Does Information in the Limit Order Book Help to Predict Returns?[1]

Geoffrey Tombeur[2] and Gunther Wuyts[3]

June 2011

[1]We would like to thank Christophe Majois, Hans Degryse and participants at the 3L workshop (Brussels) and the University of Leuven for their valuable comments on an earlier version. We gratefully acknowledge financial assistance from FWO-Flanders under contract G.0567.10.

[2]University of Leuven, Faculty of Business and Economics, Department of Accounting, Finance and Insurance, Naamsestraat 69, 3000 Leuven, Belgium. E-mail: *geoffrey.tombeur@econ.kuleuven.be.*

[3]University of Leuven, Faculty of Business and Economics, Department of Accounting, Finance and Insurance, Naamsestraat 69, 3000 Leuven, Belgium. E-mail: *gunther.wuyts@econ.kuleuven.be.*

Abstract

We investigate whether information in the public limit order book helps to predict short-term returns in an order-driven market and if so, what drives this predictability. First, when distinguishing between imbalances in buy and sell orders (flow measure) and imbalances in the shape of the limit order book at ask and bid side (a static measure), we find that imbalances in the shape of the book are more informative than flow imbalances. The main information is at the best prices, imbalances further down the book are less informative about future returns. Furthermore, order flow imbalances for market orders are more informative in predicting returns than imbalances for limit orders, but imbalances in price-improving limit orders or limit orders at the best prices do contain additional information on top of market orders. Secondly, we study what explains the return predictability. We find multiple evidence that informed trading is a key element.

JEL Codes: G10

Keywords: Limit Order Book, Order Flow Imbalances, Return Predictability, Informed Trading

1 Introduction

Nowadays order-driven markets such as Euronext, Xetra or the Tokyo Stock Exchange exhibit a large degree of transparency. Indeed, they all operate a public electronic limit order book (henceforth, 'the book') through which market participants are updated in real-time about the liquidity available in the market. They do not only observe the best bid and ask prices and their respective depths, but also prices and depths further away in the book (in full, or possibly up to a certain level). In addition, also recent behavior of participants is observable, typically at least recent trades can be seen. This high degree of transparency is the starting point and motivation for our study. We address the following question: to what extent does this large information set, provided by the book and past behavior, contain information about future movements in the price? In this way, we shed further light on a central question in finance, namely why do prices in financial markets move. The relation between return predictability, and information in the book and trading behavior, also has implications for market efficiency (see e.g. Chordia et al., 2005). The analysis of our research question is divided in two parts. First, we investigate whether public information in a limit order market helps to predict returns, and which part of the information set of traders does so. Secondly, we aim to shed light on the issue of what drives this predictability. They key question here is whether predictability is in part related to informed traders, who reveal part of their private information by their actions.

In a first part of the paper, we investigate whether short-term (i.e. five-minute) returns are predictable on basis of information that is available in a public limit order market. We explicitly make a distinction between two types of information. A first type is information on the shape (also called state) of the book. The shape of the book refers to the current level of subsequent bid and ask prices in the book, and their associated depths. The depth as well as the (sequence of) prices can be informative on future prices. A deeper book on the bid side might indicate a heavier buying interest. Likewise, if bid prices are closer to each other than ask prices, the more crowded the book is on the bid side, which again might indicate a heavier buying interest. A similar reasoning holds for the ask side. We will refer to the size of differences in ask (bid) prices as heights at the ask (bid) side. This first type of information follows the approach in Cao et al. (2009), Harris and Panchapagesan (2005) or Jain et al. (2011). We consider a broader information set than these contributions, however, by adding a second type of information: information on the flow of orders, which captures recent dynamics in order submissions and trades. In particular, we use imbalances between different categories of buy and sell orders.[1] Chordia and Subrahmanyam (2004) argue that order imbalances

[1] Orders are classified in different categories according to their aggressiveness, see Section 3.1.

have a “natural appeal” as a determinant of returns. Most studies that have examined the relation between order imbalances and returns, however, only consider trades (which result from market orders). A distinct feature of markets that operate a limit order book is that traders face the choice between submitting either a market order or a limit order. Therefore, it seems also intuitively “natural” to include limit order imbalances in the analysis as well. Our paper is thus the first to include both static information (the shape of the book) and flow information (order imbalances) from the limit order book, and to study whether both have predictive power for returns.

We empirically test for predictability of returns over five minute intervals using an extension of the methodological framework provided by Cao et al. (2009). We address predictability by considering the adjusted R^2 of a regression of returns on a set of independent variables that capture various aspects of order flow and the shape of the book. We use data from the Spanish Stock Exchange in the analysis. Our main results can be summarized as follows. When replicating the Cao et al. (2009) approach, we find that their conclusion still holds: the shape of the order book matters in explaining stock returns after controlling for lagged returns, imbalances in trades and liquidity (the spread). Most information is concentrated at the first steps in the book; information from steps further down the book becomes less and less relevant. Second, when extending the Cao et al. (2009) framework by adding information on the flow of orders, i.e. order imbalances, using a more refined order classification scheme based on order aggressiveness, we find that predictability increases, indicating that there is some information in the flow of (limit) orders not captured by the static shape of the book. To summarize, both shape of the book and order imbalances have predictive power for returns.

Having established that both flow and stock variables related to the book are relevant for predicting returns, we turn in a second part of the paper to the question of what drives this predictability, an issue that - to the best of our knowledge - has not been directly (empirically) addressed in the literature so far in the context of a limit order market. While Chordia and Subrahmanyam (2004) highlight the role of inventories of intermediaries, this reasoning does not hold in an order-driven market since inventory issues are typically assumed not the be an issue in such market type (see e.g. de Jong et al., 1996; or other papers considering spread decompositions, typically, the inventory component of the spread is not considered). One rationale for predictability of returns on the basis of book imbalances can be found in the endogeneity of order choice. The choice between a market order and a limit order, or more generally, between a more aggressive or a less aggressive order, for any trader is dependent on a trader’s patience and the shape of the limit order book (see e.g. Parlour, 1998 or Foucault et al., 2005). More impatient traders use more aggressive orders. Traders also use more aggressive orders if they find the book to be more crowded on their side and less crowded on the opposite side. The reason is that 'crowdedness' affects the limit order's non-execution risk. A buyer, for instance, is tempted to use more aggressive orders if he observes a crowded bid side or a less crowded ask side. A seller on the other hand will use less aggressive orders in this case. This may result in an overall 'switch' to more market buy orders and more (and less aggressive) limit sell orders. In the end, this may shift the prices upward. The same reasoning can be applied for downward price shifts. These patterns can cause short-term predictability even when all traders are symmetrically informed.

The question of what causes these book imbalances, however, remains, especially so if we consider the possibility of asymmetrically informed traders. Are imbalances only the result of random patterns in order arrivals or are they the result of informed traders submitting limit orders on one side of the book? If imbalances are purely random and the fundamental value has not changed, the resulting price change should only be temporarily. If it is the result of informed traders submitting limit orders on one side of the book, price changes should be permanent: informed traders drive the price to its fundamental value. This paper tries to empirically evaluate to what extent return predictability on the basis of book imbalances can be related to informed traders' behavior.

We argue that informed trading is one of the key drivers. First, we find evidence that the predictive ability of our models increases over longer horizons. This increase is larger for models that use more information from deeper down the book. This longer-run impact suggests that the imbalances in the shape are at least partly the result of informed trading. Secondly, when looking at intraday patterns of predictability we find that predictability is highest in the beginning of the trading day and then decreases during the trading day. If informed trading is more likely at the start of the trading day, when information asymmetries are larger (see e.g. Admati and Pfleiderer, 1988 or Garvay and Wu, 2009), this points to the fact that informed traders do indeed submit limit orders and that this an important explanation of the predictability of returns. Finally, we study in a cross-sectional regression framework which stock characteristics can explain return predictability by shape and order imbalances. We find that predictability is larger for stocks for which there is more informed trading and this relationship is stronger for prediction models that use more information from deeper down the book. Our result that informed trading explains (at least part of) the relation between returns, and shape of the book and order imbalances, is new and contrasts to the interpretation of Chordia et al. (2004) that inventory considerations are a key driver.

The remainder of the paper is organized as follows. Section 2 discusses previous research on predictability of short-term returns and provides two arguments for predictability based on the shape of the book. We also highlight the contributions of our analysis, compared to the literature. Section 3 provides a definition of our different variables and the methodology we use. The data set is described in Section 4. The results on predictability of returns using information in the book are presented in Section 5. Subsequently, Section 6 investigates the drivers of return predictability. Finally, Section 7 concludes.

2 Related Literature

Our paper primarily fits within the literature that aims to explain returns on basis of the shape of the limit order book or order flow. Predictability of returns based on imbalances in the flow of orders has been relatively extensively studied in the literature (see e.g. Chordia et al., 2002, 2005 and 2008; Chordia and Subrahmanyam, 2004; Lee et al., 2004). The study of Chordia et al. (2002) is the first to empirically investigate the relationship between returns and the imbalances in the flow of orders in stock markets. They examine daily market-wide order imbalances on the NYSE and returns on the S&P500. They find that contemporaneous order imbalances are positively related to returns, while lagged order imbalances are negatively related to returns. They interpret these results as price pressure caused by inventory imbalances, and inventory stabilization following price pressure. Chordia and Subrahmanyam (2004) investigate the time series relationship between daily order imbalances and returns on the individual stock level. They develop a model to show that price changes are linearly related to lagged order imbalances. A crucial assumption is again that market makers have inventory concerns. They also provide empirical evidence for this relation and find that even significant profits (although small) can be obtained by designing a trading strategy based on this relation. Chordia et al. (2005) investigate the same relation, also at the individual stock level, but for intraday returns. They take the relation between lagged order imbalances and returns to be a violation of market efficiency, and examine how long it takes for arbitrageurs to eliminate the price impact of these order imbalances. They find that the predictive ability of lagged order imbalances is strong at shorter horizons of e.g. five minutes, but disappears at longer horizons of more than thirty minutes. All these studies, however, concern the NYSE, which is a hybrid market, where inventory concerns of market makers affect the imbalance-return relationship, as argued. In response, Lee et al. (2004) investigate the relation between returns and order imbalances for the Taiwan Stock Exchange (TWSE), which is an order-driven market with no officially designated market makers. They find that lagged order imbalances are relatively poor predictors of daily returns, indicating that the exchange functions quite well in accommodating persistency in order imbalances.

