Journal of Forecasting
J. Forecast. **27**, 21–39 (2008)
Published online in Wiley InterScience
(www.interscience.wiley.com) **DOI**: 10.1002/for.1052

Forecasting Market Impact Costs and Identifying Expensive Trades

JACOB A. BIKKER,[1] LAURA SPIERDIJK,[2]*
ROY P. M. M. HOEVENAARS[3] AND
PIETER JELLE VAN DER SLUIS[4]
[1] *Supervisory Policy Division, Strategy Department, De Nederlandsche Bank, Amsterdam, The Netherlands*
[2] *Faculty of Economics and Business, Department of Economics and Econometrics, University of Groningen, Groningen, The Netherlands*
[3] *Department of Quantitative Economics, Maastricht University, Financial and Risk Policy Department, ABP, Schiphol, The Netherlands*
[4] *Department of Finance and Financial Sector Management, GTAA Fund, ABP Investments, Vrije Universiteit Amsterdam, Schiphol, The Netherlands*

ABSTRACT
Often, a relatively small group of trades causes the major part of the trading costs on an investment portfolio. Consequently, reducing the trading costs of comparatively few expensive trades would already result in substantial savings on total trading costs. Since trading costs depend to some extent on steering variables, investors can try to lower trading costs by carefully controlling these factors. As a first step in this direction, this paper focuses on the identification of expensive trades before actual trading takes place. However, forecasting market impact costs appears notoriously difficult and traditional methods fail. Therefore, we propose two alternative methods to form expectations about future trading costs. Applied to the equity trades of the world's second largest pension fund, both methods succeed in filtering out a considerable number of trades with high trading costs and substantially outperform no-skill prediction methods. Copyright © 2008 John Wiley & Sons, Ltd.

KEY WORDS market impact costs; forecasting extremes; trading cost management

INTRODUCTION

Since, in efficient markets, stock prices move in response to the release of new information, trading itself may cause prices to be revised. Loosely speaking, a buy trade tells the market that a stock is undervalued and, similarly, a sell trade indicates that a stock is overvalued. Market participants observe the information conveyed by trading and adjust their perceptions accordingly, which results

* Correspondence to: Laura Spierdijk, Faculty of Economics and Business, Department of Economics and Econometrics, University of Groningen, PO Box 800, NL-9700AV Groningen, The Netherlands. E-mail: l.spierdijk@rug.nl

in price movements. Other reasons for stock prices to move in response to trading are, for example, demand and supply imbalances and liquidity effects.

When an institutional investor sends an order to the market, it is usually not directly executed. In the case of a large order, it usually takes some time to find a counter party. Since large trades convey information, the price at trade execution is generally different than at trade initiation. Market impact costs occur when price effects cause execution prices to be less favorable than benchmark prices.

It is a well-known phenomenon that market impact costs can substantially reduce portfolio performance. A stock with a high gross return may end up with a relatively low net return when market impact costs are high. Therefore, trading costs are an important factor to consider when portfolio decisions are made.

There is a vast literature on trading costs; see Keim and Madhavan (1998) for an excellent survey. Usually, the literature distinguishes explicit and implicit trading costs. The explicit part consists of fixed costs, such as commissions, taxes, and fees. Implicit costs are built up of market impact costs (price impact), bid–ask spread, delay costs (the costs of adverse price movements that may occur when trading is postponed), and opportunity costs (the costs of not trading). Market impact costs are generally found to be the most important component of trading costs.

Often, a comparatively small group of trades causes the major part of market impact costs. For the equity trades studied in this paper, we find that only 10% of the trades determines 75% of total market impact costs. Consequently, reducing the trading costs of relatively few expensive trades would already lead to substantial savings on total trading costs. Since trading costs depend to some extent on controllable factors such as broker intermediation, investment style, trade timing, and trading venue (see, for example, Bikker *et al.*, 2007), investors may reduce trading costs by carefully controlling these factors.

Although there exists a vast literature on predicting portfolio returns, forecasting trading costs is a relatively unexplored issue. Nevertheless, the importance of forecasting trading costs has been widely recognized (see, for example, Cheng, 2003; Konstance, 2003). Since model-based forecasts of market impact costs can contribute to identifying expensive trades, they have the potential to play a crucial role in transaction costs management. One of the few studies in this field is that of Almgren *et al.* (2005), who describe a framework for forecasting market impact costs but do not evaluate the performance of their forecasting method.

This paper proposes two methods to forecast trading costs and to identify expensive trades in the future using today's available information. The first approach uses five 'buckets' to classify trades, where the buckets represent increasing levels of market impact costs. Each trade is assigned to a bucket depending on the probability that the trade will incur high market impact costs. The second method identifies expensive trades by considering the probability that market impact costs will exceed a critical level. When this probability is high, a trade is classified as potentially expensive. Applied to the pension fund data, both methods succeed in filtering out a considerable number of trades with high trading costs and substantially outperform no-skill prediction methods.

Since investors will eventually be interested in forecasts of net returns rather than predictions of trading costs only, we also discuss how to integrate the proposed costs forecasting methods with forecasts of optimal portfolio returns. We illustrate and evaluate the approach with a unique dataset containing the global equity trades executed by the world's second largest pension fund ABP in the first quarter of 2002.

The set-up of this paper is as follows. The next section describes the trading process at ABP. A description of the data, as well as some preliminary sample statistics, are given in the third section. This leads us to the fourth section, which assesses the performance of two new forecast methods to

DOI: 10.1002/for

predict future market impact costs and subsequently compares these two methods to the traditional forecasting approach. Next, we discuss the applicability of the forecasting methods in portfolio management in the fifth section. The sixth section concludes the paper.

TRADING PROCESS

The data used in this paper consists of the equity transactions of the world's second largest pension fund ABP ('Algemeen Burgerlijk Pensioenfonds') during the first quarter of 2002. ABP has about 2.4 million clients and an invested capital of approximately €190 billion,[1] corresponding to one third of total Dutch pension fund assets. In the period under consideration there were ten internal funds in ABP's equity group, apart from the externally managed funds.

Quantitative approach

Three funds followed a systematic quantitative approach for Japan, the USA, and Europe. The approach was to make bets on individual stocks, while keeping the overall portfolio sector neutral. The horizon of the fund varied from 1 to 6 months. Trading was usually done on the basis of information available up to the previous day. Variables in the model-based process were company-specific characteristics such as book-to-price ratios and analysts' forecasts, but also short-term and longer-term technical indicators, capturing short-term mean reversion and long-term momentum. New signals were usually generated at the beginning of every month. For all three funds the portfolio managers felt an urgency to trade quickly in response to the new signals, although a careful analysis of the forecasting signals of these models revealed a deterioration of its forecasting power only after 6 months.

Fundamental approach

There were seven fundamental funds. Three groups of portfolio managers each ran a European fund and a US fund. One group ran a similar Canadian fund. These funds had a fundamental macroeconomic approach to sector rotation. Here the approach was to make bets on sectors, while being neutral on stocks. These funds typically held this view for a longer-term horizon, from 6 months or more. The portfolio managers did not have the urge to trade immediately, although most of the trades were executed within 1 day. There were no views on individual companies, so the trades only comprised sector bets. Usually groups of companies were bought and sold tracking a certain sector very closely. Obviously, the groups crossed trades with each other before going to the market, but these trades are not included in the dataset.

All funds had a mandate with a maximum tracking error of 2% per annum with respect to a certain benchmark and a certain outperformance target. All funds were essentially enhanced index funds. Each of them also had a long-only constraint. The benchmarks of these mandates were the S&P500 for the USA, the MSCI Europe for Europe, the MSCI Japan for Japan, and the TSE300 for Canada.[2] Most of the trades that took place were for rebalancing purposes to keep portfolios in line with the original allocation. There were also moderate shifts due to changing tilts towards individual stocks for the quantitative portfolios and towards sectors for the fundamental approach.

[1] This is the total invested capital in December 2005.
[2] This is currently the S&P/TSX Composite.

