

Asymmetric Effects of the Limit Order Book on Price Dynamics

Tolga Cenesizoglu

Georges Dionne

Xiaozhou Zhou

December 5, 2016

Abstract

We analyze whether the information in different parts of the limit order book affect prices differently. We distinguish between slopes of lower and higher levels of the bid and ask sides and include these four slope measures as well as midquote return and trade direction in a vector autoregressive model. Slope measures of the same side based on different levels affect both short- and long-run price dynamics quite differently, in line with the predictions based on recent theoretical models such as Foucault, Kadan, and Kandel (2005) and Rosu (2009). In a high frequency day trading exercise, we show that ignoring these asymmetries costs a trader approximately 25 basis points in daily profits, suggesting that the asymmetries are important not only statistically but also economically. Our statistical results are robust to using alternative definitions of slope measures and sample periods while our economic results are robust to trading under alternative assumptions such as trading slower speeds.

Key words: Ultra-high frequency data, Hasbrouck model, Limit order book slope, High-frequency trading, Asymmetric Effect.

JEL Classification: G10, G14, G19

1 Introduction

Regardless of their original trading mechanism, almost all of the world's major exchanges now feature electronic limit order books. Some, such as Euronext Paris, have completely abandoned any form of floor trading and operate as pure electronic limit order markets, with no designated market makers. Others, such as Nasdaq, also had to adapt their trading mechanisms to reflect the growing importance of electronic limit order books originating from alternative trading systems such as Electronic Communication Networks.

As the importance of electronic limit order books in financial markets increases, so does the demand for information embedded in them. Most exchanges, such as those operated by NYSE Euronext, now offer investors access to historical and real-time data on their limit order books for a fee. Others, such as the Frankfurt Stock Exchange, make their electronic limit order book data available on their websites with a small delay. More importantly, historical and real-time data on limit order books are available at ever-increasing frequencies, thanks to recent technological advancements in electronic trading systems. For example, the Frankfurt Stock Exchange offers historical data on its electronic limit order book, including trades and quotes up to 20 levels, with millisecond time stamps. Thus, there is an immense wealth of historical and real-time information embedded in the high-frequency limit order books available to investors.

Whether information embedded in the limit order book should have any effect on future price movements is a theoretical question. Earlier microstructure models, such as those of Glosten and Milgrom (1985), Kyle (1985), Glosten (1994), and Rock (1996), treat limit orders as free options provided by uninformed investors to the market and susceptible to being picked off by better-informed investors and thus implicitly assume that the limit order book cannot be informative for future price movements. However, recent theoretical models allow informed investors to strategically choose between limit and market orders, and show that they use not only market orders, as assumed in the previous literature, but also limit orders in a rational expectations equilibrium.¹ Regardless of the channel through which information is embedded in the limit order book, the common prediction of these models is that limit orders should contain relevant information for the true value of the underlying asset and thus affect future price movements.

Hence, it is not surprising to find a growing body of empirical literature analyzing whether the information embedded in the limit order book helps predict future price movements.² However, most of the papers in the previous

¹For example, informed investors could use limit orders to avoid detection, as in Kumar and Seppi (1994), to insure themselves against the price they could obtain for their market orders, as in Chakravarty and Holden (1995), or to take advantage of their sufficiently persistent private information, as in Kaniel and Liu (2006) and Kalay and Wohl (2009). There is also more recent literature on dynamic limit order markets with strategic traders, such as the works of Foucault et al. (2005), Goettler, Parlour, and Rajan (2009), and Rosu (2009). Foucault et al. (2005) show that patient traders tend to submit limit orders, while impatient ones submit market orders in equilibrium. Rosu (2009) shows that fully strategic, symmetrically informed liquidity traders can choose between market and limit orders based on their trade-off between execution prices and waiting costs. Goettler et al. (2009) find that limit orders tend to be submitted mostly by speculators and competition among them results in their private information being reflected in the limit order book.

²Biais, Hillion, and Spatt (1995) are among the first to analyze the dynamics of limit order markets and show many interesting facts. Specifically, they show that price revisions tend to move in the direction of previous limit order flows, suggesting that the limit order book contains information relevant to future price paths. In contrast, Griffiths, Smith, Turnbull, and White (2000) find that limit orders tend to have a negative impact on prices in the Toronto Stock Exchange, because limit orders can be "picked off" by better-informed investors. This result, in turn, suggests that limit orders are placed by less-informed investors and thus do not convey much relevant information about prices. On the other hand, Cao, Hansch, and Wang (2009) provide empirical evidence based on data from the Australian Stock Exchange that the limit order book is somewhat informative, contributing approximately 22% to price discovery. They also show that order imbalances between the demand and supply schedules along the

literature, with only a few exceptions, do not base their empirical analysis on theoretical models. More importantly, none of the papers distinguishes between the information embedded in different levels of the limit order book.

In this paper, we fill this gap by examining whether the information embedded in different parts of the limit order book affects future price dynamics differently. To this end, we reconstruct the first 20 levels of the historical limit order book every millisecond for all stocks in the DAX30 index in June 2011, based on data from the Xetra electronic trading system of the Frankfurt Stock Exchange. Due to its multifaceted nature, there are many ways to summarize the information embedded in the limit order book. We focus on its slope, which is one of the most widely used variables; it is also theoretically motivated. We distinguish not only between the slopes of the bid and ask sides but also between the slopes of different levels, unlike the previous literature.

We then develop, based on recent theoretical literature, several hypotheses on potential asymmetries in the effects of different slope measures on future price dynamics. For the slopes of different sides based on same levels, we argue based on models in Kalay and Wohl (2009), Foucault et al. (2005) and Rosu (2009) that the effect of the bid-side slope on prices should be greater in magnitude than that of the ask-side slope based on the same levels, with an increase in the ask-side slope resulting in higher future prices and increase in the bid-side slope resulting in lower future prices. For the slopes of the same side based on different levels, we argue, consistent with Goettler et al. (2009), Foucault et al. (2005) and Rosu (2009) that the slopes of the lower and higher levels of the ask (bid) side should have different effects on price dynamics. However, these models do not agree on their predictions regarding the signs of these effects. On the one hand, Goettler et al. (2009) predict that an increase in the slope of the lower levels of the ask (bid) side result in higher (lower) future prices, while an increase in the slope of the higher levels of the ask (bid) side result in lower (higher) future prices. On the other hand, one can argue based on Foucault et al. (2005) and Rosu (2009) that the slope measures of the same side based on different levels should have effects of the same sign but potentially of different magnitudes.

To test these hypotheses, we follow Hasbrouck (1991) and consider data in transaction, rather than calendar, period and calculate midquote returns as well as different slope measures right after a trade. For each stock in our sample, we then estimate a separate linear vector autoregression (VAR) that includes midquote return, trade direction, and four slope measures, i.e. the bid- and ask-side slopes based on lower and higher levels. This empirical approach has several advantages over a simple regression framework. For example, it can be considered as a reduced-form linear approximation that is designed to capture the dynamics of limit order market models, and the residuals of slope measures can be interpreted as an unexpected private information shock embedded in these slope measures. More importantly, this empirical framework allows us to test the predictions for the immediate, short- and long-run effects

book are significantly related to future short-term returns, even after controlling for autocorrelations in returns, inside spread, and trade imbalance. Similarly, using data from NYSE's Trades, Orders, Reports, and Quotes, Kaniel and Liu (2006) argue that informed traders prefer limit orders to market orders and limit orders are therefore more informative than market orders. More recently, Beltran-Lopez, Giot, and Grammig (2009) also demonstrate that factors extracted from the limit order book have non-negligible information relevant to the long-run evolution of prices in the German Stock Exchange. Specifically, they find that shifts and rotations of the order book can explain between 5% to 10% of the long-run evolution of prices, depending on the liquidity of the asset. Kozhan and Salmon (2012) provide empirical evidence that variables summarizing the information in the limit order book have statistically significant power in predicting future price movements. However, they argue that this statistical relation cannot be exploited to provide economic value in a simple trading exercise.

of different slope measures on future price dynamics.

We test these predictions by comparing the coefficient estimates (and their functions) on different slope measures. Before doing so, we briefly discuss our findings on the individual effects of each slope measure on price dynamics. The ask-side slope based on lower levels has a significantly positive immediate effect on prices of all stocks in the DAX30 index. The coefficient estimates on further lags of the ask-side slope based on lower levels are mostly negative with differing levels of significance, suggesting a reversal in its positive immediate effect on price dynamics. However, the coefficient estimates on further lags are generally smaller (in magnitude) than those on the first lag. The sum of the coefficient estimates on all five lags is significantly positive for all stocks in our sample, suggesting a significant positive long run effect of the ask-side slope based on lower levels. The impulse response functions of returns to the ask-side slope measures based on lower levels are also significantly positive and it take about 40 transactions for the effect of the shock to be fully realized. The coefficient estimates on all lags of the ask-side slope based on lower levels are also jointly significantly different from zero, implying a significant, and potentially causal, overall effect of the ask-side slope based on lower levels on short-run price dynamics. The results for the effect of the ask-side slope based on higher levels are similar but weaker. Although the ask-side slope based on higher levels has a positive immediate effect on price dynamics of most stocks in our sample, its long-run effect is significantly positive for 17 and significantly negative for two out of 30 stocks in our sample. Furthermore, the ask-side slope based on higher levels has a significant overall effect on short-run price dynamics when we consider the empirical evidence across stocks jointly based on Bonferroni p-values. That said, it has significant overall effect on short-run price dynamics for half of the stocks in our sample when we consider the empirical evidence for individual stocks separately. The empirical results for the bid-side slope measures are very similar to those for the corresponding ask-side slope measures but with opposite signs.