A second and much smaller strand of the literature has considered limit order book shapes. Cao et al. (2009) are the first to present evidence that imbalances between the ask and bid side in the limit order book, observed by a trader on his screen, can help in predicting short-term (i.e. five-minute) returns. They find that for a sample of 100 Australian stocks, lagged imbalances in the shape of the book are significantly related to returns over five minute intervals, after controlling for lagged returns, trade imbalances and the spread. They use imbalances in depths and heights (see Section 3), and additionally combine depth and height information into a price impact measure. We adopt the same empirical framework for our study. Harris and Panchapagesan (2005) examine whether book information can predict returns on the New York Stock Exchange (NYSE) in order to evaluate whether specialists can benefit from book information. They summarize the book shape into two types of imbalance measures, a first based on depth and a second based on option pricing models, since limit orders may be characterized as valuable trading options (e.g. Copeland and Galai, 1983). Their results indicate that book information can significantly predict transaction price returns. In addition, these measures also predict quote revisions and the aggressiveness of liquidity provision by the specialist, suggesting that the latter uses this information to his benefit. Both Cao et al. (2009) and Harris and Panchapagesan (2005) justify the relationship between book information and returns by the nature of limit order traders. They follow Harris (1990) who distinguishes between two types of limit order traders. Pre-committed limit order traders use limit orders to reduce trading costs, but switch to market orders if their orders are unmatched, which can move the price. Value-motivated traders enter limit orders on the basis of their stock valuations. Their valuations are then impounded in the price as they are revealed to the market. Both types of limit orders are therefore informative of future prices. Jain et al. (2011) conduct a horse race among different liquidity measures, to test which of these measures has the highest predictive power for volatility, trade prices and the speed of trading on the Tokyo Stock Exchange. They find that cost-based measures of liquidity can predict future price changes better than elasticity-based measures of liquidity (slopes). Also other researchers have focused on the relationship between the shape of the book and volatility (e.g. Ahn. et al., 2001; Naes and Skeltorp, 2006; Foucault et al., 2007; Pascual and Veredas, 2010).

Our paper differs form and contributes to the above literature in a number of ways. First, we take into account information in both the shape of the limit order book and in imbalances in order flow. The literature so far has typically only considered one of the two, not both together. This allows us to compare the predictive power of a static (shape of the book) vs a flow (order imbalance) type of information available to traders. In addition, we make a distinction in order imbalance between different categories of orders. The classification of orders is based on their aggressiveness, using a scheme similar to that of Biais et al. (1995). Previous studies have focused on the predictive ability of order imbalances using only market orders. By contrast, we also take into account imbalances in limit orders as predictors of future returns since our focus is on the predictive ability of book information, and it is the flow of limit orders that finally results in book. Order imbalances might have incremental predictive power to shape imbalances as the former capture more recent trading behavior. Specifically, if informed traders have recently submitted limit orders, imbalances in the flow of limit orders might represent their private information. Thirdly, we are the first to empirically provide evidence for one possible explanation for the predictability of returns: informed traders. This explanation is also in direct contrast to the literature on order imbalances in intermediated (or hybrid) markets such as the NYSE where inventory management by market makers is considered.

Including both shape and order flow imbalances in our analysis is motivated by the literature on order-driven markets. We believe that there are two arguments that relate returns to the shape of the book, close to the rationale of pre-committed and value-motivated traders. The first argument is found in the endogenous nature of order choice and the resulting predictable order flow patterns. Models of order-driven markets typically characterize traders who submit limit orders as patient, whereas traders submitting market orders are characterized as impatient (or eager) (e.g. Glosten, 1994; Foucault et al., 2005; Van Achter, 2009). Parlour (1998) models the considerations a trader makes when deciding whether he should submit a market order or a limit order. A limit order obviously results in a better price, but it has the drawback of adding uncertainty, i.e. the risk of the order not being executed. The key determinant to consider for a trader is the probability of execution of his order. The evaluation of this trade-off depends on his patience. Parlour suggests that, due to time priority rules, the probability of execution decreases when there is already a large number of outstanding orders available at his side of the book. This is termed the 'crowding out effect'. By the same logic, when a large number of orders is available on the opposite side, chances are higher that the trader will submit a limit order, since the probability that traders on the opposite side submit market orders increases (and thereby the probability of execution of his limit order). Hence, when a trader with a buying interest enters the market and observes an excess demand, he is more likely to submit a market buy order. Sellers entering are more inclined to submit limit orders. If more future trades are buyer-initiated, the probability of a price rise (positive return) increases.

Empirical evidence for these patterns is provided by Biais et al. (1995), Ranaldo (2004) and Pascual and Veredas (2009). Ranaldo (2004) shows that a thicker (thinner) book on the same (opposite) side induces traders to submit more (less) aggressive orders. This is consistent with the evidence from Biais et al. (1995) who document traders undercutting the best prices when the depth at their side is large. The sensitivity of undercutting to the depth is related to the competition from traders on the same side.

Pascual and Veredas (2009) explicitly relate the shape of the book to order choice using a two-stage sequential ordered probit model, which allows to model impatient (market order) and patient (limit order) traders separately. They find that the frequency of patient traders increases with the spread and the thickness of the depth on the opposite side, while it decreases in the thickness of the depth on the same side. Patient traders base their order choice mostly on the shape of the book at their side, while impatient traders pay more attention to the depth at the opposite side. The aggressiveness of limit orders also decreases as the height on their side increases. Parlour (1998) notes that even in the absence of asymmetric information and with a random arrival of trader types, these trading patterns will emerge. However, any price effect that is purely due to order flow patterns should be small and short-lived only. If the fundamental value has not changed, arbitrageurs should drive the price back to its fundamental value after a certain period of time.

Our second argument to justify the return-imbalance relationship introduces asymmetric information. If information asymmetries between traders exist *and* informed traders submit limit orders, these limit orders should result in an imbalance in the shape of book (since they submit orders on one side only). It is interesting to consider then under which circumstances informed traders indeed submit limit orders. Glosten (1994), Rock (1996) and Handa et al. (2003) argue that informed traders do not submit limit orders if their private information is short-lived, because there is a large probability that their limit orders would remain unfilled. Harris (1998) conjectures that the use of market orders by informed traders depends on transaction costs on the one hand, and on the persistence of their informational advantage on the other hand. Kaniel and Liu (2006) demonstrate that, under certain conditions, informed traders are more likely to submit limit orders instead of market orders, even to an extent that limit orders convey more information than market orders. First, the higher the probability that the information is long-lived, the lower is the execution risk, favoring limit orders. Second, the larger the mispricing in the market, the higher are potential losses of non-execution, suggesting a greater favor towards market orders. Third, more uninformed traders increase the profitability of a limit order, but increase the profitability of a market order even more.

Through an experimental design, Bloomfield et al. (2005) find evidence that informed traders use both market and limit orders. They document a heavier use of market orders by informed traders at the beginning of the trading period, since mispricing is the largest then, and potential profits are greatest, consistent with the evidence from Kaniel and Liu (2006). As prices are moving closer to their true market values, however, informed traders use more limit orders, because the additional expected profits from market orders decrease. A justification for this can be found in the argument that a liquidity motivated trader faces an adverse selection problem. Limit orders suffer from the risk of being 'picked off' (Copeland and Galai, 1983; Foucault, 1999; Handa et al., 2003). For an informed trader, if he were to submit a limit order, this risk is virtually non-existent (almost by definition), while the execution risk is also much smaller (since he knows the direction of future moves in the prices). Similarly, his benefits are larger than if he were to submit a market order (he gains at least the bid-ask spread). Bloomfield et al. (2005) conjecture that this makes informed traders natural liquidity suppliers. Beber and Caglio (2005) find that informed agents act strategically and that their order submission strategies are partly dependent on market conditions. In circumstances where the probability of information based trading is high, they submit limit orders at a price further away in the book in order to hide their information. Hence, even the more distant steps may contain information.

If an imbalance in book shape is due to informed trading on one side of the book, the 'informed imbalance' should signal a currently mispriced stock, and the price should go up or down in the near future. This suggests a stronger and more long-run price effect (until new information hits the market) than in the case of price movements due to only impatience-motivated trading patterns. This reasoning does not exclusively apply to imbalances in the shape of the book; a similar reasoning can be applied to order imbalances. They may contain information on the fundamental value of stocks, as well as drive temporary price changes. Liu and Seasholes (2011) study order imbalances in relation to prices for dual-listed shares and distinguish between the efficient price of a stock and its transitory price. They suggest that some order imbalances can contain fundamental information, while other order imbalances can push prices away from their fundamental values and thereby affect a stock's transitory price. Do note that both explanations as to why our hypothesis of predictable returns might hold are not mutually exclusive. Rather, the crowding-out effect may reinforce the effect of an informational imbalance. It is therefore difficult to empirically distinguish whether predictability is related to a crowding-out effect when there are also information asymmetries in the market. When we examine the rationale for return predictability based on book information we focus on the relation between predictability and informed limit order submissions.

3 Empirical Methodology

3.1 Measuring Limit Order Book Imbalances

We distinguish between two types of book imbalances: imbalances that summarize the flow of orders during a particular time interval (flow related) and imbalances that summarize the shape of the limit order book at a particular point in time (a stock related measure). The first imbalance measures concern the shape of the limit order book, which can be easily represented visually (hence the term 'shape'). Figure 1 gives an example of the book shape using the five best limits in the book for a hypothetical stock. Bid and ask prices are shown on the vertical axis, while depths (offered and demanded) at these prices are shown on the horizontal axis. These curves could be interpreted intuitively as estimates of the supply and demand curves for a stock, where the bid side curve represents an estimate for the demand curve and the ask side curve represents an estimate for the supply curve. The curves move up and down stepwise. Therefore we will refer to every consecutive price (and depth) as a 'step' further away in the book. For instance, the first best prices and depths represent the first step in the book, while the second best prices and depths represent the second step, and so on. We summarize the shape of the book into two simple measures for each step i separately: depths $\left\{D_i^{bid}, D_i^{ask}\right\}$ and heights $\left\{H_i^{bid}, H_i^{ask}\right\}$. Heights are defined as the size of the price differences (absolute value) between the subsequent steps in the limit order book. For instance H_2^{ask} is the (absolute value of the) difference between the second and first ask price. Note that these shape measures are not cumulative: $\left\{D_i^{bid}, D_i^{ask}\right\}$ and $\left\{H_i^{bid}, H_i^{ask}\right\}$ measure depth and height *at* step i, not *up to* step i. The benefit of using these variables is that they are very intuitive to interpret.

Please insert Figure 1 around here.

We measure imbalances in the shape of the book for each step in the book separately. We use information from the first until the fifth step in the book, $i = \{1, ..., 5\}$. The scaled imbalance in depth $Dimb_{i,t}$ at each step i at time t is our second variable of interest:

$$Dimb_{i,t} = \frac{D_{i,t}^{bid} - D_{i,t}^{ask}}{D_{i,t}^{bid} + D_{i,t}^{ask}} \qquad \text{for } i = \{1, ..., 5\}. \tag{1}$$

Our third variable of interest measures the imbalance in the heights at time t at step i in the book. The scaled imbalance in height $Himb_{i,t}$ between step i and step $i-1$ at time t is defined as:

$$Himb_{i,t} = \frac{H_{i,t}^{ask} - H_{i,t}^{bid}}{H_{i,t}^{ask} + H_{i,t}^{bid}} \qquad \text{for } i = \{2, ..., 5\}. \tag{2}$$

Note that defining this variable for the first step is useless, since the only way to define a height for the best quotes is to use the midquote as the quote on the previous step, resulting in a balanced height at the first step by convention.