DOI: 10.1002/for

Before turning to the trading process at ABP, we make notice of three types of trades. A principal trade is a transaction between the pension fund and the broker, in which case the broker buys or sells stocks from or to the pension fund at a predetermined price. Hence, the risk is transferred to the broker. The broker takes on the other side of the trade and tries to execute the trade in the open market. An agency trade is a trade between the pension fund and a counter party, where the broker acts solely as an intermediate party. Thus, an agency trade involves two clients of the brokerage firm, one of them being the pension fund. Single trades apply to difficult trades that are traded separately, not necessarily with packages of other stocks. In the case of single and agency trades the risk resides with ABP. The broker represents the client (ABP) and acts in the client's best interest.

For all the trades the trading process during the first quarter 2002 was as follows. A portfolio manager formed his or her portfolio. Subsequently, he or she approached a trader at ABP. Together they discussed the proposed trade. In most cases the trader would leave out some parts of the trades (say 10%) for reasons of perceived cost reduction and would execute these as an agency or single trade elsewhere in the market. Next, the trader approached at most two of the large brokerage firms for the remaining trades and revealed some of the characteristics of the trade (volume, USA or Europe, quantitative or fundamental, sector decomposition and a judgment on the complexity of the trade). The choice of brokerage firms was based on experience of the trader. Only the largest brokers could make competitive offers in case of principal trades, although sometimes a smaller one had an edge in some market segments, for example Japan. Based on the characteristics of the trade the broker made an offer for a principal trade. The offer of the broker was compared over brokers and with the trader's own systems and experience. If the offer was acceptable, then a principal trade was executed. Otherwise, the trade was executed as an agency trade.

PRELIMINARY DATA ANALYSIS

The dataset contains detailed information on 3721 equity trades during the first quarter of 2002, with a total transaction value of €5.7 billion. Of these trades, 1962 are buys and 1759 are sells executed in Europe, the USA, Canada, and Japan. The ten internally managed equity portfolios in our sample had a total market value of €20 billion. Average trade size for buys (sells) is more than 70,000 (84,000) shares and the average value of a trade equals almost €1.5 million (respectively €1.6 million).

Data and definitions

For each transaction the data provide the execution price and the price of the stock just before the trade was passed on to the broker. Moreover, the data also tell when the trade was submitted to the broker and when it was executed. Trades that were split up into several sub-trades are considered as one single trade, if a trader at ABP decided to split up the trade. The data contain about 0.5% of such 'trade packages'. Orders split up by portfolio managers are treated as individual trades, since it is not known whether the trader eventually split up the trade in the same way as the portfolio managers did. Additionally, the data include detailed information on several trade, exchange, and stock specific characteristics. Table I provides a complete list of the variables in the dataset, including their abbreviations and definitions.[3] For a complete description of the data, we refer to Bikker *et al.* (2007).

[3] The dataset has been created on the basis of the post-trade analysis provided by ABP, in combination with additional data from Factset and Reuters. The information on the characteristics of the exchanges under consideration were obtained from the World Federation of Exchanges and the various exchanges themselves.

DOI: 10.1002/for

Table I. Potential determinants of market impact costs and their definitions

Determinants of market impact costs	Definition
momentumperc	5-day volume-weighted average return prior to trading (in %)
volatility	logarithm of 30-day individual volatility (i.e., average squared return) prior to trading (in %)
tradesize	square root of trade size relative to 30-day average daily volume prior to trading (in %)
marketcap	logarithm of market capitalization 3 months prior to trading (in billion Euro)
adv	logarithm of 30-day average daily trading volume of stock (in shares)
exprice	logarithm of execution price of stock (in Euro)
agencysingledum	0/1 variable for agency/single (1) or principal (0) trades
growthdum	0/1 variable for growth stocks
quantdum	0/1 variable for trades executed by quantitative (1) or fundamental (0) fund
preopendum	0/1 variable for trades sent to broker during pre-opening of the market
morningdum	0/1 variable for trades sent to broker during in the morning (after pre-opening)
middaydum	0/1 variable for trades sent to broker during in the afternoon
(Mondaydum)	0/1 variable for trades executed on Monday
(Tuesdaydum)	0/1 variable for trades executed on Tuesday
Wednesdaydum	0/1 variable for trades executed on Wednesday
Thursdaydum	0/1 variable for trades executed on Thursday
Fridaydum	0/1 variable for trades executed on Friday
earlymonthdum	0/1 variable for trades executed at the beginning of the month
Jandum	0/1 variable for trades executed in January
Febdum	0/1 variable for trades executed in February
(Marchdum)	0/1 variable for trades executed in March
NYSEdum	0/1 variable for trades executed on NYSE
Nasdaqdum	0/1 variable for trades executed on Nasdaq
Torontodum	0/1 variable for trades executed on Toronto Stock Exchange
Londondum	0/1 variable for trades executed on London Stock Exchange
Tokyodum	0/1 variable for trades executed on Tokyo Stock Exchange
upstairsdum	0/1 variable for trades executed on exchange with upstairs market
(dealerdum)	0/1 variable for trades executed on exchange with dealer market
LOBdum	0/1 variable for trades executed on exchange with electronic limit order book
(floordum)	0/1 variable for trades executed on exchange with trading floor
(hybriddum)	0/1 variable for trades executed on exchange with hybrid market (LOB+dealers)
mcapdom	logarithm of domestic market capitalization of the exchange on which the stock was traded (in billion Euro)
consumerdiscrdum	0/1 variable for stocks in consumer discretionary sector
consumerstdum	0/1 variable for stocks in consumer staples sector
energydum	0/1 variable for stocks in energy sector
finservdum	0/1 variable for stocks in financial services sector
healthdum	0/1 variable for stocks in health sector
ITdum	0/1 variable for stocks in IT sector
materdum	0/1 variable for stocks in materials sector
telecomdum	0/1 variable for stocks in telecommunications sector
utilitiesdum	0/1 variable for stocks in utilities sector
(industrydum)	0/1 variable for stocks in industry sector

The dummy variables in parentheses have not been included in the estimated models to avoid exact collinearity, but are included in the table for completeness. Since there are virtually no trades on Monday during the sample period, we exclude both the Monday and the Tuesday dummy. Since there is only one hybrid market in our sample (Nasdaq) for which we include a dummy in the model, we do not include the hybrid dummy in the model. Similar arguments explain why we do not include a dummy for floor trading: it coincides with the NYSE dummy.

DOI: 10.1002/for

Measuring market impact costs

Market impact costs occur when the execution price of a trade is worse than the benchmark price. Hence, in order to forecast these costs, a benchmark price has to be chosen. We opt for the pre-execution benchmark, in line with Wagner and Edwards (1993), for example. More precisely, we take as benchmark the price at the moment that the order was passed to the broker. Furthermore, we correct for market-wide price movements during trade execution, as in Chan and Lakonishok (1995, 1997). The MSCIWorld industry group indices are used as a proxy for these market movements. Thus, for a buy transaction in stock i at time t market impact costs ($\mathbf{C}_{it}^{\mathrm{B}}$) are defined as

$$\mathbf{C}_{it}^{\mathrm{B}} = \underbrace{\log\left(\mathbf{P}_{it}^{\mathrm{exe}}/\mathbf{P}_{it}^{\mathrm{pt}}\right)}_{\text{price impact}} - \underbrace{\log\left(\mathbf{M}_{it}^{\mathrm{exe}}/\mathbf{M}_{it}^{\mathrm{pt}}\right)}_{\text{market-wide price movement}} \quad (1)$$

where $\mathbf{P}_{it}^{\mathrm{exe}}$ and $\mathbf{P}_{it}^{\mathrm{pt}}$ denote the execution and pre-trade price of stock i at time t, respectively. $\mathbf{M}_{it}^{\mathrm{exe}}$ and $\mathbf{M}_{it}^{\mathrm{pt}}$ denote the value of the MSCI World industry group index corresponding to stock i at the time of the execution of the trade and at the pre-trade moment, respectively. In a similar way, we define market impact costs of sells. For both buys and sells, positive market impact cost indicates that a trade has been executed against a price worse than at trade initiation.