We now turn our attention to the empirical evidence in support of our hypotheses, starting with the short-run effects of different slope measures on price dynamics. First, the ask-side slope has a significantly different immediate effect on prices than the bid-side slope, regardless of the levels used to measure them, in line with our hypotheses. However, the empirical evidence in support of our predictions regarding the relative magnitudes of these immediate effects is weaker. To be more precise, measures of the bid-side slope have significantly stronger (greater in absolute value) immediate effects than the ask-slope based on the same corresponding levels for approximately half of the stocks in our sample, and significantly so for four stocks. Second, slope measures based on higher levels have significantly stronger immediate effects on prices than slope measures based on lower levels of the same side. However, slope measures of the same side based on different levels have immediate effects of the same sign, in line with the predictions of Foucault et al. (2005) and Rosu (2009) but in contrast to those of Goettler et al. (2009). Finally, slope measures of different sides based on the same levels as well as slope measures of the same side based on different levels have significantly different overall effects on short-run price dynamics, in line with our hypotheses.

Regarding the effects of slope measures on long-run price dynamics, we also find statistically significant evidence that is mostly in line with our hypotheses. To be more precise, we first consider the sum of the coefficients on all lags

of slope measures in the return equation as a first raw approximation of the long-run effects of slope measures on price dynamics, as argued by Dufour and Engle (2000). Long-run effects of the ask- and bid-side slopes based on lower levels are significantly different from each other in magnitude. The relative magnitudes of these long-run effects are in line with our hypotheses for half of the stocks in our sample and significantly so for only two stocks. Furthermore, there do not seem to be any statistically significant differences between the long-run effects of the ask- and bid-side slopes based on higher levels. On the other hand, there is strong empirical evidence for asymmetries between the long-run effects of slope measures based on different levels of the same side. More importantly, there is also empirical evidence that slope measures based on different levels of the same side have long-run effects of the opposite sign, in line with our hypotheses.

Although the sum of coefficients on all lags of a given slope measure can be considered as a first raw approximation for its long-run effect on prices, it might not reveal the long-run cumulative effect after several trades given that slope measures have significant dynamics of their own. Impulse response functions take the dynamics of slope variables into account and, thus, capture the long-run effect better than the simple sum of coefficient estimates. The empirical evidence based on impulse response functions is similar to that based on the sum of the coefficients and can be summarized as follows: The impulse response function to the ask-side slope is greater than that to the bid-side slope for half of the stocks, regardless of the levels used to measure them, while the opposite holds for the other half. More importantly, the difference is statistically significant at most for one stock in our sample. This in turn suggests that there is no empirical evidence whatsoever for any asymmetry between the long-run effects of different sides based on the same levels. In contrast, there is ample empirical evidence for asymmetries in the long-run effects of slope measures based on different levels of the same side, especially in the medium-run between 5 to 40 transaction periods following a shock. To be more precise, the impulse response function of returns to the ask- (bid-) side slope based on lower levels between 5 to 40 transaction periods following a shock is greater (smaller) than that to the ask- (bid-) side slope based on higher levels for more than 20 stocks, and significantly so for more than 5 (10) stocks.

Having found statistically significant evidence in support of certain asymmetries between the effects of different slope measures on price dynamics, we then show that these asymmetries can also be economically significant. We do this by comparing the performances of high-frequency day-trading strategies that ignore the information embedded in different types of asymmetries with that of an unrestricted strategy that uses this information. In other words, our unrestricted strategy uses the unrestricted model to forecast midquote returns at each transaction period while the competing (restricted) strategies employ restricted versions of this model so that a chosen pair of slope variables has symmetric effects on price dynamics. The trading strategy we consider is similar to that discussed in Kozhan and Salmon (2012) and can be summarized as follows: At each transaction period for a given stock, we take a snapshot of the limit order book right after (less than a millisecond after) observing the transaction. We then compute the forecast of the midquote return in the next transaction period based on this snapshot and a given forecasting model. We consider a forecast greater (less) than a threshold to be a buy (sell) signal and do one of the following depending on our existing position in the stock: (1) buy (short-sell) one share of the stock if we do not already have an existing

position in the stock; (2) buy (short-sell) two shares of the stock if we have an existing short (long) position in the stock, i.e. close the short (long) position and take a long (short) position of one share; (3) do nothing if we already have a long (short) position.

We compute the differences between the average daily cumulative returns of the unrestricted and restricted trading strategies. A positive difference implies that the unrestricted strategy provides, on average, higher daily cumulative returns than the trading strategy that uses the restricted forecasting model implied by a given hypothesis and suggests that the information embedded in this asymmetry is economically important. Overall, our results suggest that short-run asymmetries between different slope measures are on average economically important. More importantly, the evidence is strongest for the asymmetries between the overall short-run effects of slope measures based on different levels of the same side. To be more precise, the unrestricted strategy provides a higher average daily cumulative return than the strategies restricting the overall short-run effects of slope measures based on different levels of the same side to be the same for more than 25 stocks. Averaged over all stocks, the differences are, respectively, about 25 and 24 basis points between the unrestricted strategy and the strategies imposing these restrictions. These are in line with our hypothesis test results discussed above, which suggest that the empirical evidence for the asymmetries between the effects of slope measures based on different levels of the same side is relatively stronger than asymmetries between the effects of slope measures of different sides based on same levels. These results are robust to trading at different thresholds, at the best bid and ask prices instead of midquote prices and at slower speeds.

The rest of the paper is organized as follows. Section 2 presents the details of our data set. Section 3 discusses the theoretical motivation behind our empirical analysis. Section 4 presents the empirical model and related empirical choices. Section 5 develops our testable hypotheses based on the theoretical model and empirical model. Section 6 discusses the estimation results. Sections 7 and 8 present our main empirical results on the asymmetries in the short- and long-run effects of different slope variables on price dynamics, respectively. Section 9 discusses the robustness of our main empirical results to using alternative definitions and sample periods. Section 10 shows that the asymmetries in the effect of slope measures on price dynamics are also economically important. Section 11 concludes the paper.

2 Data

Our data are from the automated order-driven trading system Xetra operated by the Deutsche Börse Group at the Frankfurt Stock Exchange. It is the main German trading platform, accounting for more than 90% of total transactions at all German exchanges. In Xetra, there are no dedicated market makers for blue chip and other liquid stocks, unlike the NYSE, where dedicated specialists are responsible for providing liquidity to the market. Thus, all liquidity in Xetra is provided by market participants submitting limit orders.

The raw data set contains all events that are tracked and sent through the data streams. We first process the raw data set using XetraParser software, developed by Bilodeau (2013).³ We then reconstruct the first 20 levels of the

³We thank Yann Bilodeau for his comments and help in constructing the data set.

limit order book in millisecond time intervals between the normal trading hours of 9:00 a.m. and 5:30 p.m.⁴ The limit order book can change when either a trade is executed or a limit order is placed, modified, or canceled. In the unlikely event that these two types of events have the same millisecond time stamp, we need to make an assumption on the sequence of events, given that we do not observe which one arrived earlier. We assume that a trade is always executed before any other change to the limit order book with the same millisecond time stamp. Thus, we first modify the limit order book to reflect the trade execution before taking its snapshot. In other words, if a trade is executed at a given millisecond, then the snapshot of the limit order book for that millisecond already reflects the executed trade. To avoid problems due to this assumption, we ignore the state of the limit order book when a trade is executed and use its snapshot 1 millisecond after a trade.

Our data cover all stocks in the DAX30 index and all trading days in June 2011. Table 1 presents the list of stocks in the DAX30 index as of June 2011 along with some daily summary statistics from the Security Daily files in Compustat Global. We choose one month of data simply due to the sheer size of ultra high-frequency limit order books. Furthermore, as we will discuss below, we present detailed estimation results for ALV and selected results for all other stocks. We choose ALV as a representative stock because its characteristics such as market capitalization, turnover and return are similar to the average stock in the DAX30 index, as can be seen from Table 1.

[Insert Table 1 here]

Due to its multifaceted nature, there are many ways to summarize the information embedded in the limit order book. In this paper, we focus on its slope, which is theoretically motivated and is one of the most widely used variables. However, unlike previous research, we distinguish between the slopes based on lower and higher levels of the limit order book. Specifically, let $P_{l,t}^B$ and $P_{l,t}^A$ denote l^{th} best bid and ask prices, respectively, in period t . Similarly, let $D_{l_1+1,l_2,t}^B$ and $D_{l_1+1,l_2,t}^A$ denote the cumulative quantity available between levels $l_1 + 1$ and l_2 (both levels inclusive and $l_2 > l_1$) in the bid and ask sides of the limit order book, respectively. The slopes of the bid and ask sides between levels l_1 and l_2 in period t , $S_{l_1,l_2,t}^B$ and $S_{l_1,l_2,t}^A$, are defined as the change in the price relative to the cumulative quantity available between levels l_1 and l_2 :

$$S_{l_1,l_2,t}^B = \frac{P_{l_2,t}^B - P_{l_1,t}^B}{D_{l_1+1,l_2,t}^B}, \quad (1)$$

$$S_{l_1,l_2,t}^A = \frac{P_{l_2,t}^A - P_{l_1,t}^A}{D_{l_1+1,l_2,t}^A} \quad (2)$$

for $l_1 = 1, \dots, 19$ and $l_2 > l_1$. The slope of the bid side is a measure of price sensitivity to changes in quantity demanded and is always negative. A high (in absolute value) bid slope coefficient implies that the price between two

⁴During normal trading hours, there are two types of trading mechanisms: call auctions and continuous auctions. For stocks listed on the DAX 30, there are three call auctions during a trading day: the open, mid-day, and closing auctions. The prices during call auctions are not determined by trading activity but, rather, are based on a set of rules determined by the exchange. Between the call auctions, the market is organized as a continuous auction in which traders can only submit round lot-sized limit and/or market orders. The prices from the call auctions serve as the opening prices for the following continuous auctions. To avoid any bias due to the peculiar structure of the call auctions, we ignore all data corresponding to the three call auctions for a DAX 30 stock.

levels of the bid side will decrease more, on average, for a given change in quantity demanded. In other words, an increase in the bid slope suggests that the investors are only willing to buy the same total quantity at lower prices. Similarly, the slope of the ask side is a measure of price sensitivity to changes in quantity supplied and is always positive. A high ask slope coefficient implies that the price between two levels of the ask side will increase more, on average, for a given change in quantity supplied, which, in turn, suggests that the investors are only willing to sell the same total quantity at higher prices.