Our second variable of interest is the imbalance in the flow of orders. Imbalances in market orders are known to be highly autocorrelated and generate contemporaneous price pressure. Lagged order imbalances are therefore potential predictors of short-term returns. We add to the literature by also taking into account imbalances in limit orders as predictors of future returns. Using only information on market orders and treating each market order the same ignores the fact that traders, especially in a limit order market, have the opportunity to choose among different levels of aggressiveness for their orders. Intuitively, more aggressive orders are those that have a higher probability of execution and care less about a favorable execution price. The most aggressive orders consume a great deal of the liquidity available on the opposite side of the book while less aggressive orders add liquidity on their side of the book. Market orders or trades only reveal information on the flow of more aggressive orders. The flow of limit orders may contain additional information not captured by the trades, notably when informed traders submit limit orders. In addition, traders may condition their trading strategies on other traders' recent trading behavior. Therefore we look at all orders submitted by traders and take into account the choices made by traders by classifying orders into different categories of aggressiveness to investigate which types of orders are more informative on future price movements. Table 1 describes the five categories of orders that we distinguish. Our classification scheme resembles the classic scheme of Biais et al. (1995) and is very similar to that of Beber and Caglio (2005).

Please insert Table 1 around here.

We then calculate the scaled imbalance between the number of buy orders $B_{j,t}$ and the number of sell orders $S_{j,t}$ from order category j during a time interval $[t-1,t]$ as

$$Oimb_{j,t} = \frac{B_{j,t} - S_{j,t}}{B_{j,t} + S_{j,t}}. \tag{3}$$

where we refer to Subsection 3.3 for a definition of the interval. Order imbalances $Oimb_{j,t}$ are defined for all order categories of orders $j = \{1, ..., 5\}$ from Table 1. We also define order imbalances using only the trades (i.e. market orders), compiling orders from the first two categories. This measure is referred to as the trade imbalance $Trimb_t$, to clearly distinguish it from order imbalances $Oimb_{j,t}$ which can be defined for any category of orders. Most existing studies that investigate order imbalances are limited to trade imbalances and therefore do not make this distinction.

3.2 Empirical Models

We build our models on the empirical framework designed by Cao et al. (2009) by regressing returns on information that is available in the open limit order book. This information is captured by the list of imbalance measures constructed from the limit order book. We use raw returns r_t as our dependent variable, as opposed to Cao et al. (2009), who use innovations in returns. To control for autocorrelation in returns we do include five lagged returns in all our models. Our results remain virtually unchanged when innovations in returns are used. All models are estimated on a per stock base. Similar to Cao et al. (2009), we address predictability of returns by investigating the adjusted R^2 of our models. The higher it is, the better our estimates fit the actual observations: the more public information on the shape of the limit order book, or the flow of limit orders, helps in predicting. Note that by using adjusted R^2, we correct for the number of variables included in the model.

Our first set of models is similar to the base model introduced by Cao et al. (2009), which allows for a comparison with their results. It focuses primarily on the predictive ability of the order book shape. Each of the models from this set adds variables summarizing information available from an additional step k in the limit order book.

$$\begin{aligned} r_t &= \beta_0 + \beta_1 Spread_{t-1} + \beta_2 Trimb_{t-1} \\ &\quad + \sum_{i=1}^{k} \gamma_i Dimb_{i,t-1} + \sum_{i=2}^{k} \delta_i Himb_{i,t-1} + \sum_{s=1}^{5} a_s r_{t-s} + u_t. \end{aligned} \tag{4}$$

with r_t the (percentage) return on the midprice between time $t-1$ and t. The sampling frequency, i.e. the length of the interval $[t-1, t]$ is five minutes in the main part of our paper, see Subsection 3.3.

In the first model we test the predictive power of the imbalance in trades $Trimb_t$, measured over the interval $[t-1, t]$. This model is referred to as the *step 0* model since it uses only information on the flow of orders and no information on the shape from any of the steps of the book. We control for the relative quoted spread between the best bid and ask prices $Spread_t$ and add five lagged returns to the equation. Our next model is the *step 1* model (i.e. $k = 1$) which adds information available from the first step in the limit order book: the imbalance in depth offered at the best quotes $Dimb_{1,t}$. Every subsequent model (*step k* regression model) adds information available from a following step $k = \{2, ..., 5\}$ in the book: depth imbalances $Dimb_{k,t}$ and height imbalances $Himb_{k,t}$, measured at time t (the end of each interval). This results in a set of six models based on equation (4) that are estimated. We expect the predictive ability of each subsequent model to increase as more information from the book is used.

Next, we focus on the predictive ability of the order flow by extending the above set of models with more detailed information on the flow of all order submissions. Equation (4) is modified by replacing the trade imbalance by five different order imbalances $Oimb_j$ ($j = \{1, ..., 5\}$), representing the five categories of order aggressiveness from Table 1.

$$
\begin{aligned}
r_t &= \beta_0 + \beta_1 Spread_{t-1} + \sum_{j=1}^{5} \lambda_j Oimb_{j,t-1} \\
&+ \sum_{i=1}^{k} \gamma_n Dimb_{i,t-1} + \sum_{i=2}^{k} \delta_n Himb_{i,t-1} + \sum_{s=1}^{5} a_s r_{t-s} + u_t. \qquad (5)
\end{aligned}
$$

This results in five new models (*step 0 cat. j* model) replacing the *step 0* model, each adding information from a category of less aggressive orders. In the *step 1* to *step 5* models we use information on the flow of all orders as well as the shape of the order book. Our second set of models thus comprises of ten models.

The main difference between order imbalances and shape imbalances is that the former are based on flow variables and the latter on stock variables. Order imbalances should therefore capture a more dynamic aspect of the order book by looking at recent order submissions. By estimating our *step 1* to *step 5* models, we examine whether recent order submissions have any predictive power on top of the information already contained in the shape of the book.

3.3 Sampling Frequency

In the regressions above, we need to make a choice about the frequency at which these regressions are estimated; i.e. a choice about the length of the interval $[t-1, t]$. Choosing the sampling frequency very high, e.g. a trade-by-trade basis, would imply a lot of microstructure noise. On the other hand, selecting a too low frequency, e.g. hourly, may cause loss of relevant information. Following Cao et al. (2009), we choose *five-minute* time intervals which strikes a balance between these arguments. During each five-minute interval we record the number of buy orders and the number of sell orders of each category to compute $Oimb_{j,t}$. To compute $Dimb_{i,t}$ and $Himb_{i,t}$, we measure the imbalances in the shape of the limit order book at the end of each interval t. Since we include five lags of the return in our models, we lose the first five observations of each day.

The choice of the interval length matters since predictability is likely to differ for different interval lengths. Chordia et al. (2005) find that market order imbalances can predict returns over smaller time intervals of less than thirty minutes, but as the interval length increases predictability diminishes. A comparison of model estimates using different interval lengths also sheds more light on the question of whether imbalances in the limit order book are caused by informed traders submitting limit orders. Therefore, and as a robustness check to verify that our results are not driven by the choice of the sampling frequency, we also investigate other sampling frequencies.

4 Data

4.1 Institutional Setting

This study uses intraday data of stocks that are listed on the Spanish Stock Exchange (SSE), which is a purely order-driven market. Its trading activity is managed through the electronic trading platform Spanish Stock Market Interconnection System (SIBE). Investors submit their orders through brokers who are provided with real-time information on trading activity and the shape of the limit-order book by SIBE's Dissemination Information System (DIS). Continuous trading takes place from 09.00 a.m. until 05.30 p.m. and call-auctions determine the opening and closing price. The minimum order size is one share and the tick is dependent on the trade price. It is €0.05 for stocks with a price above €50 and €0.01 for stocks with a lower price (all stocks in our sample). Three types of updates in the limit order book can be distinguished: new orders, order modifications or order cancellations. Three types of orders exist. Market orders are executed against the best prices on the opposite side of the book and walk up the book until they are completely filled. Market-to-limit orders are like market orders, but do not walk down the book when the depth at the first best price is completely used. Instead, they are stored at that price as a regular limit order. Limit orders are recorded in the book at their limit price and can only be executed at that price or better. Priority of execution is time-based. The unmatched limit orders summarize into the shape of the book. The dataset contains the shape of the book from the first until the fifth step, except for the invisible part of the depth from iceberg orders. We are able to determine whether a trade involved hidden depth, though. Given that we want to study how the public limit order book, as observed by traders, helps to predict returns, we do not include hidden depth in our analysis.

4.2 Sample Description

Our sample contains all stocks from the Spanish stock exchange that were part of the Open Market (Mercado Open) and the New Market (Nuevo Mercado) from January until December 2003. Stocks from the fixing markets are not considered because they do not trade in a continuous limit order book. Stocks in the Latibex-segment, which contains Latin-American stocks that are cross-listed in Spain, are discarded as well. Four stocks are deleted from this sample because they never shown an imbalance in height at any step (heights are always one cent). Our final sample consists of 95 stocks. The sample period covers all trading days for the stock during 2003, which is 250 days for most stocks.

We combine two datafiles in our analysis. One file contains data on the limit order book and shows all updates of the first five steps in the book for each stock in the sample, time-stamped to the nearest hundredth of a second. Every update contains the five best prices at both sides of the book and their respective depths. A second file contains all trades executed during the continuous trading session. The trading data show price and size of each trade and are time-stamped to the nearest second. Preopening or postclosing orders are not included since the trading mechanism during this period is different from the one during the trading day. Both book updates and trades are indexed. The index numbers and time stamps allow for a perfect matching of trade and order book data. By comparing each update with the previous update (the previous shape of the book) and related trades, all updates are classified, either into orders from one of the five the categories from Table 1, or as a modification or cancellation.

Table 2 displays some descriptive statistics for the sample. The cross-sectional mean, standard deviation, minimum and maximum for a selection of stock characteristics is shown. In the main part of the paper, we use five-minute intervals, which given the opening hours of the Spanish Stock Exchange leads to 102 observations per day (or 25,500 observations in total for most stocks). Table 2 shows the average five-minute return (in percentage points) and its standard deviation. As could be expected, the standard deviation is quite large relative to the mean return for such a short period of time. From the hourly share volume statistics, we observe a large variation across the stocks in the sample with a skewness to the left. The median stock has an hourly volume of about 9,000 shares (€ 82,000) while the most active stock trades over 3 million shares (€ 32 million) an hour. For the least active stock this is only 70 shares (970 euros). The mean (median) average trade size is 600 (460) shares (or € 5,300 (€ 3,920)), but with a range from 30 shares (€ 460) for the stock with the smallest trades up to 2,620 shares (€ 20, 540) for the stock with the largest trades. The large difference in trading activity is again highlighted when looking at the average number of trades: the mean (median) stock trades about 36 (7) times each hour, with a minimum of one trade every four hours and a maximum of almost 450 trades each hour. The mean (median) market value of the equity for our stock sample is about € 3.5 billion (€ 1 billion), with a range from € 7 million up to € 50 billion.