Some sample statistics

To get a first impression of the magnitude of trading costs, we calculate average market impact costs. We weight each observation by the euro value of the trade, to ensure that larger trades are more important than smaller ones. Average market impact costs of buys equal 20 basis points (bp) with a standard deviation of 6 bp and those of sells 30 bp with a standard deviation of 7 bp. Including commissions, these costs equal 27 bp (6 bp standard error) for buys and 38 bp (7 bp) for sells. These price effects are relatively moderate compared to other studies (see Bikker *et al.*, 2007).

Figure 1 displays the contribution of each trade to total market impact costs. Starting with all trades sorted from cheap to expensive, the horizontal axis denotes the percentage of trades executed (in the range 0–100%). The vertical axis represents total trading costs (including commission) in millions of euros incurred by the executed trades. The 25% cheapest trades have negative market impact costs, whereas the 35% most expensive trades incur positive trading costs. The remaining 40% of medium-expensive trades have market impact costs close to zero. Together, they yield a convex and asymmetric 'market impact costs smile'. The convexity in Figure 1 implies that the 10% most expensive trades cause about 75% of total market impact costs. Consequently, the investor could realize substantial savings on total trading costs if he or she would be able to identify a few expensive trades before actual trading and to reduce the trading costs of these trades, e.g. by more careful monitoring and execution. We notice that Figure 1 is on an aggregate level and does not distinguish between differences across industry sectors, for instance. When we split up the results per sector, it becomes clear that there are differences across industries. In some sectors there are only a few cheap trades and many expensive trades or vice versa, which affects the degree of asymmetry of the market impact costs smile.[4]

We make these results more explicit by means of simulation. For this purpose, we consider an investor that has a certain skill in the range 0–100% to identify expensive trades correctly before actual trading takes place. With a skill of 100%, he or she is able to rank all trades correctly

[4] The results per industry sector are available from the authors upon request.

DOI: 10.1002/for

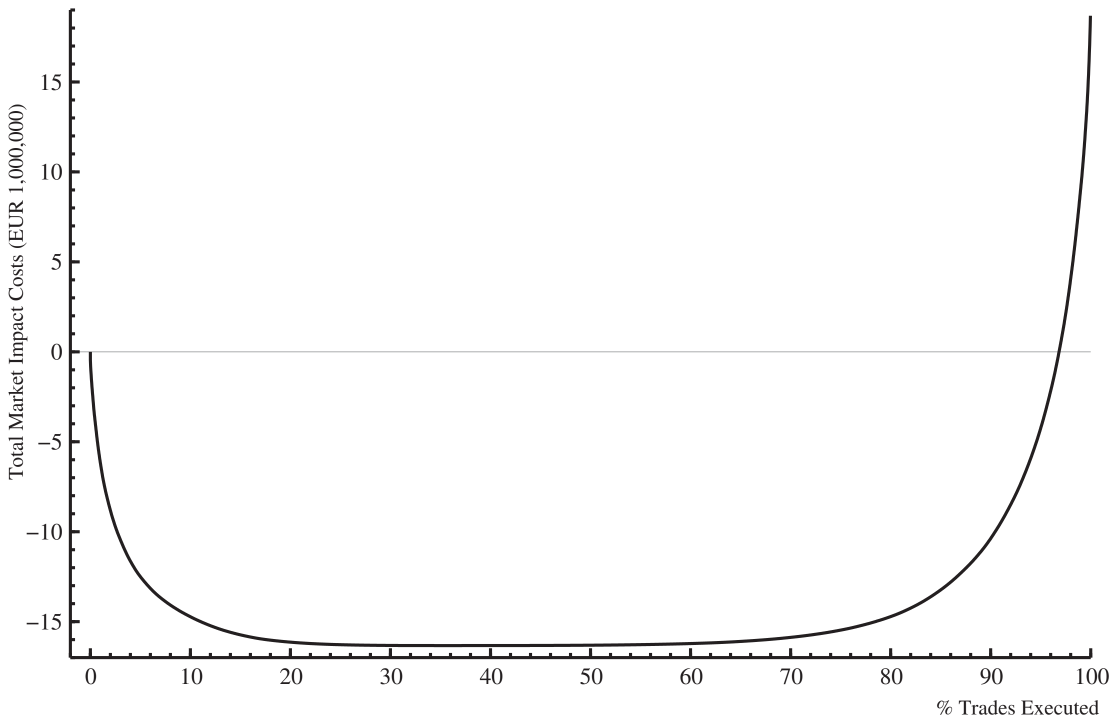

Figure 1. Trades and their contribution to total market impact costs This figure displays total market impact costs plus commission (in millions of euros) as a function of the percentage of trades executed. For example, when the 3% most expensive trades are not executed, the costs of the remaining 97% cheapest trades sum to zero

according to future trading costs ('perfect foresight'). With a 0% skill, the investor's ranking of the stocks is completely random.[5] At the same time we consider a cost reduction percentage that applies to the selected trades as a result of more careful treatment. For all skill levels and each percentage of cost reduction per trade, we simulate the corresponding total realized savings on trading costs.[6] The resulting three-dimensional graph in Figure 2(a) displays the relation between investor skill (x-axis), cost reduction per trade (y-axis), and total expected savings (z-axis). For instance, an investor skill of 20% in combination with the same cost reduction percentage per trade results in total expected savings of almost €1.4 million.

Although the extremes of no skills and perfect foresight are trivial, the simulation reveals various nontrivial patterns in Figure 2(a). The contour plot in Figure 2(b) (where each contour line represents an additional saving of €5 million) highlights the nonlinear relation between trading costs and cost reduction. Moving from southwest to northeast in Figure 2(b), the distances between the contour lines get smaller. Saving an additional €5 million requires a comparatively large improvement in either skill or cost reduction whenever these are low. By contrast, saving another €5 million needs

[5] For any skill of p% with $0 < p < 1$, p% of all trades is ranked correctly and the other $(1 - p)$% is ranked randomly over the remaining positions.

[6] We repeat this 1000 times and average the realized savings over the simulation runs to obtain the expected savings for each skill level and cost reduction percentage.