Figure 1 presents two snapshots of the limit order book for ALV on June 1, 2011 along with the corresponding ask and bid-side slope measures between the first and fifth levels and between the fifth and twentieth levels. As can be easily seen from Figure 1, both the bid and ask sides of the limit order book can take on different shapes at different times. More importantly, the slope measures can take on very different values depending on the side and levels used to measure and can change significantly even within a few hours. Table 2 presents mean and standard deviation of (log) slope measures for each stock separately in June 2011.

[Insert Figure 1 here]

[Insert Table 2 here]

3 Theoretical Background

In this section, we discuss the theoretical motivation behind our empirical analysis. To this end, we review the predictions of recent theoretical models on limit order markets regarding how different slope measures might affect price dynamics differently. We start with slope measures of different sides based on the same levels before turning our attention to slope measures of the same side but based on different levels.

Our discussion of the potential asymmetries in the effects of slope measures of different sides is mostly based on Kalay and Wohl (2009), who develop a model that makes explicit predictions on this issue. Specifically, Kalay and Wohl (2009) solve for the equilibrium in the noisy rational expectations model of Hellwig (1980) under the assumption that only informed, and not liquidity, traders can submit price-sensitive demand and/or supply schedules. In this framework, they show that buying pressure by informed traders will result in a decrease (in absolute value) of the bid-side slope as well as higher future prices of the underlying asset. Thus, one expects future prices to go up following a decrease (in absolute value) of the bid-side slope. The opposite intuition holds for the slope of the ask side. More importantly for the purposes of our paper, they also show the difference between ask and (absolute value of) bid-side slope to be negatively correlated with future price changes, suggesting that the effect of the bid-side slope on prices should not only be significantly different but also bigger in magnitude than that of the ask-side slope. However, they do not distinguish between slopes based on different levels. Assuming that their arguments extend to the slopes of both lower and higher levels, their model predicts that (1) the effect of the bid-side slopes of the lower levels should not only be significantly different but also bigger in magnitude than that of the ask-side slope of the same

levels, with an increase in slope of the lower levels of the ask side resulting in higher future prices, and an increase in slope of the lower levels of the bid side resulting in lower future prices; (2) the effect of the bid-side slopes of the higher levels should not only be significantly different but also bigger in magnitude than that of the ask-side slope of the same levels, with an increase in slope of the higher levels of the ask side resulting in higher future prices, while an increase in slope of the higher levels of the bid side results in lower future prices.

Note that the model of Kalay and Wohl (2009) and hence its empirical predictions are driven by the underlying assumption of information asymmetry. We choose to focus mostly on their model in our discussion since they explicitly derive implications of how slopes of different sides based on the same levels might affect price dynamics differently. However, one does not need to rely on the assumption of information asymmetry and can obtain similar predictions from models with alternative underlying assumptions. For example, Foucault et al. (2005) develop a dynamic model of limit order markets where traders can place limit orders that cannot be cancelled or changed. Rosu (2009) extends their model and allow traders to modify their limit orders dynamically in real time. More importantly, both of these models assume that traders are symmetrically informed but face different waiting costs and thus have varying degrees of patience. They show that the shape of the limit order book and its dynamics depend closely on the proportion of patient and impatient traders and their arrival rates. The intuition follows from the result that impatient traders tend to demand immediacy and place market orders while patient ones tend to place limit orders and do so at higher levels of the book the more patient they are. Neither of these papers derives explicit implications on how slopes of different sides based on the same levels might affect price dynamics differently. Based on their results, one can nevertheless argue that predictions similar to those in Kalay and Wohl (2009) might also hold under their assumption of differences in traders' patience. To see this, consider the case where more patient buyers arrive at the market while everything else remains the same. This has two effects on the market: (1) the bid side of the book immediately becomes flatter due to the fact that patient buyers tend to place more limit orders than market orders; (2) the buying pressure on the stock increases due to the increased presence of buyers in the market, which in turn implies higher future prices. Taking these two effects into consideration, one expects future prices to be higher following a decrease in the bid-side slope. A similar intuition but with opposite signs holds for the relation between the ask-side slope and future price movements, and one expects future prices to be lower following a decrease in the ask-side slope. If one is willing to further assume that buyers are, on average, more patient than sellers, the effect of the bid-side slope on future prices should also be stronger than that of the ask-side slope.⁵ It is then not difficult to see that these predictions are very similar to those developed under the assumption of information asymmetry.

We now turn our attention to how slopes based on different levels of the same side might affect price dynamics differently. Once again, we focus on a theoretical paper, namely that of Goettler et al. (2009), which addresses this issue explicitly, before discussing how one might obtain similar predictions based on alternative assumptions. Goettler et al. (2009) develop a model in which traders optimally choose the type of order to submit and whether to acquire

⁵Although we believe that it is reasonable to assume that buyers are, on average, more patient than sellers, whether this is the case in reality is an empirical question.

information about the asset. They solve for the equilibrium of this model and show that only a few orders in the book are stale, because traders submitting limit orders revisit the market and resubmit orders, on average, twice as often as the true value of the asset changes. Thus, orders submitted in the higher levels of the ask side suggest that the current best ask is too low and hence lead to an upward revision in expectations about the true value of the asset. On the other hand, given that the transactions prices and traders' beliefs are, on average, equal to the true value of the asset, depth at the best ask quote lead to lower prices. The opposite intuition holds for the bid side.

Goettler et al. (2009) discuss the predictions of their model in terms of the *depth* rather than the *slope* of the limit order book. Specifically, they predict that depth in the first level of the book has a different effect on future prices than depth in levels higher than the first, i.e. second and above. In unreported analysis, we considered this precise prediction and obtained results qualitatively similar to those presented. However, we should note here that depth and slope are the reciprocal of each other in their framework. This is due to their assumption that the tick sizes between any two consecutive prices in the limit order book are exactly equal to one. This assumption, of course, does not hold in reality, and one needs to take both depth and price into account to fully capture the information embedded in the limit order book. A slope measure achieves this but also requires more than one level to compute. To this end, we take their predictions one step further and analyze differences between the effects of slope measures based on lower and higher levels. Furthermore, their model does not make any predictions about the relative magnitudes of these effects, which we nevertheless analyze empirically. Taking these into account, their model predicts that (1) the slope of the lower and higher levels of the ask side should have different effects on price dynamics, with an increase in slope of the lower levels of the ask side results in higher future prices, while an increase in slope of the higher levels of the ask side resulting in lower future prices; (2) the slope of the lower and higher levels of the bid side should have different effects on price dynamics, with an increase in slope of the lower levels of the bid side results in lower future prices, while an increase in slope of the higher levels of the bid side resulting in higher future prices.

Similar to Kalay and Wohl (2009), the model in Goettler et al. (2009) is also based on the assumption of information asymmetry among traders. Once again, one does not need to rely on this assumption to obtain differences in the effect of slope measures based on different levels of the same side. For example, based on the results in Foucault et al. (2005) and Rosu (2009), one can also argue that slope measures of the same side based on different levels should have effects of different magnitudes, but not necessarily of different signs. To be more precise, assume that more patient buyers arrive in the market while everything else remains the same. As discussed above, not only does the bid side of the book immediately become flatter, but the buying pressure on the stock also increases, which in turn implies higher future prices. In addition, the slope of the higher levels decreases more than that of lower levels since more patient buyers tend to place more orders in the higher levels than in the lower ones. If one is willing to assume that the arrival rate of more patient buyers is higher, on average, than that of less patient buyers, then the effect of the bid-side slope based on higher levels should be stronger than that of the bid-side slope based on lower levels. A similar intuition but with opposite signs holds for the ask side. These are similar predictions to those based on Goettler et al. (2009), but with differences in terms of the signs of certain effects.

Several remarks are in order. First, our discussion in this section forms the basis of our testable hypotheses, which we will present in Section 5, and we discuss our empirical findings in light of these theoretical predictions. However, we empirically analyze whether different slope measures affect price dynamics differently without taking a stand on different explanations. Second, as mentioned in Section 2, our data are from Xetra, where there are no dedicated market makers for blue chip and other liquid stocks and, thus, all liquidity is provided by market participants who submit limit orders. Of course, the match between the actual trading mechanism in Xetra and those in some of the theoretical models discussed above is not perfect. These models nevertheless provide guidance for empirical analysis using data from a completely order-driven market. Finally, most of the papers discussed above implicitly assume that the effects of slope measures on price dynamics are realized immediately or in a very short period of time. Based on this assumption, we should, strictly speaking, only consider the immediate effects of slope measures on price dynamics in our empirical analysis. However, in reality, this assumption might not hold. In other words, it might take several trades for the effect of any changes in slope measures on prices to be realized. Given that different slope measures have their own dynamics, this, in turn, might cause both the immediate, overall short-run effects and the long-run effects to be quite different. Rather than testing only the immediate effect, we use these predictions to form the basis of different testable hypotheses, which then allow us to analyze the asymmetries in the effect of slope measures on prices from several different angles.

4 The Empirical Model

In this section, we first present the empirical model and related empirical choices. We then discuss how this model allows us to test our hypotheses from different angles.