The following figures from Table 2 show some insights into the liquidity and order book shape of our sample. The mean (median) stock has a spread of about nine (four) cents (or 0.72% (0.52%) of the midquote), with a minimum of almost one cent (0.11% of the midquote) and a maximum of more than €1.4 (2.90% of the midquote), which leaves much room for temporary price improvements by aggressive orders. The same image holds when looking at the average total height between the fifth step bid and ask prices. The mean (median) stock has an average total depth (bid *and* ask) on the first step of about 13,000 (3,000) shares (or € 85,000 (€ 32,000)). There is a huge dispersion, as the cross-sectional standard deviation is about twice the size of the mean and nearly ten times the size of the median. The minimum of the total depth at the first step is an average of 450 shares (€ 6,620), the maximum about 150,000 shares (€ 1.2 million). The total cumulative depth on the first five steps shows more or less the same picture.

Please insert Table 2 around here.

5 Return Predictability and Limit Order Book Information

This section presents the results for return predictability based on information in the limit order book. In the first subsection we focus on the relation between the lagged shape of the order book and returns, by implementing the approach of Cao et al. (2009) on our data sample of Spanish stocks. This could be considered as an out-of-sample test for their findings. We show that imbalances in the shape of the order book help to predict intraday returns. Next, we show that using more detailed information on the lagged order flow improves the predictability of returns. Past order flow information is not used by Cao et al. (2009), but is in general also public information for traders and can thus potentially be used by them to predict returns. Recall that we address predictability by investigating the adjusted R^2 of the regressions specified in Subsection 3.2.

5.1 Book Shape Imbalances

Estimation results of our set of models of Equation (4) are presented in Table 3. Each model adds information available from a step deeper in the limit order book. The independent variable is the return during a five minute interval. All models are estimated on a per stock basis. The second and third line show the cross-sectional mean and median of the adjusted R^2 of each model. We also test whether the adjusted R^2s from each *step k* model are significantly different from the *step k-1* models using the non-parametric Wilcoxon signed rank test. The results in Table 3 show that the main conclusion from Cao et al. (2009) still holds for our sample. Only including the trade imbalance results in a quite low predictability with an adjusted R^2 of only 1.15%. By adding the imbalance in depth on the first step the adjusted R^2 jumps to 2.65%. Adding information from the second step increases the adjusted R^2 even further, to 3.56%. After that the adjusted R^2 only improves slightly when information from the following steps is added. This is the same pattern of improvements as reported by Cao et al. (2009). Predictability in the Spanish stock market is lower, however, compared to the results of Cao et al. (2009) for Australia. The cross-sectional mean adjusted R^2 of each of our regressions range from

1.15% to 3.75%, while the adjusted R^2s from Cao et al. (2009) for the same models range from 0.56% to 7.21%. This indicates that on average returns are less predictable for our sample of 95 Spanish stocks than for their sample of 100 Australian stocks. This raises some questions as to what drives differences in predictability. The difference might be related to different stock characteristics, since their time period is not very different (although only one month, March 2000) and the market microstructure of the both exchanges is also not fundamentally different. We come back to evidence on drivers of predictability based on cross-sectional differences in Section 6.

Please insert Table 3 around here.

From the fourth line on, the average coefficient estimates across stocks are shown. Below each estimate we include the percentage of stocks for which the coefficient is significant at the 5% level. In general results from Table 3 show that the intercept is not significantly different from zero. Notice that the coefficient estimates in the different models for the same variables are quite close to each other. The trade imbalance $Trimb$, depth imbalances at the first and the second step $Dimb_1$ and $Dimb_2$ and the height imbalance at the second and third step $Himb_2$ and $Himb_3$ are the most significant. As can be judged, from the order of magnitude of the coefficients , as well as the percentage of stocks for which coefficients are significant, both depth and height imbalances become less and less significant for steps further down the book. This confirms the picture from the adjusted R^2 - our measure of predictability. Orders further down the limit order book contain few relevant information, with regard to the direction and size of future short-term price movements, that is not included in orders at the first steps.

5.2 Order Flow Imbalances

Using more information on the submission of orders in the limit order book leads to a modified version our models, as presented in Equation (5). The results of these models are shown in Table 4. Again, models are estimated on a stock by stock basis. To conserve space, results of the *step 1 to 4* models are not tabulated, but show a consistent pattern. Judging by the mean and median adjusted R^2 across stocks, predictability increases gradually for each subsequent model. The adjusted R^2's and are also higher compared to their counterparts from the models of Equation (4). The main difference between imbalances in the shape of the book and imbalances in the flow of limit orders is that the latter capture recent dynamics in order submissions, while this is not necessarily true for the shape, which offers by definition a static picture of the submitted limit orders. These results suggest that the recent flow of limit orders contains information on future prices that is not captured by the shape of the limit order book.

Please insert Table 4 around here.

Compared to Equation (4), the information on the imbalance of trades or market orders ($Trimb$) is now decomposed into two imbalance measures: the imbalance in market orders that consume more liquidity than available at the first step (aggressive market orders, $Oimb_1$), and the imbalance in market orders that consume at most the liquidity available on the first step of the book (nonaggressive market orders, $Oimb_2$). The effect of an imbalance in the most aggressive market orders is much stronger than that of less aggressive market orders. When a positive imbalance in the number of aggressive market orders exists (i.e., more aggressive buy market orders), this has a stronger effect on prices than a positive imbalance of nonaggressive market orders of the same size. Intuitively, when order imbalances are the result of informed trading, aggressive market orders may signal a relatively larger mispricing, which causes them to have a stronger effect on prices. Even if the relationship between lagged order imbalances and returns is purely driven by order splitting, aggressive market orders may be more strongly related to future prices because they indicate more aggressive market orders to come in the near future due to order splitting. A liquidity trader who splits an order, but at the same time takes more liquidity than available at the first step, is likely to have a very large order to fill, which may cause further price movements after subsequent trades.

When looking at imbalances in limit order ($Oimb_3$, $Oimb_4$ and $Oimb_5$) orders that improve the depth at the first step in the limit order book (this is $Oimb_4$) seem to have the largest impact. However, when controlling for shape imbalances the effect of limit orders that improve the depth at the first step decreases and the effect of the more aggressive limit orders, which improve the price at the first step, becomes the strongest. The effect of the least aggressive orders ($Oimb_5$) is small and mostly insignificant. When informed traders submit limit order, more aggressive limit orders could signal that their information is more short-lived, making price changes in the near future more likely. Additionally, even when no information is involved, an imbalance in more aggressive limit orders may indicate a more fierce competition on one side of the market by impatient traders, which may drive temporary future price changes.

Finally, note also that most of the estimates of the coefficients of the $Dimb_j$ and $Himb_j$ variables remain close to their values from Table 3.

5.3 Robustness

In the empirical models above (Section 3.2), we used returns as dependent variables and added lags in the regression specifications. Another approach, closer to Cao et al. (2009), is to use innovations instead. Return innovations are obtained using an autoregressive model of order five $AR(5)$:

$$r_t = \alpha_0 + \alpha_1 r_{t-1} + \alpha_2 r_{t-2} + \alpha_3 r_{t-3} + \alpha_4 r_{t-4} + \alpha_5 r_{t-5} + \varepsilon_t^* \tag{6}$$

where r_t is the return for a stock in a period $[t, t-1]$, calculated from the midquote, the α_i are the estimated coefficients and ε_t^* is the innovation in return. Innovations in return are then regressed on information that is available in the open limit order book, just as before. For instance, Equation 4 becomes then:

$$\varepsilon_t^* = \beta_0 + \beta_1 Spread_{t-1} + \beta_2 Trimb_{t-1} + \sum_{i=1}^{k} \gamma_i Dimb_{i,t-1} + \sum_{i=2}^{k} \delta_i Himb_{i,t-1} + \eta_t.$$

where obviously, we do no longer add lagged returns. Adjusting all our regressions in a similar way and redoing the analyses in the previous two subsections, does not change any of our results.

As another robustness check we look at the results of the 35 stocks who were part of the Spanish national blue chip index, the IBEX-35, during the sample period. Again, our key results remain valid. The detailed results are available upon request.

6 What Explains Return Predictability on the Basis of the Limit Order Book?

In this section, we aim to gain further insights in the determinants of return predictability. Our main question of interest is whether return predictability is due to informed trading in the limit order book. We proceed in a number of steps. In the first subsection, we discuss predictability at longer horizons. If informed traders use limit orders in their trading strategies, imbalances reflect information on the true value of the security. Returns should therefore also be predictable at longer horizons. In the second subsection intraday differences in predictability are considered. If predictability is larger at times when information asymmetries are larger, this is likely to be related to informed traders. Since their information is most valuable at times when information asymmetries are the largest, informed traders are most likely to submit orders during these times. Finally, the last subsection looks in detail at determinants of predictability in a cross-sectional regression framework.

6.1 Predictability at Longer Horizons

A first interesting question is whether returns are also predictable over longer horizons on the basis of the same kind of limit order book information. Chordia et al. (2005) find that lagged market order imbalances lose their predictive ability over longer horizons, as the coefficient estimate, t-statistic and adjusted R^2 converge to zero. They conjecture that autocorrelated market order imbalances cause the specialist to alter his prices because of inventory concerns. Arbitrageurs are able to estimate order imbalances and their influence, but it takes a few minutes. As a reaction, they engage in countervailing trades, which removes predictability of returns. Chordia et al. (2005) argue that these countervailing trades make the market efficient after a period of thirty minutes or less. An implication of their finding that lagged order imbalances are not informative of future prices at longer horizons, is that informed traders do not engage in order splitting, at least with regard to market orders. Market order imbalances (or trade imbalances, using our terminology) are not 'informed' in the sense that they do not carry any valuable information on future prices.

In order to evaluate whether our imbalance measures are 'informed', we estimate our models from Equation (5) for longer time horizons: ten, fifteen and thirty minutes. Results are presented in Table 5. Panel A shows the cross-sectional mean and median adjusted R^2 of the models at these different sampling frequencies. Each row in the table corresponds to a model from Equation (5). We use the Wilcoxon signed rank test to examine whether the median difference between the adjusted R^2 of the model for a given sampling frequency and the smaller sampling frequency (from the previous column) is significantly different from zero. In general, the pattern shown is one of increasing adjusted R^2s for longer time horizons. The mean adjusted R^2 is increasing for almost each model, for each horizon. Up to fifteen minutes, the median adjusted R^2s increases with the time horizon, indicating that predictability increases with the interval length. For intervals of thirty minutes, the median adjusted R^2 drops again for the *step 0* and *step 1* models, while it increases further for the *step 2* to *step 5* models (although the differences are not significant). Generally, the pattern of increasing adjusted R^2s at longer horizons is clearer and more consistent for the models that use more information from steps further down the book. Predictability on the basis of information comprised in the shape of the limit order book is more persistent. This finding indicates that imbalances in the shape of the limit order book are, at least partially, due to informed trading. Some caution should be warranted, however, when considering the differences in these adjusted R^2s, because they are not entirely comparable. The units of measurement of returns are dependent on the interval length over which returns are measured.