DOI: 10.1002/for

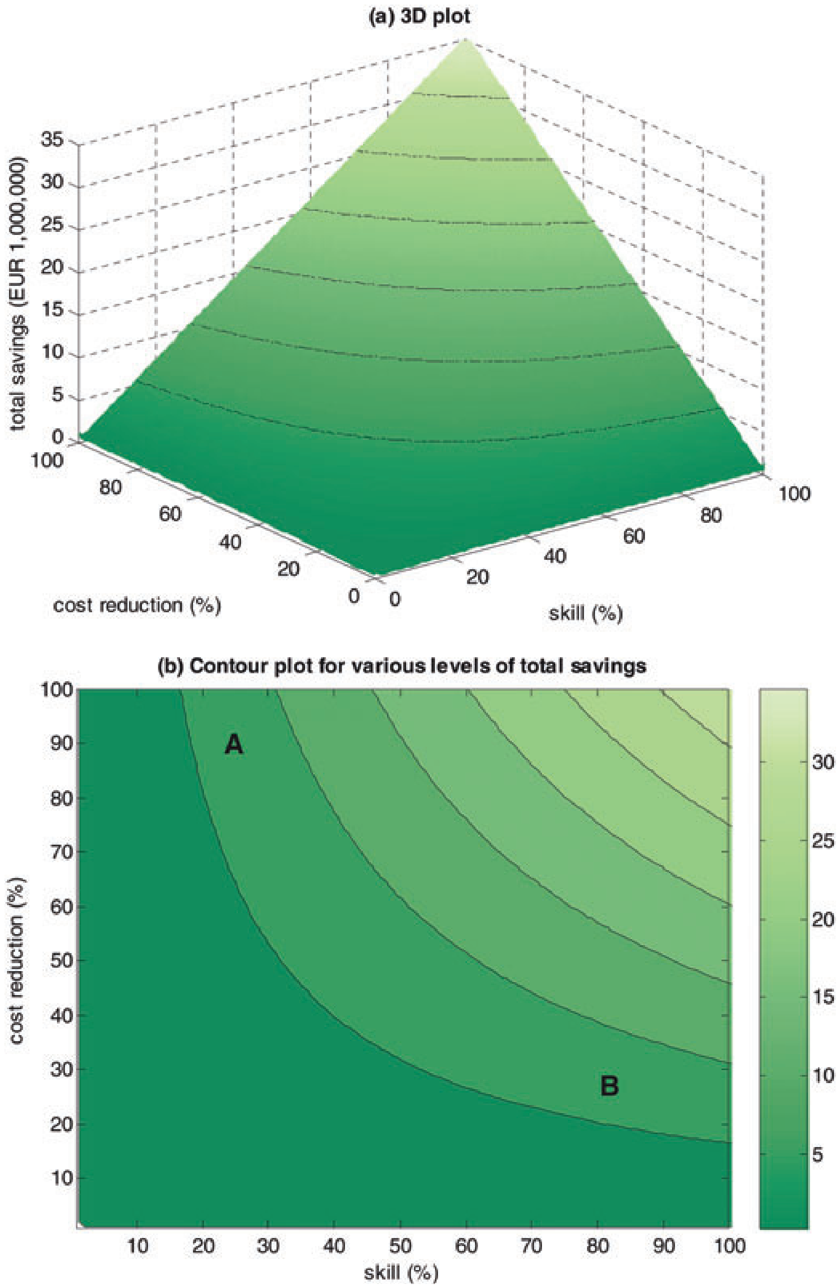

Figure 2. Relation between investor skills, cost reduction, and savings. (a) Total (expected) savings on market impact costs including commission (in millions of euros) for each percentage of investor skills and cost reduction per trade. (b) Corresponding contour plot for various levels of total savings

DOI: 10.1002/for

to be combined with a much smaller improvement once these are already high. Also, when investor skills are low and the percentage of cost reduction is high (point A), relatively more cost reduction than skill improvement is needed to arrive at lower market impact costs. Similarly, with high investor skills and low cost reduction (point B), relatively more skill improvement than cost reduction is required to reduce market impact costs. Also, the contour lines in Figure 2(b) show that relatively low skill values in combination with a substantial cost reduction per trade yield the same savings on total trading costs as comparatively high skill values and low cost reduction percentages.

FORECASTING MARKET IMPACT COSTS

Market impact costs usually depend on various trade, exchange, and stock specific characteristics. To formalize this, we assume that the market impact costs of a buy are determined by *N* factors (say, $\mathbf{X}_1, \ldots, \mathbf{X}_N$) and a random noise term ε. Since trading costs of buys and sells usually show different behavior, we follow the literature and consider separate models for them. Without loss of generality, we confine ourselves here and in the sequel to buys. We deal in exactly the same way with sells, using similar notation. Thus, we assume that market impact costs (exclusive of commission) of buys ($\mathbf{C}^{\mathrm{B}}$) satisfy

$$\mathbf{C}^{\mathrm{B}} = \beta_0^{\mathrm{B}} + \sum_{j=1}^{N} \beta_j^{\mathrm{B}} X_j + \varepsilon^{\mathrm{B}},\ \mathrm{E}\left(\varepsilon^{\mathrm{B}} \,|\, \mathbf{X}_1, \ldots, \mathbf{X}_N\right) = 0 \tag{2}$$

For the *N* factors we initially consider the variables listed in Table I.

The model in equation (2) explains market impact costs of buys and sells from various, to some extent controllable, factors. In this section we go one step further and use the model to forecast future market impact costs. Forecasts of market impact costs can be used to identify expensive trades before actual trading takes place. When forecast trading costs exceed a certain critical level, the investor may decide to change the type of broker intermediation, trade timing or trading venue, or to monitor the trade more carefully during execution. Alternatively, forecasts of trading costs can be combined with forecasts of returns to obtain estimates of future 'net' returns.

Clearly, forecasts need to be sufficiently 'accurate' to contribute to effective trading cost management. When forecast market impact costs appear to be underestimated, the respective trade may not have been given the additional monitoring that could have avoided (part of) these high trading costs. Hence, inaccurate forecasts may lead to a costly 'missed detection'. On the other hand, when forecast market impact costs appear to be overestimated, the trade may have been given unnecessary attention during the trading process, resulting in wasted costs. Hence, inaccurate forecasts can also result in a costly 'false alarm'. Therefore, good forecasts strike a balance between missed detection and false alarm. Obviously, the optimal balance between these two rates will depend on individual investor preferences.

Throughout, we evaluate the forecasting power of the model in equation (2) both in-sample and out-of-sample. To do so, we divide the data sample into an in-sample part (the first 2 months of trade, about 75% of the total sample) and an out-of-sample period (the final month of trade, corresponding to 25% of the sample). Subsequently, we use the in-sample data to perform a model selection procedure and to estimate an appropriate market impact cost model for both buys and sells. That is, we estimate the linear regression model in equation (2) using ordinary least squares. We apply the

DOI: 10.1002/for

Table II. Market impact costs and its determinants

Variable	Buys		Sells	
	Coeff.	*t*-value	Coeff.	*t*-value
intercept	–45.6	–1.8	–2.3	–0.1
momentumperc	**24.5**	10.6	**–21.4**	–9.4
tradesizertdv	**5.6**	2.7	**7.7**	2.7
marketcap	3.6	1.7	**–8.9**	–3.5
agencysingledum	**48.6**	6.3	–8.5	–0.8
quantdum	**–52.2**	–4.5	**56.6**	4.3
preopendum	**80.3**	6.7	**–71.7**	–3.9
morningdum	**81.1**	6.6	**–86.9**	–4.9
Wednesdaydum	**–62.1**	–9.0	**62.9**	6.2
Thursdaydum	**–89.2**	–8.1	**87.6**	5.9
Fridaydum	**–117.7**	–11.3	111.1	8.6
earlymonthdum	16.3	1.4	**46.2**	3.1
Febdum	**46.2**	5.9	**–40.2**	–4.2
NYSEdum	**–56.1**	–6.6	**66.8**	6.6
Nasdaqdum	**–101.4**	–7.7	**140.0**	10.4
Adj. R^2	0.17		0.23	

This table displays a list of determinants of market impact costs and their estimated coefficients (together with White's (1980) heteroskedasticity-robust t-values) based on the model of equation (2) for buys. The same model was estimated for sells. Coefficients in bold face are significant at the 5% level.

Akaike criterion to delete any redundant explanatory variables from the regression model. Table II displays the estimated coefficients of the final model, including *t*-values and adjusted R^2s. For a detailed interpretation of the model coefficients, we refer to Bikker *et al.* (2007). Given the final model specification, the in-sample forecasts correspond to the forecast trading costs for the in-sample trades. To obtain out-of-sample forecasts, we use an expanding window. That is, we estimate the model using all trades up to rebalancing k. Subsequently, we forecast trading costs for all trades during rebalancing $k + 1$. We do this sequentially for all rebalancings in the out-of-sample period.

Numeric forecasts of market impact costs

There are several ways to forecast future market impact costs. For the moment we assume that the investor's goal is to obtain pre-trade cost estimates for every trade to be executed. The usual way to do this, is by means of expected market impact costs. Since the noise term in specification (2) has mean zero, expected trading costs of buys equal $E(\mathbf{C}^B \mid \mathbf{X}_1, \ldots, \mathbf{X}_N) = \beta_0^B + \Sigma_{j=1}^N \beta_j^B \mathbf{X}_j$. Given estimates of the coefficients β_j^B based on historical data, we easily calculate expected trading costs given factors $\mathbf{X}_1, \ldots, \mathbf{X}_N$.