Our empirical model is based on the study by Hasbrouck (1991), which shows that a VAR for the interactions between returns and trade directions is consistent with stylized market microstructure models such as that of Glosten and Milgrom (1985). Specifically, Hasbrouck (1991) suggests that the following VAR model be used to analyze the effects of information embedded in trades on prices:

$$r_t = \sum_{\tau=1}^{\infty} \alpha_{r,\tau} r_{t-\tau} + \sum_{\tau=1}^{\infty} \alpha_{x,\tau} x_{t-\tau} + \varepsilon_{r,t}, \quad (3a)$$

$$x_t = \sum_{\tau=1}^{\infty} \beta_{r,\tau} r_{t-\tau} + \sum_{\tau=1}^{\infty} \beta_{x,\tau} x_{t-\tau} + \varepsilon_{x,t}, \quad (3b)$$

where t indexes trades; x_t is the sign of the trade in period t (+1 for a trade initiated by a buyer and -1 for a trade initiated by a seller); r_t is the midquote return defined as the change in the average of the best bid and ask quotes between periods $t-1$ and t , that is, $r_t = \Delta q_t = q_t - q_{t-1}$; and q_t is the simple average of the best bid and ask quotes in period t . This is a very general and flexible model that nests many of the standard microstructure models as special cases. The disturbances in this framework, $\varepsilon_{r,t}$ and $\varepsilon_{x,t}$, are generally modeled as white noise processes and can be

interpreted as public information embedded in unexpected returns and private information embedded in unexpected trades, respectively.

We assume that the dynamics of a limit order market similar to those discussed in Section 3 can also be approximated by a linear VAR system similar to that proposed by Hasbrouck (1991). Specifically, we include slope measures based on different levels of the ask and bid sides as additional state variables in the above VAR, which yields:⁶

$$r_t = \sum_{\tau=1}^{\infty} \alpha_{r,\tau} r_{t-\tau} + \sum_{\tau=1}^{\infty} \alpha_{x,\tau} x_{t-\tau} + \sum_{\tau=1}^{\infty} \boldsymbol{\alpha}'_{z,\tau} \mathbf{z}_{t-\tau} + \varepsilon_{r,t}, \quad (4a)$$

$$x_t = \sum_{\tau=1}^{\infty} \beta_{r,\tau} r_{t-\tau} + \sum_{\tau=1}^{\infty} \beta_{x,\tau} x_{t-\tau} + \sum_{\tau=1}^{\infty} \boldsymbol{\beta}'_{z,\tau} \mathbf{z}_{t-\tau} + \varepsilon_{x,t}, \quad (4b)$$

$$\mathbf{z}_t = \sum_{\tau=1}^{\infty} \boldsymbol{\gamma}_{r,\tau} r_{t-\tau} + \sum_{\tau=1}^{\infty} \boldsymbol{\gamma}_{x,\tau} x_{t-\tau} + \sum_{\tau=1}^{\infty} \boldsymbol{\gamma}'_{z,\tau} \mathbf{z}_{t-\tau} + \varepsilon_{z,t}, \quad (4c)$$

where \mathbf{z}_t is a vector that includes the slope measures of interest.

To implement this model empirically, we need to make some empirical choices regarding the sampling approach, the slope measures and the truncation point for the infinite sums. Regarding the sampling approach, we can measure limit order book variables, including the best bid and ask prices, every millisecond. However, a trade can only be matched to a millisecond interval and thus one needs to decide whether to take a snapshot of the limit order book right before or right after a trade. The theory does not provide much guidance on this issue. We therefore follow the previous literature, e.g. Hasbrouck (1991) and Dufour and Engle (2000), and measure the limit order book variables right after a trade.⁷ This sampling approach implies that the midquote return and limit order book variables in period t are observed right after (less than a millisecond after) the trade in period t and its direction. Hence, we include the trade direction in period t to control for its contemporaneous effect on returns and limit order book variables in the estimated version of Equation (4).

Regarding the truncation issues, we follow Hasbrouck (1991) and Dufour and Engle (2000) and truncate the infinite sums in Equation (4) at five lags, assuming that five lags are sufficient to capture the dynamics of the variables of interest.⁸ Furthermore, the timing convention discussed above is reflected in the starting points of the summations in the estimated version of Equation (4). To be more precise, the summations for trade direction in the equations for the returns and limit order book variables start at zero instead of one. The estimated version of the model is then as

⁶We use bold letters to distinguish vectors and matrices from scalars.

⁷We also considered the alternative sampling approach of measuring the limit order book variables right before a trade. Our results remain qualitatively similar.

⁸We consider different lag structures, up to a maximum of eight lags. To be consistent with the previous literature, we present results based on a lag structure of five lags. The results based on the model estimated using different numbers of lags are similar to those presented in the paper and are available from the authors upon request.

follows:

$$r_t = \sum_{\tau=1}^5 \alpha_{r,\tau} r_{t-\tau} + \sum_{\tau=0}^5 \alpha_{x,\tau} x_{t-\tau} + \sum_{\tau=1}^5 \alpha'_{z,\tau} \mathbf{z}_{t-\tau} + \varepsilon_{r,t}, \quad (5a)$$

$$x_t = \sum_{\tau=1}^5 \beta_{r,\tau} r_{t-\tau} + \sum_{\tau=1}^5 \beta_{x,\tau} x_{t-\tau} + \sum_{\tau=1}^5 \beta'_{z,\tau} \mathbf{z}_{t-\tau} + \varepsilon_{x,t}, \quad (5b)$$

$$\mathbf{z}_t = \sum_{\tau=1}^5 \gamma_{r,\tau} r_{t-\tau} + \sum_{\tau=0}^5 \gamma_{x,\tau} x_{t-\tau} + \sum_{\tau=1}^5 \gamma'_{z,\tau} \mathbf{z}_{t-\tau} + \varepsilon_{z,t}. \quad (5c)$$

We include four slope measures: the first two are the bid- and ask-side slopes between their corresponding first and fifth levels, $S_{1,5,t}^B$ and $S_{1,5,t}^A$, which we use to capture the slopes of lower levels, and the other two are the bid- and ask-side slopes between their corresponding fifth and twentieth levels, $S_{5,20,t}^B$ and $S_{5,20,t}^A$, which we use to capture the slopes of higher levels. This empirical choice is motivated by two factors. First, the first level is undoubtedly the most frequently updated one, and we thus want to include this information in our definition of the slope of the lower level. Second, levels of the limit order book higher than 10 are less frequently updated and might have stale information. We want to minimize the effect of this stale information by including levels between five and ten, which are still updated quite frequently, in our definition of higher levels. In Section 9, we show that our results are robust to using alternative definitions of lower and higher levels.

We estimate the empirical specification in Equation 5 via ordinary least squares (OLS) with heteroskedasticity and autocorrelation consistent standard errors a la Newey and West (1987)). This specification has several advantages compared to a regression framework. First, it can be considered a reduced-form linear approximation that is designed to capture the dynamics of limit order market models discussed in the introduction. It is also used by Brogaard, Hendershott, and Riordan (2016) in a similar fashion to understand the role of limit orders in price discovery. Second, Goettler et al. (2009) argue that competition among speculators results in their private information being partially revealed in the limit order book. Hence, $\varepsilon_{z,t}$ in this framework can be interpreted as an unexpected private information shock embedded in the limit order book. Lastly, this empirical specification is flexible and allows us to analyze the effect of slope measures on prices from different angles. To be more precise, the coefficient estimates on the first lag of slope variables can be interpreted as their immediate effect on prices while controlling for their own lags and other lagged information, similar to a simple linear regression framework. When we also consider the coefficient estimates on further lags of these slope variables, we have an idea about their overall short-run dynamic effect. This empirical approach also allows us to analyze the long-run effect of limit-order book related information on prices, which is not possible based on a simple regression framework.

Note that the terms short-, medium- and long-run in this context have different interpretation than a standard VAR with monthly data, for example. We refer to one or two transactions after a shock as the short-run, to any period between two and 20 transaction periods as the medium run, and to any period more than 20 transactions from a shock as the long-run. To analyze the long-run effect, we consider the sum of the coefficient estimates on all lags

of a slope variable as a first raw approximation of its long-run impact on prices, as argued by Dufour and Engle (2000). However, the simple sum of coefficient estimates might not reveal the long-run cumulative effect because slope measures have significant dynamics of their own. Thus, we also compute the impulse response functions of returns to slope measures to analyze the asymmetries in the long-run cumulative effects of shocks to slope measures on returns. We do this based on the simulation approach discussed in Hamilton (1994). More precisely, we first simulate the estimated VAR for a long enough period, setting all residual terms to zero to obtain its steady state. Starting with the steady state, we then simulate the VAR once more, but this time assuming that the initial residual of slope measure of interest is equal to the standard deviation of this slope measure while all the other residuals (initial or future) remain at zero. The difference between these two simulations of the VAR is the impulse response function of the returns to a one standard deviation shock to the slope variable of interest. Note that the impulse response functions computed based on this approach are known as nonorthogonalized impulse response functions and are different than orthogonalized impulse response functions, which are much more commonly used in the literature. Unlike the orthogonalized impulse response functions, the ordering of variables does not play a role in computation of nonorthogonalized impulse response functions. To compute the standard errors of these impulse response functions, we follow the Monte Carlo simulation approach discussed in Hamilton (1994). To this end, we first draw from the asymptotic distribution of the coefficient estimates, which is a multivariate normal distribution, and calculate the impulse response function based on this random draw of coefficient estimates. We repeat this 1000 times and calculate the lower and upper confidence bands as the 5% and 95% quantiles of these 1000 repetitions.