Please insert Table 5 around here.

Panel B of Table 5 shows the coefficient estimates of the *step 5* model applied to different sampling frequencies (to conserve space, we do not report the other models). Below the estimates, the percentage of stocks for which the coefficients is significant at the 5% level, is shown. This percentage drops for almost every coefficient as the length of the horizon increases. Before turning to conclusions, for a sampling frequency with longer time intervals, we remark that (1) the number of observations in each sample decreases with a factor 1/2, 1/3 and 1/6 respectively, reducing the statistical power of the analysis, (2) due to excluding the first five observations for each day we lose more information from the more predictable start of the trading day (see subsection 6.2). To take these considerations into account we simply redo the analysis for a five-minute sampling frequency, but using the same numbers of observations and exclude the same time interval at the start of the trading day[2]. The estimation results of the models using these smaller samples are shown in columns 4, 6 and 8 of Panel B.

The differences in statistical significance of coefficients between the five minute sample, and the samples at longer horizons become much smaller now. Imbalances in aggressive market orders ($Oimb_1$), however, are economically, as well as statistically, clearly less significant for longer time horizons, consistent with the idea of order splitting of Chordia et al. (2005). The same finding holds for aggressive (price improving) limit orders. The coefficient of $Oimb_3$ becomes smaller and statistically less significant. Imbalances in aggressive limit orders are not primarily information driven, but may be caused by traders jumping the queue when the limit order book is crowded. For the other order imbalances, coefficients are increasing for longer horizons, but again they become statistically less significant. Overall, we conclude that the recent flow of orders helps to predict returns better, the shorter the time horizon. Predictability based on the recent order flow is less likely to be due to informed traders.

For most of the shape imbalance measures, the opposite pattern is observed: depth and height imbalances are more strongly related to future returns over longer horizons. All coefficient estimates are increasing with the horizon and are generally statistically more significant. Even limit order further down the book seem to contain more valuable information on future prices for longer horizons, which is consistent with informed traders submitting limit orders. If informed traders submit limit orders further away from the best quotes, this is likely to be more informative on prices at longer horizons (long-lived information). These findings suggest that predictability is related to information in the book. If predictability would be only a short-term phenomenon, it would be entirely the consequence of random trading patterns and impatient traders crowding each other out.

6.2 Time of Day Patterns in Predictability

A second important question, that can help to shed light on the rationale of return predictability, is whether returns are more predictable in one part of the day than another

[2]We reduce the sampling period to (approximately) the first six months, the first four months and the first second months, wile excluding the first ten, first fifteen and first thirty observations of each day respectively, such that we retain the same number of observations for each of the models compared to a ten, fifteen and thirty minute sampling frequency.

part. If predictability is higher at times when information asymmetries are larger, this might be an indication of informed limit order trading. Informed traders are more active when they can obtain the largest trading profits, i.e., when mispricing is the largest. In order to answer the question of intraday patterns in predictability we classify all observations for a stock into four smaller samples based on the time of the day: 09.30 a.m. until 11.30 a.m., 11.30 a.m. until 01.30 p.m., 01.30 p.m. until 03.30 p.m. and 03.30 p.m. until 05.30 p.m.. Our main criterion of a better predictability during one part of the day is a higher cross-sectional mean and median adjusted R^2, tabulated for the six models for these four subsamples in Table 6. Each row in the table corresponds to a model from Equation (5). We use the Wilcoxon signed rank test to examine whether the median difference between the adjusted R^2 of the model for a given subsample and the first period subsample is significantly different from zero.

Please insert Table 6 around here.

For the *step 0 cat. 1* to *cat. 3* models, that contain only the most aggressive orders, the mean and median adjusted R^2 seem to be the highest at the end of the trading day and the lowest at the beginning of the trading day. The median difference between the first and last subperiod is even significantly different from zero for the *cat. 1* and *cat. 2* models. For the *step 0 cat. 4* and *cat. 5* models, that contain also less aggressive (limit) orders, the mean is higher at the start of the trading day and the pattern of the median adjusted R^2 is close to a U-shape: the adjusted R^2 is higher at the start and end of the trading day. For the *step 1* to *step 5* models, the mean adjusted R^2 shows a near U-shape (or an 'inverse J-shape' in fact), with the a higher adjusted R^2 at the start of the trading day, decreasing in the second and third subperiod and slightly increasing again in the last subperiod. The median adjusted R^2 is gradually decreasing over the course of the trading day. The signed rank test indicates that the median difference between the first subperiod and later subperiods is significantly different from zero for both afternoon periods and even already significant at the ten percent level in the second subperiod, for the *step 4* and *step 5* models. Overall, our models that contain only information on the recent flow of market orders can predict returns better at the end of the trading day, while the fit of our models that contain also information on the shape of the order book, is significantly better at the start of the trading day.

One explanation for the better fit of the models based on market order imbalances may be the behavior of traders with a deadline approaching. If, for instance, large institutional traders need to trade before the end of the trading day, persistency of (market) order imbalances may become more prominent toward the end of the trading day.

The better fit of the models based on shape imbalances at the start of the trading day can be related to patterns in information asymmetries throughout the trading day.

Information asymmetries decrease during the trading day since all market participants learn about the fundamental value of the stock through the trading process. In addition, informed traders prefer to hide among liquidity traders, which concentrate their trades at the open and the close. This causes increased levels of informed trading at the open and close of the trading day. That is why in most financial markets U-shaped patterns of liquidity and trading activity are observed over the course of the day (Admati and Pfleiderer, 1988; Garvey and Wu, 2009). A similar reasoning can be applied to our finding of a decreasing pattern of predictability. If informed traders are more active when their informational advantage is larger *and* they reside to limit orders as a part of their trading strategy, the book should be more informative at the start of the trading day. Our findings are consistent with informed traders submitting information to the limit order book. This is in line with the results of Beber and Caglio (2005), who find that informed traders use less aggressive and more limit orders, in order to hide their information.

However, our finding contradicts Kaniel and Liu (2006) and Bloomfield et al. (2005) who find that informed traders use more market orders when their informational advantage is high, at the beginning of the trading day, and more limit orders when their informational advantage is low, toward the end of the trading day.

6.3 Determinants of Predictability

The two previous subsections have provided already a first idea on whether return predictability is at least partially driven by informed trading. In this subsection we investigate in detail for which stocks predictability is highest. In particular, we are most interested whether predictability is higher for stocks for which information asymmetries are larger and thus for which informed trading is more likely. To measure informed trading we apply the methodology from de Jong et al. (1996) to compute the adverse selection component of the spread in an order-drive market (see Appendix A for more detail on the estimation). To obtain reliable estimates, we need to ensure to have enough observations within each day. Therefore, we impose a restriction of at least 50 trades for each day before a trading day is included in the estimation sample. If more than half of the total number of trading days for a stock is deleted, we do not estimate our informed trading measure for this stock and delete the stock entirely from our sample for the subsequent analysis. Due to this requirement we retain only the 47 stocks with the highest number of daily trades. Do note that all results from a regression analysis with such a relatively low number of observations should be interpreted with some caution. Relaxing the criteria just specified to keep more stocks in the sample, does not alter our results, however.

Given the longer time period needed to estimate our measures of informed trading and predictability, we opt for a cross-sectional regression. More specifically, we regress $\overline{R}_i^2$, the adjusted R^2 of the different models above (see equation (5)) with subscript i referring to the regression of stock i, on a number of stock characteristics. The other stock characteristics that we include are the logarithm of the average hourly trading volume $Log(Volume_i)$, to control for the level of liquidity, and the standard deviation of the five-minute midquote return $\sigma_i(r_t)$,to control for the level of risk associated with the stock. The model to be estimated is then as follows (with u_i the error term):

$$\overline{R}_i^2 = \beta_0 + \beta_1 InfTrading_i + \beta_2 Log(Volume_i) + \beta_4 \sigma_i(r_t) + u_i \tag{7}$$

Including other control variables, such as the average bid-ask spread or market capitalization, does not change our results. We do not include all variables together, however, because of relatively high correlation between them.

Results are presented in Table 7. Each row in the table corresponds to a different model from equation (5). Below the coefficient estimates we show t-statistics in parentheses. Coefficients which are significant at the 10%, 5% level and 1% level are indicated with ***, ** and *, respectively. The last column reports the adjusted R^2 of the regression.[3]

Please insert Table 7 around here.

Predictability shows a positive relationship with the adverse selection component of the spread $InfTrading_i$ for all models, except the *step 0 cat. 1* model, which uses only information on the flow of the aggressive market orders. For our *step 0* prediction models the cross-sectional relationship between predictability and informed trading grows stronger for models that incorporate more information on less aggressive orders. This suggests that informed traders do submit less aggressive orders, which reflect their information. For our prediction models including flow as well as shape information (*step 1* to *step 5*) we observe the same pattern: predictability becomes more strongly related to the adverse selection component of the spread when more information from further down the book is included in the models. For stocks that experience relatively more informed trading, the shape of the book (especially when steps deeper in the book are included) helps to predict returns better, compared to stocks with relatively less informed trading. Overall, informed trading is positively related to predictability based on book information. It is remarkable though, that our informed trading measure becomes insignificant for the *step 1* and *step 2* models. Predictability based on information from the first and second best prices and depths, might be related more to the crowding-out effect than

[3]Note that Equation (7) is estimated using OLS. Alternatively, a censored model could be estimated, since theoretically the adjusted R^2 is bounded between a value slightly below zero and one. Estimating the model in this way does not change our results.

information submitted in the book.

Table 7 further shows that returns are less predictable for higher volume stocks. For stocks that experience more trading activity the market is more efficient: information gets impounded into prices more quickly, reducing predictability based on lagged book information. Volatility of the midquote is negatively related to predictability, but insignificant.

7 Conclusion

In this paper we investigate whether information provided on trading screens by the public limit order book can help to predict short-run returns. We find evidence that both the recent flow of incoming orders and the current shape of the limit order book are informative on future prices. For the flow of orders, we find that imbalances in market orders are more informative than imbalances in limit orders. Imbalances in the most aggressive limit orders, however, contain useful additional information on top of market orders. Shape imbalances are more informative than flow imbalances, but they become less informative if they summarize information from further down the book.