Figure 3 displays scatter diagrams of realized and forecast trading costs, together with regression lines that capture the relation between forecasts and realizations. Clearly, predicted trading costs differ considerably from realized costs. This is confirmed by formal error measures, such as *Theil's inequality coefficient*, which takes the value one in the no-skill case and equals zero with perfect foresight. In-sample it takes the value 0.61 for buys and 0.58 for sells. Out-of-sample it equals 0.67 and 0.68, respectively. Several other error measures are displayed in Table III. Similar to Theil's inequality measure, the mean absolute relative error is a scale-independent measure for the forecast error which should be as close to zero as possible. The mean squared error is another measure for

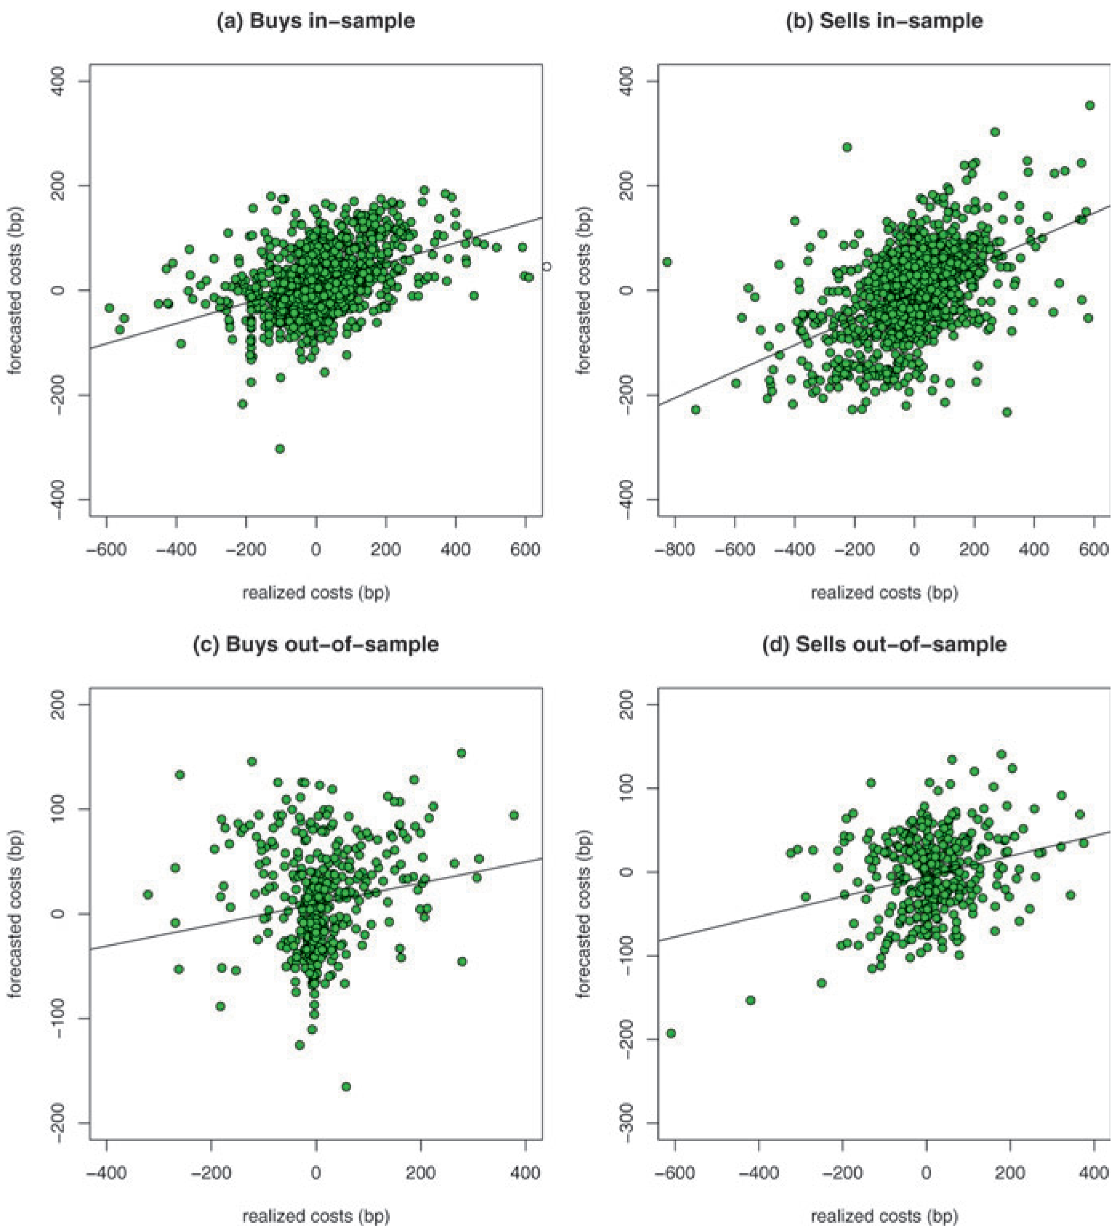

Figure 3. Realized and forecast market impact costs This figure displays realized and forecast market impact costs for buys and sells (both in-sample and out-of-sample), together with regression lines that express forecast trading costs as a function of realized costs

DOI: 10.1002/for

Table III. Quality measures for in-sample and out-of-sample forecasts

	Buys		Sells	
	In-sample	Out-of-sample	In-sample	Out-of-sample
Theil's *U*	0.61	0.67	0.57	0.68
Mean absolute percentage error	4.1	8.6	3.2	3.5
Mean squared error	12,876	9,283	17,398	11,425
Bias part (%)	0.0	0.0	0.0	2.5
Variance part (%)	39.0	21.6	33.1	31.8
Covariance part (%)	61.1	78.6	66.9	66.0
Naive hit ratio (%)	59.2	52.7	50.5	59.1
Hit ratio (%)	67.4	57.8	66.3	51.6
Correlation to realized costs	0.44	0.19	0.50	0.27

This table displays various quality measures for the in-sample and out-of-sample forecasts of market impact costs based on the standard forecast method derived from the regression model.

the prediction error, but it depends on the scale of the data. Therefore, we also report its decomposition into bias, variance, and covariance percentages which sum up to 100%. The bias percentage tells how far the mean of the forecast is from the mean of the actual series, whereas the variance percentage measures the variation of the forecast relative to the variation of the actual costs. The covariance percentage measures the remaining unsystematic forecasting errors. Ideally, the bias and variance proportions should be small, so that most of the discrepancy between forecast and realized market impact costs is idiosyncratic. The hit ratio counts the percentage of forecasts with the correct sign. Table III also displays the 'naive' hit ratio (obtained by assigning each trade to the most likely category). Finally, Table III reports the correlation between the forecast and realized trading costs, which should ideally be as close to one as possible. For more information on the error measures and their definitions, we refer to the Appendix.

As expected, the performance of the forecasts is generally better in-sample than out-of-sample. Table III shows that both the in-sample and out-of-sample forecasts have a small bias proportion, but a considerable variance percentage. However, the variance proportion is still much smaller than the covariance percentage. Hence, the forecasts succeed well in capturing the mean of the actual market impact costs, but are less successful in capturing the variation of these costs. Only for sells in the out-of-sample period is the hit ratio of the quantile model lower than the naive hit ratio. The correlations between forecast and realized trading costs reflect to what extent the model indeed forecasts higher trading costs for stocks that actually incur high costs. Both in-sample and out-of-sample, the correlation is significantly positive, reflecting a modest positive relation between forecast and realized trading costs.

Since market impact costs reflect the price movements of a stock during trade execution, the difficulty of forecasting these costs does not come as a complete surprise. Moreover, the out-of-sample month differs substantially from the in-sample period, which also complicates forecasting. That is, January 2002 was bearish and February was quite flat, but the out-of-sample month of March was bullish. Nevertheless, even for the turbulent out-of-sample month the upward slopes of the regression lines in Figure 3 reflect fairly positive correlations between realized and forecast trading costs. In-sample the correlations equal 0.44 for buys and 0.50 for sells. Out-of-sample they take values 0.19 and 0.27, respectively. All correlations are significant at a 5% significance level.