5 Testable Hypotheses

As mentioned above, our empirical framework allows us to analyze potential asymmetries in the immediate, short- and long-run effects of slope measures on prices. We start our discussion with how to test for asymmetries in the immediate effects of different slope measures on prices. We do this by testing the equality of the coefficient estimates on the first lags of different slope measures. Our estimation results, which we will discuss in detail in Section 6, suggest that these coefficient estimates are positive for the ask-side and negative for the bid-side slope measures regardless of the levels used to measure them, for almost all stocks in our sample. This is in line with the predictions of most theoretical models including those discussed in Section 3, with the exception of Goettler et al. (2009), which predicts that ask- and bid-side slopes based on higher levels have negative and positive effects, respectively. It then makes more sense to test the equality of the relative magnitudes of these coefficient estimates, i.e. their absolute values, when analyzing the asymmetries in the immediate effects of slopes measures of different sides. Therefore, our first set of testable hypotheses on the asymmetries in the immediate effects of slope measures of different sides (H1a for lower levels and H2a for higher levels) and the same side but based on different levels (H3a for the ask side and H4a for the bid side)

can be written mathematically as follows:

$$\begin{aligned}
H1a & : \alpha_{z,1}(S_{1,5,t}^A) = (1 - 2I_{1a})\alpha_{z,1}(S_{1,5,t}^B); \\
H2a & : \alpha_{z,1}(S_{5,20,t}^A) = (1 - 2I_{2a})\alpha_{z,1}(S_{5,20,t}^B); \\
H3a & : \alpha_{z,1}(S_{1,5,t}^A) = \alpha_{z,1}(S_{5,20,t}^A); \\
H4a & : \alpha_{z,1}(S_{1,5,t}^B) = \alpha_{z,1}(S_{5,20,t}^B);
\end{aligned}$$

where $\alpha_{z,i}(\cdot)$ denote the element of $\alpha_{z,i}$ corresponding to the slope measure of interest in parentheses. The terms multiplying the coefficients on the bid-side slope measures, i.e. $(1 - 2I_{1a})$ and $(1 - 2I_{2a})$, allow us to compare the magnitudes of the coefficient estimates if they are positive for the ask-side and negative for the bid-side slope. To be more precise, I_{1a} and I_{2a} are binary variables defined as $I_{1a} = 1_{\{\alpha_{z,1}(S_{1,5,t}^A) > 0, \alpha_{z,1}(S_{1,5,t}^B) < 0\}}$ and $I_{2a} = 1_{\{\alpha_{z,1}(S_{5,20,t}^A) > 0, \alpha_{z,1}(S_{5,20,t}^B) < 0\}}$ where $1_{\{x,y\}}$ is an indicator function that takes the value of one if both conditions x and y are satisfied and zero otherwise.

We now turn our attention to how to test for the asymmetries in the overall short-run effects of different slope measures on prices. We do this by jointly testing the equality of the coefficient estimates on all five lags of different slope measures. To be consistent with our discussion above, we take the sign of the estimated coefficients into account when comparing the overall short-run effects of slope measures of different sides. To be more precise, if the coefficient estimate on a given lag of the ask-side slope measure is positive and that on the corresponding lag of the bid-side slope measure is negative, we compare their relative magnitudes instead of their values. Therefore, our second set of testable hypotheses can then be written as follows:

$$\begin{aligned}
H1b & : \alpha_{z,\tau}(S_{1,5,t}^A) = (1 - 2I_{1b,\tau})\alpha_{z,\tau}(S_{1,5,t}^B) \quad \text{for } \tau = 1, 2, \dots, 5; \\
H2b & : \alpha_{z,\tau}(S_{5,20,t}^A) = (1 - 2I_{2b,\tau})\alpha_{z,\tau}(S_{5,20,t}^B) \quad \text{for } \tau = 1, 2, \dots, 5; \\
H3b & : \alpha_{z,\tau}(S_{1,5,t}^A) = \alpha_{z,\tau}(S_{5,20,t}^A) \quad \text{for } \tau = 1, 2, \dots, 5; \\
H4b & : \alpha_{z,\tau}(S_{1,5,t}^B) = \alpha_{z,\tau}(S_{5,20,t}^B) \quad \text{for } \tau = 1, 2, \dots, 5;
\end{aligned}$$

where $I_{1b,\tau} = 1_{\{\alpha_{z,\tau}(S_{1,5,t}^A) > 0, \alpha_{z,\tau}(S_{1,5,t}^B) < 0\}}$ and $I_{2b,\tau} = 1_{\{\alpha_{z,\tau}(S_{5,20,t}^A) > 0, \alpha_{z,\tau}(S_{5,20,t}^B) < 0\}}$ for $\tau = 1, 2, \dots, 5$.

We then test for the asymmetries in the long-run effects of different slope measures on prices by comparing the sum of the coefficient estimates on all five lags, which, as mentioned above, can be interpreted as a first raw approximation of the long run impact on prices of slope measures. Once again, to be consistent with our discussion above, we take the signs of these sums into account when comparing the long-run effects of slope measures of different sides. In other words, if the sum of the estimated coefficients on all five lags of a given pair of the ask- and bid-side slope measures are positive and negative, respectively, we then test the equality of their relative magnitudes instead of their values.

Taking this into account, our third set of testable hypotheses can then be written as follows:

$$\begin{aligned}
H1c & : \sum_{\tau=1}^5 \alpha_{z,\tau}(S_{1,5,t}^A) = (1 - 2I_{1c}) \sum_{\tau=1}^5 \alpha_{z,\tau}(S_{1,5,t}^B); \\
H2c & : \sum_{\tau=1}^5 \alpha_{z,\tau}(S_{5,20,t}^A) = (1 - 2I_{2c}) \sum_{\tau=1}^5 \alpha_{z,\tau}(S_{5,20,t}^B); \\
H3c & : \sum_{\tau=1}^5 \alpha_{z,\tau}(S_{1,5,t}^A) = \sum_{\tau=1}^5 \alpha_{z,\tau}(S_{5,20,t}^A); \\
H4c & : \sum_{\tau=1}^5 \alpha_{z,\tau}(S_{1,5,t}^B) = \sum_{\tau=1}^5 \alpha_{z,\tau}(S_{5,20,t}^B);
\end{aligned}$$

where $I_{1c} = 1_{\{\sum_{\tau=1}^5 \alpha_{z,\tau}(S_{1,5,t}^A) > 0, \sum_{\tau=1}^5 \alpha_{z,\tau}(S_{1,5,t}^B) < 0\}}$ and $I_{2c} = 1_{\{\sum_{\tau=1}^5 \alpha_{z,\tau}(S_{5,20,t}^A) > 0, \sum_{\tau=1}^5 \alpha_{z,\tau}(S_{5,20,t}^B) < 0\}}$.

Finally, we also consider the impulse response functions of returns to test for the asymmetries in the long-run effects of different slope measures on prices. As mentioned above, the sum of coefficients on all lags of a given slope measure might not reveal the long-run cumulative effect after several trades given that slope measures have significant dynamics of their own. Impulse response functions take the dynamics of slope variables into account and, thus, capture the long-run effect better than the simple sum of coefficient estimates. Once again, when comparing the impulse response functions to slope measure of different sides, we take their signs into account. If the impulse response function to given ask- and bid-side slope measures at a given transaction period following a shock are positive and negative respectively, then we compare their relative magnitudes.

6 Estimation Results

In this section, we discuss the estimation results on each slope variable separately and do not make any comparisons between estimation results for different slope measures. Rather, we present these comparisons in the next section, together with our hypothesis test results.

As mentioned above, we present detailed estimation results for ALV and selected results for all other stocks. Table 3 presents the coefficient estimates on slope variables in the return equation for ALV.

[Insert Table 3 here]

We start our discussion with the effects of the ask-side slopes on prices. The coefficient estimates on the first lags of the ask-side slopes based on lower and higher levels are both significantly positive, suggesting that the immediate effect of the ask-side slope on returns is significantly positive regardless of the levels used to measure it. These results are significant both statistically and economically. To understand the economic magnitude of these coefficient estimates, one needs to take into account the fact that we consider the logarithm of slope measures in the VAR. For example, the log stock price increases by 0.16 basis points ($0.16b.p. = 0.259 \times 0.631$) at the next trade, following a

one standard deviation increase in the slope of the ask-side between the first and fifth levels. This is a considerably significant effect given that the mean log return between two transaction periods is practically zero.

In contrast, the coefficient estimates on higher lags of the ask-side slopes are mostly negative with differing levels of significance. This suggests a certain degree of reversal in the positive immediate effect of the ask-side slope. However, the coefficient estimates on higher lags are smaller in magnitude compared to the immediate effect. Thus, the sums of the coefficient estimates on all lags of both ask-side slope measures remain significantly positive, providing preliminary empirical evidence that the ask-side slope has a positive long run effect on prices, regardless of levels used to measure it.

The F-statistics show that coefficient estimates on all lags are jointly significantly different from zero. This in turn has two important implications. First, different measures of the ask-side slope not only have significant immediate effects on returns but also significant overall effects when one takes further lags into account. Second, these F-statistics can be interpreted as tests of causality in the sense of Granger (1969), suggesting that different measures of the ask-side slope cause prices to change.

Figure 2 presents the impulse response functions of returns to the different slope measures for ALV. The ask-side slopes have positive long-run cumulative effects on prices, regardless of the levels used to measure them, in line with approximation based on the sum of the coefficients. Furthermore, the long-run effects of the different slope measures for ALV have the same sign as their immediate effect, and are significantly different from zero for all horizons at the 5% level. To understand the magnitude of these long-run effects, note that the impulse response functions can be economically interpreted as the change in the log price of the stock in basis points as a function of transaction periods in response to a one-standard deviation positive shock to the slope variable of interest. For example, the log stock price of ALV increases by more than 0.6 basis points after 100 transactions following a one-standard deviation positive shock to the ask-side slope based on lower levels. After about 40 transactions, this long-term effect of the ask-side slope based on lower levels is almost fully incorporated into prices, while it takes slightly longer for the effect of the ask-side slope based on higher levels to be fully incorporated.

