Prices, *Journal of Financial Economics*, 31, pp. 319-379.
- [29] Ho, T.S.Y., Stoll, H., 1981, Optimal Dealer Pricing under Transactions and Return Uncertainty, *Journal of Financial Economics*, 9, pp. 47-73.
- [30] Huberman, G., Stanzl, W., 2001, Quasi Arbitrage and Price Manipulation, unpublished working paper, Yale School of Management/Columbia Business School.
- [31] Karpoff, J.M., 1988, Costly Short Sales and the Correlation of Returns with Volume, *Journal of Financial Research*, 11, pp. 173-188.
- [32] Kavajecz, K.A., Odders-White, E.R., 2001, An Examination of Changes in Specialists' Posted Price Schedules, *The Review of Financial Studies*, 14, pp. 681-704.

- [33] Keim D.B., Madhavan, A., 1996, The Upstairs Market for Large-Block Transactions: Analysis and Measurement of Price Effects, *The Review of Financial Studies*, 9, pp. 1-36.
- [34] Kempf, A., Korn, O., 1999, Market Depth and Order Size, *Journal of Financial Markets*, 2, pp. 29-48.
- [35] Koenker, R., 1981, A Note on Studentizing A Test for Heteroskedasticity, *Journal of Econometrics*, 17, pp. 107-112.
- [36] Koop, G., Hashem Pesaran, M.H., Potter, S.M., 1996, Impulse Response Functions in Nonlinear Multivariate Models, *Journal of Econometrics*, 74, pp. 119-147.
- [37] Kraus, A., Stoll, H., 1972, Price Impacts of Block Trading on the New York Stock Exchange, *Journal of Finance*, 27, pp. 569-588.
- [38] Lee, M.J., 1996, *Methods of Moments and Semiparametric Econometrics for Limited Dependent Variable Models*, Springer, New York.
- [39] Lee, C.M.C, Ready, M.J., 1991, Inferring Trade Direction from Intraday Data, *Journal of Finance*, 46, pp. 733-746.
- [40] Madhavan, A., Smidt, S., 1991, A Bayesian Model of Intraday Specialist Pricing, *Journal of Financial Economics*, 30, pp. 99-134.
- [41] Robinson, P.M., 1988a, Root- N Consistent Semiparametric Regression, *Econometrica*, 56, pp. 931-954.
- [42] Robinson, P.M., 1988b, Semiparametric Econometrics, *Journal of Econometrics*, 3, pp. 35-51.
- [43] Seppi, D.J., 1990, Equilibrium Trading and Asymmetric Information, *Journal of Finance*, 45, pp. 73-94.
- [44] Shanno, D.F., 1970, Conditioning of Quasi-Newton Methods for Function Minimization, *Mathematics of Computation*, 24, pp. 647-657.

- [45] Spierdijk, L., 2002, An Empirical Analysis of the Role of the Trading Intensity in Information Dissemination on the NYSE, Discussion Paper Tilburg University and CentER 2002-30, *forthcoming Journal of Empirical Finance*.
- [46] Spierdijk, L., Nijman, Th. E., Van Soest, A.H.O., 2002, The Price Impact of Trades in Illiquid Stocks in Periods of High and Low Market Activity, Discussion Paper Tilburg University and CentER nr. 2002-29.
- [47] Stoll, H.R, 1989, Inferring the Components of the Bid-Ask Spread: THEory and Empirical Tests, *Journal of Finance*, 44, pp. 115-134.
- [48] Whang, Y-J., Andrews, D., 1993, Tests of Specification for Parametric and Semiparametric Models, *Journal of Applied Econometrics*, 57, pp. 277-318.
- [49] White, H., 1980, A Heteroskedasticity-Consistent Covariance Matrix Estimator and a Direct Test for Heteroskedasticity, *Econometrica*, 48, pp. 817-838.
- [50] Zebedee, A.A., 2001, The Impact of a Trade on National Best Bid and Offer Quotes: a New Approach to Modeling Irregularly Spaced Data, *Journal of Multinational Management*, 11, pp. 363-383.

Notes

¹We find similar plots for the functions $v_{t-k} \rightarrow \hat{f}_{n,1}(v_t, \dots, v_{t-m})$ for $k = 2, \dots, m$, although the impact of trading volume on $f(\cdot)$ dies out for higher values of m .

²With ‘concave’ we mean concave for buys. As Figure 2 shows, the function is convex for sells. However, for convenience we will simply say that the function is ‘concave’.

³These percentages express the absolute difference between the price impact in the partially linear and the log-linear model, as % of the price impact in the partially linear model.