In a second part of the paper, we investigate what elements can explain this predictability. We show that for longer time horizons the adjusted R^2s grow larger, and this pattern is clearer and more consistent for the models that use more information from steps further down the book. This finding indicates that imbalances in the shape of the limit order book are, at least partially, due to informed trading. Practically the same pattern is observed when looking at the magnitude and significance of the coefficients. Shape imbalances are more strongly related to future returns over longer horizons. This suggests that limit orders submitted to the book contain information on the fundamental value of a security. The fact that predictability is larger at times when informed trading is larger, i.e. the start of the trading day, also points to this fact. Finally, our cross-sectional analysis directly shows that predictability is the largest for stocks with more asymmetric information. Overall, our results indicate that information on the fundamental value of securities slips into the book because informed traders submit limit orders, even at more distant steps. Traders can therefore learn on the value of a security by observing the shape of the book, and the recent flow of orders submitted to the book.

Appendix A: Spread Decomposition

The adverse selection component of the spread in a limit order market, which is used in the cross-sectional analysis from Section 6, is estimated using the de Jong et al. (1996) model (also used in e.g. Ahn et al., 2002). De Jong et al. (1996) estimate price effects of trading on the Paris Bourse using an extension from the Glosten (1994) model. Ahn et al. (2002) apply the same model on a set of stocks from the Tokyo Stock Exchange. The following equation is estimated:

$$\triangle p_{\tau+1} = c + R_0 \triangle Q_{\tau+1} + R_1 \triangle (q_{\tau+1} Q_{\tau+1}) + e_0 Q_\tau + e_1 q_\tau Q_\tau + u_{\tau+1} \tag{8}$$

with p_τ the logarithm of the transaction price at time τ, Q_τ the sign of the transaction (+1 for a buyer-initiated transaction, -1 for a seller-initiated transaction), q_τ the transaction size and $\triangle$ the difference operator. The time subscript τ refers to transaction time (as opposed to t which refers to 5-minute intervals). The model is estimated using OLS with Newey-West standard errors. De Jong et al. (1996) measure transaction size in multiples of Normal Market Size (NMS), the 'minimum marketable quantity', the minimum size for which market makers registered at London's SEAQ International are obliged to display bid and ask quotes (de Jong et al., 1995). Ahn et al. (2002) measure transaction size in multiples of the minimum trading unit. Since the minimum trading unit in our sample is one share, we measure transaction size in number of shares. The transaction size is censored at the 99.5 percentile. Overnight returns and opening prices are excluded. Combining the trade file with the limit order book updates file, we can classify each trade exactly as either buyer-initiated (a transaction at the ask side) or seller-initiated (a transaction at the bid side) (except for very rare occasions, which are excluded from the sample).

The spread in this model, $R_0 + R_1 q$, is dependent on the transaction size q. It can be divided in an order processing component equal to $c_0 + c_1 q$ and an adverse selection component, equal to $(R_0 - c_0) + (R_1 - c_1)q$. For our analysis we set q equal to the 99th percentile and use the adverse selection component relative to the spread.

References

Admati, A., P. Pfleiderer, 1988. A Theory of Intraday Patterns: Volume and Price Variability. *Review of Financial Studies*, 1, pp. 3-40.

Ahn, H-J., K-H. Bae and K. Chan. 2001. Limit orders, depth and volatility: evidence from the stock exchange of Hong Kong. *Journal of Finance* 56, pp. 769-790.

Ahn, H.-J., J. Cai, Y. Hamao, R. Ho. 2002. The Components of the Bid–Ask Spread in a Limit-Order Market: Evidence from the Tokyo Stock Exchange. *Journal of Empirical Finance*. 9. pp. 399-430.

Beber, A. and C. Caglio. 2005. Order submission strategies and information: Empirical evidence from the NYSE. *Working paper.*

Biais, B., P. Hillion and C. Spatt. 1995. An empirical analysis of the limit order book and the order flow in the Paris Bourse. *Journal of Finance* 50, pp. 1655-1689.

Bloomfield, R. M. O'Hara and G. Saar. 2005. The make or take decision in an electronic market: evidence on the evolution of liquidity. *Journal of Financial Economics* 75, pp. 165-199.

Cao, C., O. Hansch and X. Wang. 2009. The information content of an open limit-order book. *Journal of Futures Markets* 29, pp. 16-41.

Chordia, T., R. Roll and A. Subrahmanyam. 2002. Order imbalance, liquidity and market returns. *Journal of Financial Economics* 65, pp. 111-130.

Chordia, T., R. Roll and A. Subrahmanyam. 2005. Evidence on the speed of convergence to market efficiency. *Journal of Financial Economics* 76, pp. 271-292.

Chordia, T., R. Roll and A. Subrahmanyam. 2008. Liquidity and market efficiency. *Journal of Financial Economics* 87, pp. 249-268.

Chordia, T. and A. Subrahmanyam. 2004. Order imbalance and individual stock returns: Theory and evidence. *Journal of Financial Economics* 72, pp. 485-518.

Copeland, T. and D. Galai. 1983. Information effects on the bid-ask spread. *Journal of Finance* 38, pp. 1457-1469.

de Jong, F., T. Nijman and A. Roell. 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.

de Jong, F., T. Nijman and A. Roell. 1996. Price Effects of Trading and Components of the Bid-Ask Spread on the Paris Bourse, *Journal of Empirical Finance*, 3, pp. 193-213.

Foucault, T. 1999. Order flow and trading costs in a dynamic limit order market. *Journal of Financial Markets* 2, pp. 99-134.

Foucault, T., O. Kadan and E. Kandel. 2005. Limit order book as a market for liquidity. *Review of Financial Studies* 18, pp. 1171-1217.

Foucault, T., S. Moinas and E. Theissen. 2007. Does anonimity matter in electronic limit order markets? *Review of Financial Studies* 20, pp. 1707-1747.

Garvey, R., F. Wu, 2009, Intraday Time and Order Execution Quality Dimensions, *Journal of Financial Markets*, 12, pp. 203-228.

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

Handa, P., R. Schwartz and A. Tiwari. 2003. Quote setting and price formation in order-driven markets. *Journal of Financial Markets* 6, pp. 461–489.

Harris, L. 1990. Liquidity, Trading Rules and Electronic Trading Systems. *New York University Salomon Center Monograph Series in Finance and Economics*, Monograph 1990-4.

Harris, L. 1998. Optimal dynamic order submission strategies in some stylized trading problems. *Financial markets, institutions and instruments* 7, pp. 1-76.

Harris, L. and V. Panchapagesan. 2005. The information content of the limit order book: evidence from NYSE specialist trading decisions. *Journal of Financial Markets* 8, pp. 25-67.

Jain, P., P. Jain and T.H. McInish. 2011. The predictive power of limit order book for future volatility, trade price and speed of trading. *Working Paper.*

Kaniel, R. and H. Liu. 2006. So what orders do informed traders use? *Journal of business* 79: 84 - 110.

Lee, Y., Y. Yiu, R. Roll and A Subrahmanyam. 2004. Order imbalances and market efficiency: Evidence from the Taiwan Stock Exchange. *Journal of Financial and Quantitative Analysis* 39, pp. 327-349.

Liu, C. and M. Seasholes. 2011. Trading imbalances, dual-listed shares, and transitory prices. *Working Paper.*

Naes, R. and J.A. Skjeltorp. 2006. Order book characteristics and the volume-volatility relation: Empirical evidence from a limit order market. *Journal of Financial Markets* 9, pp. 408-432.

Parlour, C. 1998. Dynamics in limit order markets. *Review of Financial Studies* 11, pp. 789 – 816.

Pascual, R and D. Veredas. 2009. What pieces of limit order book information matter in explaining order choice by patient and impatient traders? *Quantitative Finance* 9, pp. 527-545.

Pascual, R and D. Veredas. 2010. Does the open limit order book matter in explaining informational volatility? *Journal of Financial Econometrics* 8, pp. 57-87.

Ranaldo, A. 2004. Order aggressiveness in limit order book markets. *Journal of Financial Markets* 7, pp. 53-74.

Rock, K. 1996. The specialist's order book and price anomalies. *Review of Financial Studies* 10, pp. 103-15.

Van Achter, M. 2009. A Dynamic Limit Order Market with Diversity in Trading Horizons, *Working Paper*.

Figure 1: Shape of the book for a hypothetical stock

Note: This figure presents the shape of the book for a hypothetical stock. Bid and ask prices are shown on the vertical axis and share depth is shown on the horizontal axis.