[Insert Figure 2 here]

We now turn our attention to the effects of the bid-side slopes on prices. The results are similar to those for the corresponding ask-side slopes but with opposite signs. Briefly, the coefficient estimate on the first lag of the bid-side slope is significantly negative regardless of the levels used to measure it, suggesting a negative impact response of returns to changes in the bid-side slope. This is also economically significant, with the log stock price decreasing by 0.17 basis points ($-0.17b.p. = -0.268 \times 0.631$) at the next trade, following a one-standard deviation increase in the slope of the bid-side between the first and fifth levels. The coefficient estimates on higher lags of the bid-side slopes are mostly positive, suggesting a degree of reversal. However, the sum of the coefficient estimates on all lags remains negative, providing preliminary evidence on the negative long-run effect on return of an increase in different measures of the bid-side slope. The F-statistics suggest that different measures of the bid-side slope also have significant overall

effects and cause prices to change in the sense of Granger (1969). The bid-side slopes have significantly negative long-run cumulative effects on prices, regardless of the levels used to measure them, and it takes about 50 to 60 transactions for these effects to be fully incorporated in prices.

Although estimation results exhibit variability depending on the stock, the results presented for ALV in Table 3 and Figure 2 are quite representative of what we obtain for other stocks. For the sake of brevity, we present a selected set of results for all other stocks in our sample, instead of full estimation results. Table 4 presents the coefficient estimates on the first lags of different slope variables, i.e. $\alpha_{z,1}$, the sum of the coefficient estimates on all lags, $\sum_{\tau=1}^5 \alpha_{z,\tau}$, and the F-statistic testing whether the the coefficient estimates on all lags are jointly equal to zero.

[Insert Table 4 here]

We start our discussion again with the ask-side slope measures. First of all, the Bonferroni p-values, which provide a very conservative way of testing the null hypotheses jointly across all stocks, suggest that the ask-side slopes have significantly positive immediate and long-run effects on prices of all stocks in the DAX30 index when considered jointly. This is also true for almost all stocks when considered separately. Second, the results for the ask-side slope based on lower levels are quite similar for different stocks. The results for the ask-side slope based on higher levels exhibit more heterogeneity across stocks. For example, when considered separately, the immediate and long-run effects of the ask-side slopes based on higher levels are not significantly different from zero for three and eleven stocks, respectively. Furthermore, the long-run effect of the ask-side slope based on higher levels is negative for several stocks and significantly so for few stocks such as DTE and MUV2.

[Insert Figure 3 here]

Figure 3 presents the impulse response functions of returns to slope measures. Although we computed the confidence bands for these impulse response functions as discussed in Section 4, we do not present them in Figure 3 for ease of exposition. Below we summarize our findings based on the confidence bands. The ask-side slope based on lower levels has significantly positive long-run effects on prices for all stocks in our sample. It takes about 40 transactions for the long-term effect of the ask-side slope based on lower levels to be fully incorporated into prices. The impulse response function of returns to the ask-side slope based on higher levels is significantly positive for all the transaction periods considered and most stocks in our sample. There are a few exceptions, such as DPW, DTE, LIN and MUV2, for which these impulse functions are negative in the medium- and long-terms after about 20 transaction periods following a shock. However, this is significantly so for only in the medium term for DTE. Furthermore, it also takes longer for the long-term effect of the ask-side slope based on higher levels to be fully incorporated into prices. There are even few stocks for which the long-term effect is not fully incorporated into prices after 100 transactions.

Turning our attention to the bid-side variables reveals similar results, again with opposite signs. Briefly, the bid-side slopes have significantly negative immediate and long-run effects for all stocks when considered jointly and for most stocks when considered separately. There are few exceptions. For example, the long-run effect of the bid-side

slope based on higher levels is significantly positive for DPW and LIN. Furthermore, the impulse response functions of returns to the bid-side slope based on lower levels are significantly negative for all horizons and stocks, and it takes about 40 transactions for the full effect of a shock to be realized. The impulse response functions of returns to the bid-side slope based on higher levels are significantly negative for most stocks. There are, once again, few exceptions. For example, the impulse response functions are positive in the medium to long term for DB1, DPW and LIN and in the short to medium term for DTE. However, the impulse response function is significantly positive only for DB1 after about 60 transaction periods following a shock.

7 Asymmetries in the Short-run Effects of Limit Order Book Slopes on Prices

In this section, we first discuss our findings on the asymmetries between immediate effects of different slope measures before turning our attention to the asymmetries between their overall short-run effects.

Table 5 presents the F-statistics for H1a, H2a, H3a, and H4a as well as the number of stocks for which we reject the null hypotheses at the 5% level and Bonferroni p-values. We start our discussion with the results for hypotheses H1a and H2a. When we consider the empirical evidence across all stocks based on Bonferroni p-values, we find that the ask-side slope has a significantly different immediate effect on prices than the bid-side slope does, regardless of the levels used to measure them. When the empirical evidence is considered separately for each stock, the immediate effects of the ask- and bid-side slopes based on lower (higher) levels are significantly different from each other for nine (seven) stocks. Furthermore, the immediate effects of the ask- and bid-side slopes based on higher levels for few stocks such as DTE, FME and LIN are significantly different from each other, whereas this is not true for their slope measures based on lower levels.

[Insert Table 5 here]

These results point to strong asymmetries between the immediate effects on price dynamics of the bid- and ask-side slopes regardless of the levels used to measure them. However, the evidence in support of any predictions regarding the relative magnitudes of these immediate effects discussed in Section 3 is rather mixed. To be more precise, the coefficient estimates in Table 4 show that the bid-side slope based on lower levels has significantly stronger, i.e. greater in absolute value, immediate effects than the corresponding ask-side slope for 18 of the 30 stocks in our sample, and significantly so for four stocks. Of course, the opposite holds for the other 12 stocks in our sample and significantly so for five of them. Similar results hold for the relative magnitudes of the bid- and ask-side slope measures based on higher levels.

We turn our attention to the results for our hypotheses H3a and H4a. When considered jointly across all stocks based on Bonferroni p-values, slope measures of the same side have significantly different immediate effects on prices depending on the levels used to measure them. This is also true for most of the stocks in our sample when we consider

the empirical evidence for each stock separately. Furthermore, these results along with the coefficient estimates in Table 4 suggest that slope measures based on higher levels have significantly stronger (greater in absolute value) immediate effects on prices slope measures based on lower levels. However, the coefficient estimates in Table 4 also suggest that slope measures of the same side based on different levels have immediate effects of the same sign. This is in contrast with the predictions based on Goettler et al. (2009) and more in line with those based on Rosu (2009) and Foucault et al. (2005) as discussed in Section 3.

We now consider the asymmetries between the overall dynamic effects of different slope measures. Table 6 presents the F-statistics corresponding to hypotheses H1b, H2b, H3b, and H4b. These results provide strong empirical evidence in support of asymmetries not only between the overall effects of slope measures of different sides based on the same levels but also between the overall effects of slope measures of the same side based on different levels. Several remarks regarding our results for hypotheses H1b and H2b are in order. In H1b and H2b, as discussed in Section 5, we consider the relative magnitudes of a given pair of coefficient estimates only if they are positive for the ask-side and negative for the bid-side slope measures. To be more precise, assume that the estimated coefficients on the first lags of the ask- and bid-side slope measures for a given stock are positive and negative, respectively, while the opposite is true for the coefficient estimates on all other lags. This corresponds to a reversal in the immediate effect of slope measures, as discussed in Section 6. In H1b and H2b, we would then simply test the equality of the coefficient estimates on all lags except those on the first lags, for which we would consider their relative magnitudes. This in turn renders the results more significant than otherwise because we test the equality of coefficient estimates with opposite signs. Given that some reversal of the immediate effect is observed for most of the stocks in our sample, this in turn explains the highly significant results for H1b and H2b. In unreported analysis, we tested the equality of the relative magnitudes of the coefficient estimates regardless of their signs and found that the absolute values of the coefficient estimates on all lags of the ask- and bid-side slope measures based on lower (higher) levels are significantly different from each other for thirteen (nine) stocks in our sample.

[Insert Table 6 here]

Overall, our results point to strong asymmetries in the short-run effects of different slope measures on price dynamics. The empirical evidence for the asymmetries in the short-run effects is much stronger between slope measures of the same side based on different levels. More importantly, the short-run effects of slope measures of the same side based on different levels have the same sign. However, we do not find much empirical evidence supporting the prediction that the bid-side slope has a stronger effect on prices than does the ask-side slope based on the same levels.

8 Asymmetries in the Long-run Effects of Limit Order Book Slopes on Prices

In this section, we focus on the asymmetries in the long-run effects of different slope measures on future price dynamics. We start with our results based on the sum of the estimated coefficients on all lags of slope measures before analyzing the impulse response functions of returns to slope measures.

Table 7 presents the F-statistics corresponding to our hypotheses H1c, H2c, H3c and H4c. When we consider the empirical evidence for H1c jointly across all stocks, the long-run effects of the ask- and bid-side slopes based on lower levels are significantly different from each other in magnitude. When we consider the empirical evidence for each stock separately, the long-run effects of the ask- and bid-side slopes based on the lower levels are significantly different from each other for four stocks. This provides some, although not very strong, empirical evidence of asymmetries between the long-run effects of the ask- and bid-side slopes based on lower levels. Nonetheless, the results in Table 4 show that the sum of the coefficient estimates for the ask-side slope based on lower levels is greater in magnitude than that for the corresponding bid-side slope for half of the stocks in our sample and significantly so for only two stocks, while the opposite holds for the other half and significantly so for two stocks. We find similar results for our hypothesis H2c. More precisely, the long-run effects of the ask- and bid-side slopes based on higher levels are significantly different from each other in magnitude when considered jointly across all stocks, and significantly so for six stocks when each stock is considered separately. The results in Table 4 show that the sum of the coefficient estimates for the ask-side slope based on higher levels is greater in magnitude than that for the corresponding bid-side slope for 12 stocks and significantly so for only two stocks, while the opposite holds for the remaining 18 stocks, and significantly so for four stocks.