DOI: 10.1002/for

Bucket classification approach

Although the forecast quality of the specification in equation (2) is limited, the model at least succeeds in forecasting higher trading costs for stocks that actually experienced higher costs of trading.[7] Therefore, we expect to be more successful in classifying market impact costs in terms of 'high' or 'low', rather than forecasting exact numeric values.

To predict trading costs in terms of 'high' or 'low', we distinguish five buckets with predefined boundaries. We use the probability to exceed a certain level of market impact costs on a trade (i.e., $P(\mathbf{C}^B > T \mid \mathbf{X}_1, \ldots, \mathbf{X}_N)$) to predict the bucket in which market impact costs will fall. The higher the probability that a trade will cause high trading costs, the higher the bucket we will predict for that trade. We take the same buckets for buys and sells and define them in such a way that we have five buckets with increasing levels of market impact costs: bucket 1 ('no costs', $(-\infty, 0]$ bp), bucket 2 ('low costs', $(0, 20]$ bp), bucket 3 ('average costs', $(20, 50]$ bp), bucket 4 ('high costs', $(50, 80]$ bp), bucket 5 ('severe costs', $(80, \infty)$ bp). Given $p_0 = 0$ and $p_5 = 1$, we set four critical 'cut-off probabilities' p_1, p_2, p_3, and p$_4$ to assign the trades to one of the five buckets. If the 'excess probability' satisfies $p_i < \mathrm{IP}(\mathbf{C}^B > T \mid \mathbf{X}_1, \ldots, \mathbf{X}_N) < p_i + 1$ for a certain critical level T, we predict[8] that a buy will fall in bucket $i + 1$.

We can easily calculate the excess probability corresponding to the regression model in equation (2), provided that we know the distribution of the error term. If we denote the distribution function of the noise term by $F^B(x) = P(\varepsilon^B \le x)$, the excess probability for buys according to model (2) can be written as

$$P\left(\mathrm{C}^B > T \mid \mathbf{X}_1, \ldots, \mathbf{X}_N\right) = 1 - F^B\left(T - \beta_0^B - \sum_{j=1}^{N} \beta_j^B \mathbf{X}_j\right) \quad (3)$$

The distribution of the noise term has to be known in advance to calculate this probability. As usual, the assumption of normality seems obvious and convenient, but is nevertheless likely to be restrictive. Therefore, we take the empirical distribution of the noise term based on the in-sample period, which avoids any parametric assumptions. This means that we calculate the excess probability as the fraction of trades in the in-sample period for which the residuals[9] exceed the value in the parentheses of $F^B(\cdot)$ in the excess probability.

In practice, the choice of the critical level and cut-off probabilities will depend on the investor's preference regarding the balance between the false alarm and the missed detection rate. Here we use the in-sample period to determine appropriate values of the critical level T and the cut-off probabilities $p_1, \ldots, p_4$. We set $T = 80$ and $p_1 = 0.10$, $p_2 = 0.175$, $p_3 = 0.275$, and $p_4 = 0.50$. For sells we proceed in a similar way and set $T = 80$ and $p_1 = 0.09$, $p_2 = 0.15$, $p_3 = 0.25$, and $p_4 = 0.40$.

Again we use an expanding estimator and evaluate the quality of the bucket forecasting approach. The upper panel of Table IV displays the classification results, both in absolute numbers and percentages. Ideally, the percentages on the diagonals of the second and fourth panel at the right-hand side of Table IV should be as close as possible to 100%. The higher they are, the more trades are classified in the correct buckets. Misclassification occurs when off-diagonal elements in Table IV

[7] We notice that a correlation between realized and forecast trading costs of $x\%$ corresponds to an investor skill of approximately $x\%$ as well.

[8] Alternatively, we could use linear discriminant analysis or an ordered probit/logit model for this classification problem. However, we obtain better results with the current method.

[9] The residuals are defined as $e^B = \mathbf{C}^B - \hat{\beta}_0^B - \Sigma_{j=1}^N \hat{\beta}_j^B \mathbf{X}_j$.

DOI: 10.1002/for

Table IV. Classification results based on bucket approach

	In-sample					Out-of-sample					
BUYS											
#	Realized					#	Realized				
Predicted	1	2	3	4	5	Predicted	1	2	3	4	5
1	204	31	32	8	22	1	70	21	10	4	4
2	185	50	61	36	30	2	75	31	11	5	6
3	103	68	69	37	68	3	41	18	10	9	15
4	87	42	60	37	75	4	29	6	9	7	23
5	35	10	19	15	121	5	26	6	5	2	14
Total	614	201	241	133	316	Total	241	82	45	27	62
%	Realized					%	Realized				
Predicted	1	2	3	4	5	Predicted	1	2	3	4	5
1	33.2	15.4	13.3	6.0	7.0	1	29.0	25.6	22.2	14.8	6.5
2	30.1	24.9	25.3	27.1	9.5	2	31.1	37.8	24.4	18.5	9.7
3	16.8	33.8	28.6	27.8	21.5	3	17.0	22.0	22.2	33.3	24.2
4	14.2	20.9	24.9	27.8	23.7	4	12.0	7.3	20.0	25.9	37.1
5	5.7	5.0	7.9	11.3	38.3	5	10.8	7.3	11.1	7.4	22.6
Total	100	100	100	100	100	Total	100	100	100	100	100
SELLS											
#	Realized					#	Realized				
Predicted	1	2	3	4	5	Predicted	1	2	3	4	5
1	198	11	8	7	12	1	31	5	6	5	3
2	188	49	24	21	36	2	33	20	14	9	14
3	175	64	36	18	46	3	43	21	18	20	20
4	81	33	31	26	67	4	42	7	9	10	22
5	52	12	36	21	121	5	9	5	1	7	12
Total	694	169	135	93	282	Total	158	58	48	51	71
%	Realized					%	Realized				
Predicted	1	2	3	4	5	Predicted	1	2	3	4	5
1	28.5	6.5	5.9	7.5	4.3	1	19.6	8.6	12.5	9.8	4.2
2	27.1	29.0	17.8	22.6	12.8	2	20.9	34.5	29.2	17.6	19.7
3	25.2	37.9	26.7	19.4	16.3	3	27.2	36.2	37.5	39.2	28.2
4	11.7	19.5	23.0	28.0	23.8	4	26.6	12.1	18.8	19.6	31.0
5	7.5	7.1	26.7	22.6	42.9	5	5.7	8.6	2.1	13.7	16.9
Total	100	100	100	100	100	Total	100	100	100	100	100

The above panel displays the classification results for the bucket approach. For both buys and sells it reports the amounts and percentages of trades with realized costs in bucket i and forecast costs in bucket j (i, j = 1, 2, 3, 4, 5).

	BUYS		SELLS		NAIVE
	In-sample	Out-of-sample	In-sample	Out-of-sample	
Correctly class.	32.0	28.9	31.3	23.6	20
Low predicted high	21.3	20.7	20.6	29.2	40
High predicted low	21.4	21.3	20.3	25.4	40
Seriously misclass.	21.4	20.9	20.5	27.8	32
>2 buckets misclass.	32.7	34.4	37.8	42.5	48

The above panel reports five 'overall' quality measures for the bucket classification method.

DOI: 10.1002/for

are not equal to zero. The lower panel of Table IV displays several measures related to the overall classification quality. We consider the percentage of (1) correctly classified trades, (2) trades with no or low market impact costs that are predicted to have high or severe trading costs, (3) trades with high or severe market impact costs that are predicted to have no or low trading costs, (4) seriously misclassified trades, which are defined as trades with no or low costs classified as high or severe or vice versa, and (5) trades misclassified two or more buckets away from the correct bucket. We compare the resulting percentages to the 'no-skill' or 'naive' model assigning a trade to bucket $i = 1, \ldots, 5$ with probability 1/5. For both buys and sells, the bucket approach strongly outperforms the no-skill method on all five criteria. In particular, the important category of trades with high or severe trading costs that are wrongly classified as having no or low costs is only around 20–25%, versus 40% in the no-skill model. In line with expectations, the performance of the bucket approach relative to the naive model is in-sample more convincing than out-of-sample. However, the out-of-sample performance is still very good.