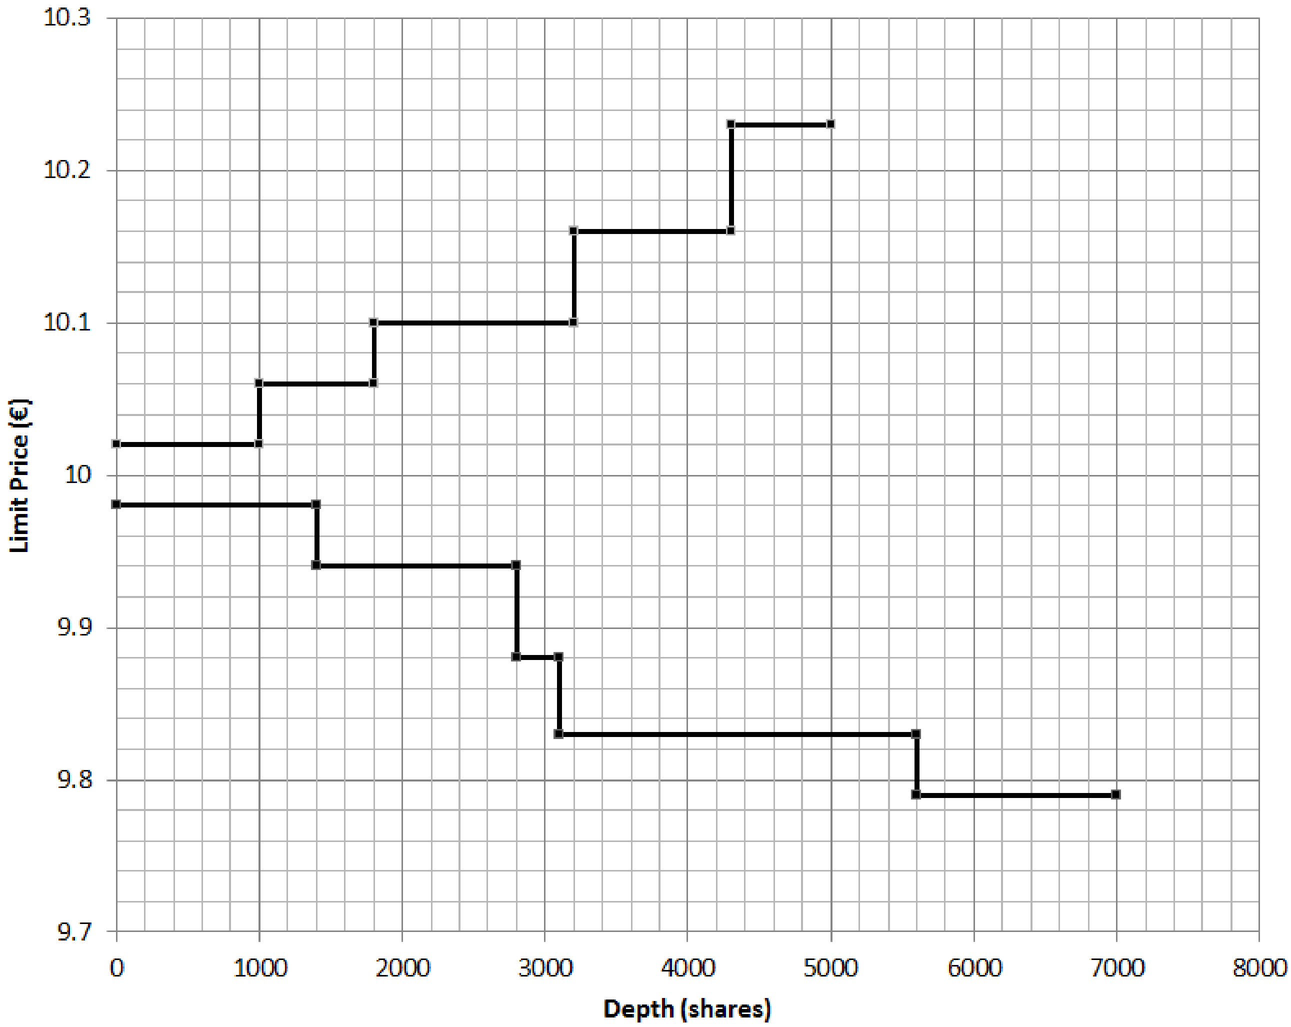

Table 1: Order Classification

Note: This table describes how order submissions are classified into five different categories on the basis of their aggressiveness.

Category	**Description of order type**
Category 1	Market orders that consume depth beyond the best quotes ('aggressive market orders')
Category 2	Market orders that consume at most the depth available on the best quotes ('nonaggressive market orders')
Category 3	Limit orders that improve the best quotes
Category 4	Limit orders that add depth to the best quotes
Category 5	Limit orders that add depth to quotes beyond the best quotes

Table 2: Descriptive Statistics

Note: This table presents descriptive statistics for each of the 95 stocks in our sample: the cross-sectional mean, median, standard deviation, minimum and maximum of selected variables. $\mu(r_t)$ is the mean five-minute midquote return (in percentage points), $\sigma(r_t)$ is the standard deviation of this variable, $Vol.$ (sh.) is the mean share volume per hour (in thousands of shares), $Vol.$ (EUR) is the mean euro volume per hour (in thousands of euros), $TradeSize$ (sh.) is the median trade size (in thousands of shares), $TradeSize$ (EUR) is the median trade size (in thousands of euros), $NrTrades$ is the mean number of trades per hour, MV (EUR) is the mean market capitalisation (in millions of euros), $Spread$ (cents) is the mean spread (in cents), $Spread$ (%) is the mean spread relative to the midquote (in percentage points), $Height_5$ (cents) is the mean total height from the fifth ask price to the fifth bid price (in cents), $Height_5$ (%) is the mean total height from the fifth ask price to the fifth bid price relative to the midquote (in percentage points), $Depth_1$ (sh.) is the mean total depth at the best bid and ask prices (in thousands of shares), $Depth_1$ (EUR) is the mean total monetized depth at the best bid and ask prices (in thousands of euros), $Depth_5$ (sh.) is the mean total cumulative depth from the first best until the fifth best bid and ask prices (in thousands of shares), $Depth_5$ (EUR) iis the mean total cumulative depth from the first best until the fifth best bid and ask prices (in thousands of euros).

	$\mu(r_t)$	$\sigma(r_t)$	$Vol.$ (sh.)	$Vol.$ (EUR)
Mean	0.0008	0.1628	122.64	1261.16
Median	0.0006	0.1511	9.32	82.74
St. Dev.	0.0015	0.0580	453.93	4216.95
Min.	-0.0033	0.0586	0.07	0.97
Max.	0.0047	0.4165	3205.26	32011.76
	$TradeSize$ (sh.)	$TradeSize$ (EUR)	$NrTrades$	MV (EUR)
Mean	0.60	5.30	35.93	3494.85
Median	0.46	3.92	6.99	990.00
St. Dev.	0.56	4.44	72.84	7814.86
Min.	0.03	0.46	0.25	7.00
Max.	2.62	20.54	444.61	50476.00
	$Spread$ (cents)	$Spread$ (%)	$Height_5$ (cents)	$Height_5$ (%)
Mean	9.12	0.72%	45.82	3.99%
Median	3.95	0.52%	20.54	2.62%
St. Dev.	17.70	0.63%	91.04	4.11%
Min.	1.07	0.11%	9.08	0.65%
Max.	140.34	2.90%	702.78	28.27%
	$Depth_1$ (sh.)	$Depth_1$ (EUR)	$Depth_5$ (sh.)	$Depth_5$ (EUR)
Mean	13.19	85.19	81.78	516.44
Median	3.06	31.99	19.28	193.51
St. Dev.	29.06	188.63	175.52	1187.43
Min.	0.45	6.62	2.83	45.27
Max.	149.24	1232.20	968.04	7982.73

Table 3: Estimation Results Order Book Shape Models

Note: This table presents estimates of the different models from equation (4). The independent variable is the return measured during a five minute interval (expressed in percentage points). Explanatory variables are defined in Subsection 3.1. The models are estimated on a per stock basis for the 95 stocks in our sample. White heteroskedasticity-consistent standard errors are used. The second and third line show the cross-sectional mean and median of the adjusted R^2 of each model. If the Wilcoxon signed rank test indicates that the median difference between the adjusted R^2 from a given model and the adjusted R^2 from the previous 'step' model is significantly different from zero at the 10%, 5% level and 1% level this is indicated with ***, ** and *, respectively. From the fourth line on, the cross-sectional average coefficient estimates are shown. Coefficients which are significant at the 10%, 5% level and 1% level are indicated with ***, ** and *, respectively. Below each estimate we include the percentage of stocks for which the coefficient is significant at the 5% level. In the bottom three lines we indicate whether a constant, the relative quoted spread and lagged returns are included in the models. Between parentheses we show the dominant sign of the coefficients for the different stocks ("+" for significantly positive, "-" for significantly negative and "/" for not significant).

	Step 0	**Step 1**	**Step 2**	**Step 3**	**Step 4**	**Step 5**
mean adj. R^2	1.15%	2.65%	3.56%	3.71%	3.74%	3.75%
median adj. R^2	0.95%	2.46%***	3.48%***	3.51%***	3.51%***	3.50%***
$Trimb_{t-1}$	**0.020*** **	**0.019*** **	**0.018*** **	**0.017*** **	**0.017*** **	**0.017*** **
	93%	94%	94%	94%	94%	94%
$Dimb_{1,t-1}$		**0.034*** **	**0.033*** **	**0.033*** **	**0.034*** **	**0.034*** **
		100%	100%	100%	100%	100%
$Dimb_{2,t-1}$			**0.009*** **	**0.009*** **	**0.009*** **	**0.009*** **
			62%	62%	64%	62%
$Himb_{2,t-1}$			**0.065*** **	**0.064*** **	**0.064*** **	**0.064*** **
			98%	98%	98%	98%
$Dimb_{3,t-1}$				**0.004***	**0.003**	**0.003**
				38%	36%	34%
$Himb_{3,t-1}$				**0.021*** **	**0.020*** **	**0.020*** **
				61%	61%	61%
$Dimb_{4,t-1}$					**0.001**	**0.001**
					16%	14%
$Himb_{4,t-1}$					**0.009**	**0.009**
					31%	31%
$Dimb_{5,t-1}$						**0.001**
						11%
$Himb_{5,t-1}$						**0.007**
						18%
Constant	Yes (/)	Yes (/)	Yes (/)	Yes (/)	Yes (/)	Yes (/)
Spread	Yes (/)	Yes (/)	Yes (/)	Yes (/)	Yes (/)	Yes (/)
r_{t-1} to r_{t-5}	Yes (-)	Yes (-)	Yes (-)	Yes (-)	Yes (-)	Yes (-)

Table 4: Estimation Results Order Flow Models

Note: This table presents estimates of the different models from equation (5). The independent variable is the return measured during a five minute interval (expressed in percentage points). Explanatory variables are defined in Subsection 3.1. The models are estimated on a per stock basis for the 95 stocks in our sample. White heteroskedasticity-consistent standard errors are used. The second and third line show the cross-sectional mean and median of the adjusted R^2 of each model. If the Wilcoxon signed rank test indicates that the median difference between the adjusted R^2 from a given model and the adjusted R^2 from the previous 'step' model is significantly different from zero at the 10%, 5% level and 1% level this is indicated with ***, ** and *, respectively. From the fourth line on, the cross-sectional average coefficient estimates are shown. Coefficients which are significant at the 10%, 5% level and 1% level are indicated with ***, ** and *, respectively. Below each estimate we include the percentage of stocks for which the coefficient is significant at the 5% level. In the bottom three lines we indicate whether a constant, the relative quoted spread and lagged returns are included in the models. Between parentheses we show the dominant sign of the coefficients for the different stocks ("+" for significantly positive, "-" for significantly negative and "/" for not significant).