[Insert Table 7 here]

Results for H3c and H4c reveal strong empirical evidence of asymmetries between the long-run effects of slope measures of the same side based on different levels. To be more precise, the F-statistics for H3c and H4c show that the long-run effects of slope measures based on lower levels, regardless of the side considered, is significantly different from those based on higher levels for all stocks when considered jointly and for almost all stocks when the empirical evidence for each stock is considered separately. Results in Table 4 suggest that the long-run effects of the ask- and bid-side slopes are positive and negative, respectively, for most stocks in our sample regardless of the levels used to measure them. That said, there is evidence, albeit weak, based on these long-run effects in support of the predictions consistent with Goettler et al. (2009), unlike the empirical evidence, or lack thereof, based on immediate effects. Specifically, the ask-side slope based on higher levels has a negative long-run effect on prices for nine stocks and significantly so for two stocks, while the bid-side slope based on higher levels has a positive long-run effect for six stocks and significantly so for two stocks. More importantly, the empirical evidence regarding the relative magnitudes of these long-run effects is much clearer than the empirical evidence regarding the relative magnitudes of

their immediate effects. More precisely, slope measures based on lower levels have long-run effects on prices that are much stronger (in absolute value) than slope measures based on higher levels for all stocks in our sample, regardless of the side considered. In an unreported analysis, we find that this is also significantly so for almost all stocks in our sample for both the bid and ask sides.

We now turn our attention to the asymmetries between the impulse response functions of returns to different slope measures. For different pairs of slope measures, Figure 4 presents the number of stocks for which the impulse response function of returns to one of the slope measures is (significantly) greater than that to the other slope measure as a function of transaction periods. We start with the asymmetries between the impulse response functions to slope measures based on the same levels but different sides. As discussed in Section 5, we compare the relative magnitudes of these impulse response functions if they are positive for the ask-side slope measures and negative for the bid-side slope measures at a given transaction period. The impulse response function to the ask-side slope is greater than that to the bid-side slope for half of the stocks, regardless of the levels used to measure them, while the opposite holds for the other half. More importantly, the difference is statistically significant for at most one stock in our sample. This in turn suggests that there is no empirical evidence for any asymmetry between the long-run effects of different sides based on the same levels.

[Insert Figure 4 here]

In contrast, there is ample empirical evidence of asymmetries in the long-run effects of slope measures based on different levels of the same side, especially in the medium-run between 5 to 40 transaction periods following a shock. Furthermore, the impulse response functions presented in Figure 3 show that the slope measures of the same side based on different levels almost always have the same sign, with only a few exceptions. To be more precise, the impulse response function of returns to the ask-side slope based on lower levels between 5 to 40 transaction periods following a shock is greater than that to the ask-side slope based on higher levels for more than 20 stocks and significantly so for more than 5 stocks. Although these numbers decrease in the long run, there are about 20 stocks for which this holds and significantly so for about 5 stocks. The opposite holds for about 10 stocks, and significantly so for few stocks but only in the short run and not at all in the long run. The evidence is somewhat stronger for the asymmetries between the long-run effects of the bid-side slope measures based on different levels. The impulse response function of returns to the bid-side slope based on lower levels between 5 to 40 transaction periods following a shock is smaller than that to the bid-side slope based on higher levels for more than 20 stocks and significantly so for more than 10 stocks. These numbers also decrease in the long run but there are about 20 stocks for which this holds, and significantly so for about 5 stocks. Once again, there is very limited evidence that the opposite holds significantly.

Overall, our results on the asymmetries between the long-run effects of slope measures are similar to those on the asymmetries between their short-run effects. Although we find empirical evidence for asymmetries in the long-run effects of slope measures of different sides based on the same levels, it is not very strong. Furthermore, there is not much consistent empirical evidence supporting the predictions on the relative magnitudes of the long-run effects of

these slope measures. We also find strong evidence for asymmetries in the long-run effects of slope measures of the same side based on different levels. More importantly, the long-run effects of slope measures of the same side based on different levels have the same sign.

9 Robustness Checks

In our main set of results, we distinguish between lower and higher levels of the limit order book based on the first five and fifth to twentieth levels, respectively. In this section, we consider four alternative definitions of lower and higher levels: (1) we change the cutoff level to ten and define lower and higher levels based on the first ten and tenth to twentieth levels, respectively; (2,3) we then exclude the first level from the definition of lower levels and define it to be levels between the second and fifth (or tenth), and we define higher levels as levels between the fifth (or tenth) and twentieth; (4) we finally change the cutoff level to two and define lower and higher levels based on the first two and second to twentieth levels, respectively. Furthermore, we have three months of data available: July 2010, May 2011 and June 2011. In our main set of results, we choose to focus on a single month, namely June 2011, due to the fact it is similar to the average of these three months in terms of returns, volatility and trading activity. We now show that our hypothesis test results are indeed robust to using alternative definitions of lower and higher levels and to using alternative sample periods.

[Insert Table 8 here]

Table 8 presents the number of stocks for which we reject H1a to H4c as well as the Bonferroni p-values for testing these hypotheses jointly across stocks. We find that our conclusions based on our main sets of results remain unchanged regardless of the sample period and the definition of lower and higher levels. To be more precise, we reject most of the null hypotheses when we consider the empirical evidence jointly across all stocks based on Bonferroni p-values. This in turn suggests that there are statistically significant asymmetries between the effects of different slope measures on price dynamics. More importantly, when we consider the empirical evidence for each stock separately, the empirical evidence is much stronger for the asymmetries in the effects of slope measures based on different levels of the same side compared to the asymmetries in the effects of slope measures based on the same levels of different sides. The slope measures based on different levels of the same side have significantly different short- and long-run effects on price dynamics for about 25 out of 30 stocks on average while the same is true for about 8 stocks when we compare the effects of slope measures based on the same levels of different sides. Although not presented, we find similar empirical evidence to that presented in Section 8 when we analyze the impulse response functions of returns to different slope measures based on alternative definitions and sample periods. More importantly, our results regarding the signs and relative magnitudes of these effects are very similar to those discussed in Sections 7 and 8.

10 Economic Value of Asymmetric Effects

Our empirical results discussed so far provide statistically significant evidence in support of certain asymmetries between the effects of different slope measures on price dynamics. In this section, we show that these asymmetries can also be economically significant. We do this by comparing the performances of high-frequency day-trading strategies that ignore the information embedded in different types of asymmetries to that of an unrestricted strategy that uses this information. To be more precise, our unrestricted strategy employs the return equation of the VAR system, Equation (5a), to forecast midquote returns at each transaction period while the competing strategies employ restricted versions of this equation so that a chosen pair of slope variables has symmetric effects on price dynamics.⁹

To be consistent with the idea of high-frequency trading, we consider the possibility of trading at every transaction period and, thus, focus on the economic value of asymmetries in the short-run, rather than long-run, effects of different slope variables. To this end, we impose the restrictions implied by hypotheses 1a-4a and 1b-4b one at a time on Equation (5a) and use this version as our forecasting models in these so-called restricted strategies. The restricted strategies thus ignore the information embedded in a certain type of asymmetry in forecasting price movements. Comparing the performances of the restricted strategies and the unrestricted strategy allows us to evaluate the economic value of different asymmetries. To give a more concrete example, in the first restricted strategy, we restrict the coefficients on the first lags of the ask and bid slope measures based on lower levels to be identical. This forces the ask- and bid-side slope measures based on lower levels to have symmetric immediate effects on prices while allowing for all other potential asymmetries in the effects of slope measures on prices. If the performance of the unrestricted strategy is higher than that of the restricted strategy, we then argue that the asymmetry between the immediate effects of the ask- and bid-side slope measures based on lower levels is economically important.

The trading strategy we consider is similar to that discussed in Kozhan and Salmon (2012) and can be summarized as follows: At each transaction period t for a given stock, we take a snapshot of the limit order book right after (less than a millisecond after) observing the transaction in line with our empirical model discussed in Section 4. We then compute the forecast of the midquote return in the next transaction period $t + 1$, \hat{r}_{t+1} , based on this snapshot and a given forecasting model and reevaluate our existing position if the return forecast is either greater than a nonnegative threshold, κ , or less than $-\kappa$. To be more precise, we consider a forecast greater than κ , i.e. $\hat{r}_{t+1} > \kappa$, to be a buy signal and do one of the following depending on our existing position in the stock: (1) buy one share of the stock if we do not already have an existing position in the stock; (2) buy two shares of the stock if we have an existing short position in the stock, i.e. close the short position and take a long position of one share; (3) do nothing if we already have a long position. Similarly, we consider a forecast less than $-\kappa$, i.e. $\hat{r}_{t+1} < -\kappa$, to be a sell signal and do one of the following depending on our existing position in the stock: (1) short-sell one share of the stock if we do not

⁹In all these forecasting models, we choose to exclude the contemporaneous effect of trade direction on returns. As discussed above, any contemporaneous effect of the trade direction on returns is not a causal relation, even if the trade direction is observed right before the return is calculated. Nevertheless, we also considered trading strategies based on forecasting models that include the contemporaneous effect of trade direction on returns. Our results based on this VAR model are similar to those based on the restricted VAR model and, thus, are not presented here in detail but available upon request.

already have an existing position in the stock; (2) short-sell two shares of the stock if we have an existing long position in the stock, i.e. close the long position and take a short position of one share; (3) do nothing if we already have a short position. We implicitly assume that our trading does not alter the dynamics of the relation between returns, trade directions, and limit order book. We believe that this is a reasonable assumption because we consider trading at most two shares at a time and holding a short or long position of one share at any point in time. The effect of our trading and position should be negligible given that the trading volume and limit order book depth of stocks in our sample are quite large. We also close any existing position at the end of each trading day and, thus, do not hold an overnight position, in line with the idea of day trading. We evaluate the performance of each strategy based on its average cumulative daily return.