Identifying expensive trades: probability method

Another way of dealing with future market impact costs is to identify trades that have a high chance of being (too) expensive. That is, we assume that a trade is identified as expensive when the excess probability exceeds a certain critical level; i.e., when $P(\mathbf{C}^B > T \mid \mathbf{X}_1, \ldots, \mathbf{X}_N) \geq p$, for certain investor-specific values of the critical level T and cut-off probability p. The difference between this 'probability method' and the previous two approaches is that we no longer forecast a level or range of market impact costs for each trade, but only identify those that are likely to be expensive. Again we use the empirical distribution to calculate the excess probability.

To assess how well the probability method discriminates between cheap and expensive trades, we set the critical level at $T = 50$. For cut-off probabilities $p = 0.025, 0.05, \ldots, 0.975, 1$, we plot the corresponding false alarm rates against the detection rates. The solid curves in Figure 4 highlight several relevant quantities for buys and sells, both in-sample and out-of-sample. The distances below (to the left of) the solid curves reflect the detection (false alarm) rate. Moreover, the distances above (to the right of) the curves correspond to the missed detection (correct no-detection) rate. For instance, $p = 0.25$ results in false alarm and detection rates of, respectively, 54% and 83% (in-sample) and 47% and 82% (out-of-sample) for buys. For sells these rates equal 43% and 75% (in-sample) and 46% and 64% (out-of-sample). In the case of perfect foresight, the detection rate equals 100% at each false alarm rate. The dashed 45° lines in Figure 4 reflect the no-skill forecasts that are obtained by randomly assigning p% of all trades to the group of trades with costs higher than 50 bp. For $0 \leq p \leq 1$, this results in false alarm and detection rates of p%. The 45° lines are obtained by considering the entire range of possible values for the cut-off probability p. From Figure 4 we see that at each level of the false alarm rate the probability method has a higher detection rate than the no-skill method.

The surface of the areas below each solid curve can be viewed as the proportion of correct forecasts across all possible thresholds. This surface ranges from 0 to 100%, where the value 50% corresponds to the no-skill value (which equals the surface of the area under the 45° line) and 100% to perfect foresight. It serves as a summary statistic of the discrimination power of the model. For buys, its value equals 76% for the in-sample period and 77% for the out-of-sample period. For sells, it takes the values 78% and 68%, respectively.[10] Hence the probability method substantially outperforms

[10] We have calculated the surfaces using the trapezoidal rule for numerical integration based on 40 points.

DOI: 10.1002/for

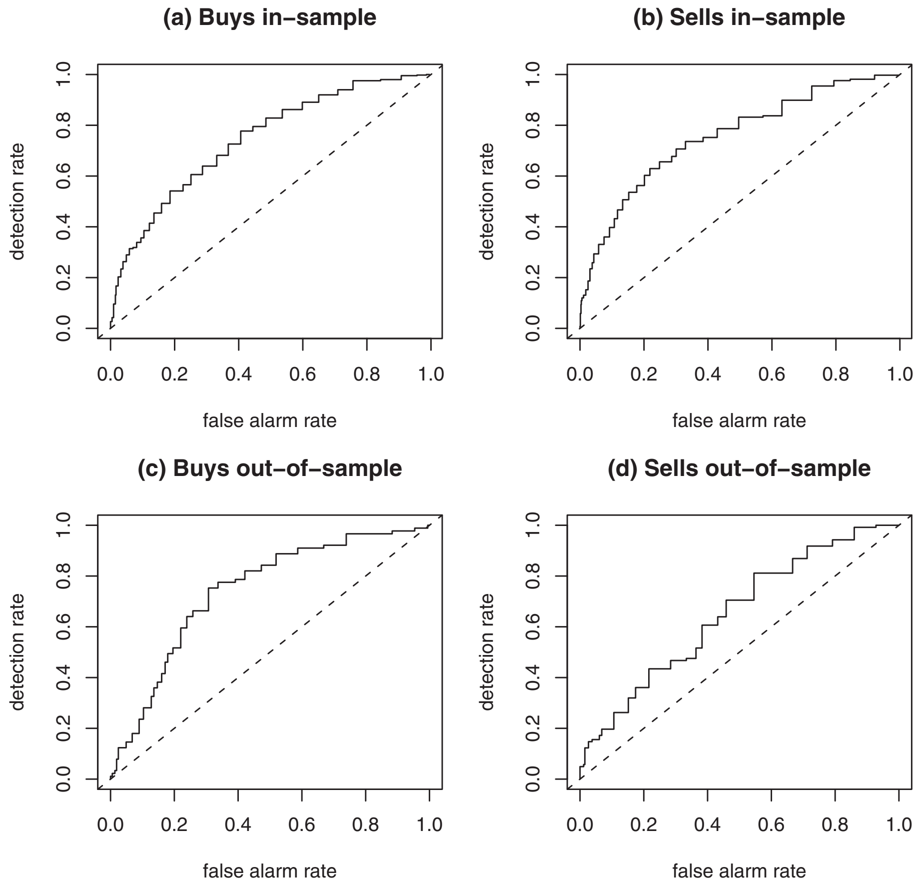

Figure 4. Discrimination ability of forecasts for various cut-off probabilities For cut-off probabilities $p = 0.025$, 0.05, . . . , 0.975, 1, the solid curves in this plot exhibit the false alarm rate against the detection rate for buys and sells, both in-sample and out-of-sample. The cut-off probability $p = 1$ corresponds to zero false alarm and detection rates, whereas these rates take the value one for $p = 0$. The distances below (to the left of) the solid curve reflect the detection (false alarm) rate. Moreover, the distances above (to the right of) the curve correspond to the missed detection (correct no-detection) rate. The dashed 45° lines reflect the no-skill forecasts that are obtained by randomly assigning p% of all trades to the group of trades with more than 50 bp. For $0 \leq p \leq 1$, this results in false alarm and detection rates of p%. The 45° lines are obtained by considering the entire range of possible values for p.

the naive forecasting approach across all values of the cut-off probability. Moreover, the results are robust across different levels of the critical level T.

In practice, the investor will choose a probability level on the basis of his or her preferences regarding the balance between the false alarm and the missed detection rate. However, the above analysis of forecast quality provides a convenient measure of discrimination that does not depend on arbitrary threshold levels.

DOI: 10.1002/for

INTEGRATING MARKET IMPACT COSTS IN THE EQUITY INVESTMENT PROCESS

Portfolio management is an economic decision problem dealing with the trade-off between expected returns, risk, and transaction costs. In this section we discuss how our forecasting model for trading costs as given in equation (2) can be employed for making better-informed decisions in portfolio management. We do not strive to discuss the above economic decision problem in detail here, as an excellent treatment can be found in Grinold and Kahn (1999). Here we limit ourselves to the concepts and focus on our contribution to a better specification of the cost function.

As advocated by Grinold and Kahn (1999), we opt for a full integration of forecasts of transaction costs, particularly market impact costs, in portfolio construction at the same stage of the optimization process where risk and return forecasts are used. In this way the components return, risk, and costs are optimized simultaneously. We note that the above costs not only refer to transaction costs but also to either self-imposed constraints or constraints that are imposed by the regulator. This integrated approach is especially relevant in case restrictions are added to the portfolio allocation problem. In such cases a sequential approach, rather than an integrated approach, could lead to less optimal portfolio allocations. For example, a long–short equity manager may carefully construct a market neutral portfolio by imposing zero exposures to certain style factors. If subsequently certain stocks are allocated different weights on the basis of the transaction cost analysis, the final portfolio may be far from market-neutral.