			Step 0			**Step 5**
	Cat. 1	**Cat. 2**	**Cat. 3**	**Cat. 4**	**Cat. 5**	
mean adj. R^2	0.96%	1.22%	1.37%	1.62%	1.64%	3.99%
median adj. R^2	0.77%	1.02%***	1.18%***	1.52%***	1.53%***	3.77%***
$Oimb_{1,t}$	**0.037*** **	**0.038*** **	**0.041*** **	**0.041*** **	**0.041*** **	**0.037*** **
	80%	80%	84%	84%	83%	79%
$Oimb_{2,t}$		**0.018*** **	**0.019*** **	**0.019*** **	**0.019*** **	**0.016*** **
		88%	88%	88%	88%	89%
$Oimb_{3,t}$			**0.015*** **	**0.013*** **	**0.013*** **	**0.012*** **
			76%	67%	66%	67%
$Oimb_{4,t}$				**0.017*** **	**0.017*** **	**0.011*** **
				83%	83%	76%
$Oimb_{5,t}$					**0.001**	**0.001**
					29%	18%
$Dimb_{1,t}$						**0.033*** **
						100%
$Dimb_{2,t}$						**0.008*** **
						63%
$Himb_{2,t}$						**0.062*** **
						97%
$Dimb_{3,t}$						**0.003**
						35%
$Himb_{3,t}$						**0.019*** **
						60%
$Dimb_{4,t}$						**0.001**
						15%
$Himb_{4,t}$						**0.009**
						28%
$Dimb_{5,t}$						**0.001**
						9%
$Himb_{5,t}$						**0.006**
						17%
Constant	Yes (/)	Yes (/)	Yes (/)	Yes (/)	Yes (/)	Yes (/)
Spread	Yes (/)	Yes (/)	Yes (/)	Yes (/)	Yes (/)	Yes (/)
r_{t-1} to r_{t-5}	Yes (-)	Yes (-)	Yes (-)	Yes (-)	Yes (-)	Yes (-)

Table 5: Different Time Horizons

Note: This table deals with predictability at longer horizons. Panel A presents the cross-sectional mean and median of the adjusted R^2 (in percentage) of the models from equation (5), using four different sampling frequencies: five, ten, fifteen and thirty minutes. ***, ** and * denote that the Wilcoxon signed rank test indicates that the median difference between the adjusted R^2 of the model for a given sampling frequency and the sampling frequency from the preceding column is significantly different from zero at the 10%, 5% level and 1% level, respectively. Panel B presents estimates of equation (5), for the step 5 model using four different sampling frequencies: five, ten, fifteen and thirty minutes. For the five minute sampling frequency, we use three different sample sizes that have an equal amount of observations as the ten, fifteen and thirty minute samples (in these samples data are also excluded from the first 50, 75 and 150 minutes of the trading day respectively, for comparability). The independent variable is the return measured during a five minute interval (expressed in percentage points). Explanatory variables are defined in Subsection 3.1. The models are estimated on a per stock basis for the 95 stocks in our sample. White heteroskedasticity-consistent standard errors are used. The cross-sectional average coefficient estimates are shown. Coefficients which are significant at the 10%, 5% level and 1% level are indicated with ***, ** and *, respectively. Below each estimate we include the percentage of stocks for which the coefficient is significant at the 5% level. In the bottom three lines we indicate whether a constant, the relative quoted spread and lagged returns are included in the models. Between parentheses we show the dominant sign of the coefficients for the different stocks ("+" for significantly positive, "-" for significantly negative and "/" for not significant).

Panel A: Adjusted R^2 at different horizons

		mean			
Interval		**5 min.**	**10 min.**	**15 min.**	**30 min.**
	Cat. 1	0.96%	1.04%	1.14%	1.21%
	Cat. 2	1.22%	1.25%	1.35%	1.39%
Step 0	**Cat. 3**	1.37%	1.37%	1.51%	1.50%
	Cat. 4	1.62%	1.63%	1.78%	1.75%
	Cat. 5	1.64%	1.65%	1.80%	1.77%
Step 1		3.00%	3.08%	3.22%	3.20%
Step 2		3.82%	3.98%	4.17%	4.20%
Step 3		3.95%	4.16%	4.35%	4.39%
Step 4		3.98%	4.20%	4.41%	4.48%
Step 5		3.99%	4.22%	4.44%	4.52%
		median			
Interval		**5 min.**	**10 min.**	**15 min.**	**30 min.**
	Cat. 1	0.77%	0.84%**	0.88%***	0.76%
	Cat. 2	1.02%	1.03%	1.15%	0.92%
Step 0	**Cat. 3**	1.18%	1.16%	1.27%	1.03%
	Cat. 4	1.52%	1.49%	1.67%	1.37%
	Cat. 5	1.53%	1.49%	1.69%	1.36%
Step 1		2.87%	2.84%	3.16%	2.85%
Step 2		3.66%	3.80%	4.14%*	4.28%
Step 3		3.73%	4.06%*	4.29%**	4.44%
Step 4		3.76%	4.11%**	4.41%**	4.47%
Step 5		3.77%	4.12%**	4.39%**	4.52%

Table 5 (continued)

Panel B: Coefficient Estimates at Different Horizons						
Horizon	10 min.	5 min.	15 min.	5 min.	30 min.	5 min.
Sample Nr Obs	Full	10' Equiv.	Full	15' Equiv.	Full	30' Equiv.
$Oimb_{1,t-1}$	**0.032****	**0.039****	**0.033***	**0.039****	**0.033**	**0.038***
	52%	66%	45%	54%	24%	47%
$Oimb_{2,t-1}$	**0.016*****	**0.016*****	**0.019*****	**0.016*****	**0.022***	**0.016*****
	74%	87%	74%	76%	53%	73%
$Oimb_{3,t-1}$	**0.009**	**0.011***	**0.011**	**0.012**	**0.008**	**0.012**
	32%	46%	31%	35%	14%	32%
$Oimb_{4,t-1}$	**0.013*****	**0.011*****	**0.016****	**0.011****	**0.018**	**0.011****
	61%	67%	58%	59%	39%	49%
$Oimb_{5,t-1}$	**0.001**	**0.001**	**0.002**	**0.000**	**0.005**	**0.000**
	12%	11%	11%	7%	5%	8%
$Dimb_{1,t-1}$	**0.046*****	**0.033*****	**0.054*****	**0.032*****	**0.071*****	**0.031*****
	100%	100%	99%	99%	98%	99%
$Dimb_{2,t-1}$	**0.013*****	**0.009*****	**0.017*****	**0.009*****	**0.025****	**0.009****
	65%	54%	64%	49%	52%	45%
$Himb_{2,t-1}$	**0.076*****	**0.062*****	**0.083*****	**0.059*****	**0.134*****	**0.084*****
	93%	94%	88%	94%	80%	92%
$Dimb_{3,t-1}$	**0.005**	**0.003**	**0.007**	**0.003**	**0.008**	**0.003**
	26%	24%	21%	27%	17%	21%
$Himb_{3,t-1}$	**0.034*****	**0.017****	**0.039****	**0.017****	**0.067**	**0.015****
	58%	47%	55%	47%	44%	43%
$Dimb_{4,t-1}$	**0.002**	**0.002**	**0.004**	**0.002**	**0.006**	**0.001**
	13%	8%	13%	6%	12%	5%
$Himb_{4,t-1}$	**0.014**	**0.010**	**0.017**	**0.006**	**0.042**	**0.004**
	34%	24%	26%	19%	26%	17%
$Dimb_{5,t-1}$	**0.002**	**0.001**	**0.003**	**0.001**	**0.004**	**0.001**
	14%	6%	11%	5%	5%	5%
$Himb_{5,t-1}$	**0.005**	**0.005**	**0.012**	**0.003**	**0.015**	**0.002**
	18%	17%	17%	16%	7%	14%
Constant	Yes (/)	Yes (/)	Yes (/)	Yes (/)	Yes (/)	Yes (/)
Spread	Yes (/)	Yes (/)	Yes (/)	Yes (/)	Yes (/)	Yes (/)
r_{t-1} to r_{t-5}	Yes (-)	Yes (-)	Yes (-)	Yes (-)	Yes (-)	Yes (-)

Table 6: Time Of Day Patterns

Note: This table presents the cross-sectional mean and median of the adjusted R^2 (in percentage) of the models from equation (5), performed on a set of four subsamples. Observations from each stock are classified into one of four subsamples according to the time of the day. The models are estimated on a per stock basis for the 95 stocks in our sample. ***, ** and * denote that the Wilcoxon signed rank test indicates that the median difference between the adjusted R^2 of the model for a given subsample and the first period subsample is significantly different from zero at the 10%, 5% level and 1% level, respectively.

		mean			
		09.30 a.m. - 11.30 a.m.	**11.30 a.m. - 01.30 p.m.**	**01.30 p.m. - 03.30 p.m.**	**03.30 p.m. - 05.30 p.m.**
	Cat. 1	1.18%	1.27%	1.37%	1.34%
	Cat. 2	1.47%	1.52%	1.65%	1.64%
Step 0	**Cat. 3**	1.69%	1.66%	1.79%	1.76%
	Cat. 4	2.04%	1.92%	2.02%	1.99%
	Cat. 5	2.07%	1.94%	2.03%	2.02%
Step 1		3.56%	3.54%	3.40%	3.25%
Step 2		4.49%	4.33%	4.10%	4.11%
Step 3		4.66%	4.47%	4.19%	4.26%
Step 4		4.70%	4.49%	4.20%	4.30%
Step 5		4.72%	4.51%	4.21%	4.32%
		median			
		09.30 a.m. - 11.30 a.m.	**11.30 a.m. - 01.30 p.m.**	**01.30 p.m. - 03.30 p.m.**	**03.30 p.m. - 05.30 p.m.**
	Cat. 1	0.99%	1.03%	0.87%	1.03%*
	Cat. 2	1.18%	1.26%	1.25%	1.27%*
Step 0	**Cat. 3**	1.31%	1.37%	1.36%	1.52%
	Cat. 4	1.81%	1.64%	1.61%	1.81%
	Cat. 5	1.84%	1.63%	1.62%	1.81%
Step 1		3.37%	3.17%	3.24%*	2.83%**
Step 2		4.54%	4.19%	3.69%***	3.65%***
Step 3		4.54%	4.32%	3.86%***	3.75%***
Step 4		4.63%	4.32%*	3.88%***	3.81%***
Step 5		4.65%	4.36%*	3.89%***	3.89%***

Table 7: Cross-Sectional Determinants of Predictability

Note: This table presents estimates of equation (7), a cross-sectional model with the adjusted R^2 of the different models from equation (5) as the dependent variable, each row corresponds to one model. $InfTrading$ is the adverse selection component of the spread estimated using the methodology of de Jong et al. (1996), $Log(Volume)$ is the logarithm of the average hourly volume expressed in euros, $\sigma(r_t)$ is the standard deviation of the average five minute midquote return. Below the coefficient estimates t-statistics are shown. Coefficients which are significant at the 10%, 5% level and 1% level are indicated with *, ** and ***, respectively. The last column reports the adjusted R^2 of the regression.

		$Constant$	$InfTrading$	$Log(Volume)$	$\sigma(r_t)$	Adj. R^2
	Cat. 1	0.038***	-0.014	-0.002***	-0.017	27%
		5.14	-1.55	-3.59	-1.42	
	Cat. 2	0.034***	0.004	-0.002***	-0.013	23%
		4.62	0.43	-4.01	-1.14	
Step 0	**Cat. 3**	0.034***	0.020**	-0.002***	-0.018	28%
		4.09	1.89	-4.44	-1.35	
	Cat. 4	0.034***	0.028**	-0.002***	-0.011	27%
		3.74	2.47	-4.22	-0.79	
	Cat. 5	0.033***	0.030***	-0.002***	-0.010	27%
		3.59	2.60	-4.15	-0.70	
Step 1		0.061***	0.001	-0.002*	0.040	8%
		3.07	0.04	-1.95	1.25	
Step 2		0.078***	0.039	-0.004***	0.002	13%
		3.65	1.47	-2.95	0.07	
Step 3		0.071***	0.056**	-0.004***	-0.005	12%
		3.18	2.00	-2.74	-0.14	
Step 4		0.068***	0.061**	-0.004***	-0.007	12%
		3.03	2.16	-2.65	-0.18	
Step 5		0.067***	0.061**	-0.004***	-0.006	12%
		2.98	2.15	-2.59	-0.16	