In short, the trade-off between return, risk, and costs can be formalized by expressing the following investor utility function:[11]

$$U = w'\mu - (\lambda/2)w'\Sigma w - \kappa c(w) \tag{4}$$

Here w are the portfolio weights, μ are the expected returns, Σ is the covariance matrix of the returns and $c(\cdot)$ is the cost function. Furthermore, $\lambda > 0$ and $\kappa > 0$ are, respectively, the risk aversion and the cost aversion parameter.

The risk–return trade-off problem has been well understood since the landmark analysis by Markowitz (1952) on mean–variance efficiency and Sharpe's (1964) CAPM model. The incorporation of costs in the development of optimal strategies complicates the analysis and is still an underdeveloped area in financial economics. We refer to Grinold (2006) for a recent attempt of integrating costs in the theory of portfolio management. The costs function $c(\cdot)$ is usually specified as a function of w only. For example, Keim and Madhaven (1996) find that market impact costs increase with three-halves power of the size of the trade, i.e., $c(w) \propto w^{3/2}$. Grinold (2006) works with $c(w) \propto w^2$. The contribution of our paper is that we come up with a broader specification of $c(\cdot)$. That is, our function $c(\cdot)$ is not only a function of the weights w but also of the risk factors $\mathbf{X}_1, \ldots, \mathbf{X}_N$ in our model in equation (2). We note that some of these variables, such as volatility and momentum, are exogenous to the decision maker. Other determinants, such as the day of the week and the agency–principal dummy, are under the control of the decision maker, so these in turn can be integrated in the utility framework in equation (4) above.

Estimation errors may dominate portfolio construction in practice, particularly when the uncertainty about the market impact cost estimate is much larger than the mean. More accurate specification of the cost function $c(\cdot)$ in equation (4) will improve portfolio optimization. Additionally, the

[11] We follow the usual mean–variance framework in our specification of the expected utility function.

DOI: 10.1002/for

market impact cost model can be used to identify which factors (e.g., percentage of average daily volume) attribute most to market impact. Such insights can be used to control the active bets in a similar way that other investment constraints do.

The cost forecasts can also be used to identify potentially expensive trades in terms of market impact costs. Such forecasts have a signaling function in the trade-monitoring phase. A different trading strategy can be adopted for trades that are likely to turn out expensive.

CONCLUSIONS

When a relatively small group of trades causes the major part of the market impact costs of an investment portfolio, a reduction of the trading costs of comparatively few expensive trades would already result in substantial savings on total trading costs. For the global equity trades analyzed in this paper, executed by the world's second largest pension fund ABP in the first quarter of 2002, we find that only 10% of the trades causes 75% of total trading costs. Simulations emphasize that there is a strong nonlinear trade-off between trading costs and the number of trades executed.

Since trading costs depend to some extent on controllable factors such as broker intermediation, trade timing, and trading venue, investors may reduce the costs of trading by carefully controlling these factors. As a first step in this direction, this paper has proposed two methods to identify potentially expensive trades. The bucket classification method predicts the category trading costs will fall into. The probability approach detects expensive trades on the basis of the probability that market impact costs will exceed a particular level. Applied to the equity trades executed by ABP, the proposed methods succeed in identifying a considerable number of expensive trades and substantially outperform no-skill prediction methods. However, we emphasize that the forecasting power of any model at a certain moment in time does not necessarily carry over to the future. Continuous updating of forecasting models, both in terms of functional form and model coefficients, will be necessary to adapt to changes in the market environment.

The results of this paper illustrate the productive role that model-based forecasts can play in trading cost management. This role can be extended further when the cost forecasts are integrated with the risk–return optimization process of portfolio allocations.

ACKNOWLEDGEMENTS

The authors would like to thank the participants of the research seminars at De Nederlandsche Bank, University of Twente, and Delft University of Technology. We are also grateful to the participants of the 9th Conference of the Swiss Society for Financial Market Research in Zürich and the Mathematical PhD Day in Utrecht. In particular, we would like thank Alain Durré for his useful feedback. The usual disclaimer applies. The views expressed in this paper are not necessarily shared by DNB and ABP or its subsidiaries.

REFERENCES

Almgren R, Thum C, Hauptmann E, Li H. 2005. Equity market impact. *Latin Risk* September: 21–28.

Bikker JA, Spierdijk L, Van der Sluis PJ. 2007. Market impact costs of institutional equity trades. *Journal of International Money and Finance* **26**: 974–1000.

DOI: 10.1002/for

Chan LKC, Lakonishok J. 1995. The behavior of stock prices around institutional trades. *Journal of Finance* **50**: 1147–1174.

Chan LKC, Lakonishok J. 1997. Institutional equity trading costs: NYSE versus Nasdaq. *Journal of Finance* **52**: 713–735.

Cheng M. 2003. Pretrade cost analysis and management of implementation shortfall. *AIMR Conference Proceedings* July: 26–34.

Grinold RC. 2006. Implementation efficiency. *Investment Insights 4.06*. Barclays Global Investors: San Francisco, CA.

Grinold RC, Kahn RN. 1999. *Active Portfolio Management: A Quantitative Approach for Producing Superior Returns and Selecting Superior Returns and Controlling Risk*. McGraw-Hill: New York.

Keim DB, Madhavan A. 1996. The upstairs market for large-block transactions: analysis and measurement of price effects. *Review of Financial Studies* **9**: 1–36.

Keim DB, Madhavan A. 1998. The cost of institutional equity trades. *Financial Analysts Journal* **54**: 50–69.

Konstance MS. 2003. Trading cost analysis and management. *AIMR Conference Proceedings* July: 8–12.

Markowitz HM. 1952. Portfolio selection. *Journal of Finance* **7**: 77–91.

Sharpe WF. 1964. Capital asset prices: a theory of market equilibrium under conditions of risk. *Journal of Finance* **19**: 425–442.

Wagner WH, Edwards M. 1993. Best execution. *Financial Analysts Journal* **49**: 65–71.

White H. 1980. Heteroskedasticity consistent covariance matrix estimator and a direct test for heteroskedasticity, *Econometrica* **48**: 817–838.

APPENDIX: FORECAST ERROR MEASURES

Given a sample of observations $C_1, \ldots, C_n$ and corresponding forecasts $\hat{C}_1, \ldots, \hat{C}_N$ the mean absolute percentage error (MAPE) is defined as

$$\text{MAPE} = \frac{1}{n}\sum_{i=1}^{n}\left|\frac{\hat{C}_i - C_i}{C_i}\right| \tag{A.1}$$

Furthermore, the mean squared error (MSE) is calculated as

$$\text{MSE} = \frac{1}{n}\sum_{i=1}^{n}\left(\hat{C}_i - C_i\right)^2 \tag{A.2}$$

The bias, variance, and covariance proportions of the MSE are given by

$$\text{BP} = \frac{\left(\bar{\hat{C}} - \bar{C}\right)}{\sum_{i=1}^{n}\left(\hat{C}_i - C_i\right)^2 \Big/ n}, \text{VP} = \frac{\left(s_{\hat{C}} - s_C\right)^2}{\sum_{i=1}^{n}\left(\hat{C}_i - C_i\right)^2 \Big/ n}, \text{CP} = \frac{2(1-\hat{\rho})\, s_{\hat{C}} s_C}{\sum_{i=1}^{n}\left(\hat{C}_i - C_i\right)^2 \Big/ n}$$

where $\bar{C}, \bar{\hat{C}}$, s_C, $s_{\hat{C}}$ are the sample means and variances of $C_1, \ldots, C_n$ and $\hat{C}_1, \ldots, \hat{C}_n$, respectively. The sample correlation between the series of observed values and forecasts is denoted by $\hat{\rho}$. Finally, Theil's inequality coefficient is obtained as

$$U = \frac{\sqrt{\frac{1}{n}\sum_{i=1}^{n}\left(\hat{C}_i - C_i\right)^2}}{\sqrt{\frac{1}{n}\sum_{i=1}^{n}\hat{C}_i^2 + \frac{1}{n}\sqrt{\sum_{i=1}^{n}C_i^2}}} \tag{A.3}$$

DOI: 10.1002/for