Several remarks are in order regarding our main set of results, which are presented in Table 9. First, as in Kozhan and Salmon (2012), κ can be considered a parameter to filter out potentially weak signals, and one trades less frequently as higher values of κ are considered. In our main set of results, we consider a κ of zero. In other words, we only consider the sign, not the strength, of our forecasts. This allows us to analyze the economic value of different asymmetries in forecasting the direction of prices. In robustness checks, we also consider alternative values for κ and, thus, take the strength of our forecasts into account in our trading strategies. Second, our main empirical results are based on the assumption that we can trade at midquote prices, instead of the more realistic assumption of trading at the best bid and ask prices. We make this assumption mainly to be consistent with our empirical models, which forecast midquote returns. This assumption also allows us to analyze the economic value of the asymmetries without having to worry about transaction costs. Nevertheless, we test the robustness of our results to trading at the best bid and ask, rather than midquote, prices. Third, we calculate our trading signals based on the snapshots of the limit order book right after (less than a millisecond after) a transaction and assume that we can trade at the prices in that snapshot of the limit order book. However, in reality, we can neither observe a transaction (and the snapshot of the limit order book right after it) instantaneously nor trade at prices observed in this snapshot of the limit order book. This is due to delays in different legs of the trading process where information flows from the exchange to the trader's system and is processed by the trader and an order is then sent from the trader's system to the exchange. Thus, a trader will trade based on information and prices different from those in the observed snapshot of the limit order book. Ideally, one would like to examine the effects of these two legs of the information delay on the economics values of asymmetries separately. However, this is relatively difficult. Instead, we take a very simplistic approach and assume that we can only trade at prices observed in the snapshot of the limit order book 500 milliseconds after a transaction. Finally, our main results are in-sample results in the sense that we use the same data period to estimate parameters of forecasting models and to test the relative performances of trading strategies based on these forecasting models.

[Insert Table 9 here]

Table 9 presents the differences between the average daily cumulative returns of the unrestricted and restricted trading strategies. A positive number implies that the unrestricted strategy provides, on average, higher daily cumulative

returns than the trading strategy that uses the restricted forecasting model implied by hypotheses in the corresponding column heading. In other words, a positive number suggests that the information embedded in a given asymmetry is economically important. Overall, our results suggest that the eight short-run asymmetries considered are all economically important since return differences are positive for most of the stocks in our sample. The Bonferroni p-values of the positive return differences suggest that they are also statistically significant for all asymmetries at the 5% significance level when the empirical evidence is considered jointly. They are also statistically significant for between 9 and 14 stocks depending on the asymmetry when the empirical evidence for each stock is considered separately.

The evidence is strongest for the asymmetries between the overall short-run effects of slope measures based on different levels of the same side. To be more precise, the unrestricted strategy provides a higher average daily cumulative return than the strategies restricting the overall short-run effects of slope measures based on different levels of the same side to be the same, i.e. strategies imposing the restrictions implied by H3b and H4b, for more than 25 stocks. Averaged over all stocks, the differences are, respectively, about 25 and 24 basis points between the unrestricted strategy and the strategies imposing the restrictions implied by H3b and H4b. A difference of 25 basis points in daily returns is economically quite important given that the average daily return for all stocks in DAX 30 in June 2011 ranges between -47 and 65 basis points with an overall average of one basis point (see Table 2). Furthermore, the return differences are statistically significant for 10 to 14 stocks with corresponding Bonferroni p-values of 0.001 and 0.014. These are in line with our hypothesis test results discussed in Section 7, which suggest that the empirical evidence for the asymmetries between the effects of slope measures based on different levels of the same side is relatively stronger than those between the effects of slope measures of different sides based on the same levels.

10.1 Robustness Checks

In this section, we discuss the robustness of our main empirical results to making alternative empirical choices.

Trading at Different Thresholds

We first show that asymmetries continue to be economically important when we take into the account not only the sign but also the strength of forecasts by considering alternative values of the threshold parameter κ . For each stock, we consider values of κ that correspond to 10%, 20%, 30%, 40% and 50% of the standard deviation of its tick-by-tick returns in June 2011. As mentioned above, κ can be considered a parameter to filter out potentially weak signals, and one trades less frequently as higher values of κ are considered. To be more precise, we can trade potentially at each observed transaction period if we completely ignore the trading signal. For example, in our main set of results where we set κ to zero, we trade, on average, approximately 30% to 35% of the time, i.e. 30% to 35% of all transaction periods, for a given stock in a given day. On the other hand, when we set κ to 50% of the corresponding standard deviation, we trade less than 1% of the time. Thus, choosing higher values of κ also allows us to analyze the effect of the trading frequency on our results in addition to the effect of jointly considering the sign and strength of our

forecasts. In unreported analysis, we find that the number of trades indeed decreases on average when we consider higher values of κ , as expected. Furthermore, the absolute performances of all trading strategies, also not presented, decrease as we trade less and less frequently at higher values of κ .

[Insert Table 10 here]

Table 10 presents the difference between the daily cumulative returns of the unrestricted strategy and the trading strategy that uses the restricted forecasting model implied by hypotheses in the corresponding column heading, averaged over all stocks and trading days in June 2011. All return differences are positive and statistically significant at 5%, suggesting that the eight short-run asymmetries considered continue to be economically important when we consider alternative values of κ . There is a small tendency for the return differences to increase as we consider higher values of κ , suggesting that taking both the strength of the signal and its sign provides the day trader with further economic profits. Finally, the evidence seems to be strongest for the asymmetries between the immediate effects of ask- and bid-side slope measures based on lower levels. Nevertheless, the evidence for the asymmetries between the overall short-run effects of slope measures based on different levels of the same side remains strong when we consider higher values of κ .

Trading at the Best Bid and Ask Prices

We now consider the effect of trading at the best bid and ask prices corresponding to the direction of the required position. Before presenting their relative performances when we trade at the best bid and ask prices, several remarks are in order regarding the absolute performances of different strategies. Although we do not present our results for the sake of brevity, we find that daily cumulative returns of the unrestricted and restricted strategies are all, on average, positive when we assume that we can trade at midquote prices. However, these profits all become, on average, quite negative when we assume that we trade at the best bid and ask prices. This in turn suggests that the positive profits of all strategies are wiped out by transaction costs, in line with the findings in Kozhan and Salmon (2012).

[Insert Table 11 here]

Nevertheless, Table 11 shows that the return differences between the unrestricted and restricted strategies are positive not only on average but also for most firms in our sample when we trade at the best bid and ask prices with a κ of zero. This is significantly so for all asymmetries at the 5% significance level when the empirical evidence for all stocks is considered jointly and for seven to seventeen stocks depending on the asymmetry considered when the empirical evidence for each stock is considered separately. In unreported analysis, we also find that these results mostly hold when we consider higher values of κ up to 25% of the standard deviation of tick-by-tick returns for a given stock.

Trading at a Slower Speed

As mentioned above, in our main sets of results, we assume that we can observe a transaction and the associated snapshot of the limit order book and trade at prices in this snapshot without any delay. However, any trader who wants to implement such a trading strategy faces several delays in reality. First, the trader observes the transaction and the associated snapshot of the limit order book with a delay, which corresponds to the time it takes for this information to travel from the exchange to the trader's system. Second, whether the trader is a human or a computer, it takes time to process this information. Finally, if the trader decides to trade based on this information, it takes time for the order to arrive at and to be executed by the exchange's system.

With each delay, the market conditions can undoubtedly change and the trader's information can become stale quite rapidly in today's extremely fast markets. As a result, the trader would make decision based on stale information. Ideally, one should distinguish between the first leg of the trading process during which the information travels from the exchange's systems to the trader's system (and is processed by the trader) and the second leg, during which the information flows in the other direction. It would be intriguing to analyze the effects of these two legs of the information delay on the economic values of asymmetries separately, yet this is relatively difficult because these delays are random and can change significantly from trader to trader depending on the physical proximity of the trader to the exchange and the information processing speeds.

We thus adopt an approach similar to the one discussed in Kozhan and Tham (2012), which can be considered as a first approximation to the effect of these delays on the economic value of asymmetries. Specifically, we assume that there is no delay in the first leg of the trading process and that the trader can observe a transaction (and the associated snapshot of the limit order book) and react to it instantaneously. The violation of this assumption in reality might make the trader change the forecasting model due to the associated delay but does not alter the content of the information observed. In the second leg of the trading process, we assume that there is a delay of 500 milliseconds. In other words, once a trader places an order, it takes 500 milliseconds to be executed by the exchange's systems. Along with our first assumption, this implies that the orders will be executed at prices in the snapshot of the limit order book observed exactly 500 milliseconds after a transaction. As mentioned in Section 4, the average latency on the Xetra system, i.e. the average time required for an order to travel from the trading participant's system across the network to its backend and for confirmation of its receipt to be sent back to the participant, is about 13 milliseconds. Thus, a delay of 500 milliseconds can be considered as relatively long compared to the average latency in the system. However, one might also consider this delay of 500 milliseconds as the total time for both legs of the trading process. Nevertheless, we also considered a delay of 100 milliseconds, and our results remain similar to those presented in Table 12.

[Insert Table 12 here]

Our results can be summarized as follows: In unreported analysis, we find that the absolute performance of all trading strategies when we trade with a latency of 500 milliseconds is much lower than that based on a latency of zero milliseconds, in line with our expectations. More importantly, Table 12 presents the differences between the

average daily cumulative returns of the unrestricted and restricted trading strategies when we trade at midquote prices with a κ of zero and a latency of 500 milliseconds, and shows that our main set of results is mostly robust to trading at slower speeds.¹⁰ For example, the number of stocks for which asymmetries are both economically important and statistically significant is similar to our main set of results. Furthermore, the evidence still seems to be strongest for the asymmetries between the overall short-run effects of slope measures based on different levels of the same side. However, as expected, the return differences are, on average, smaller than those in our main set of results, which is also reflected in the Bonferroni p-values of positive return differences.

11 Conclusion

In this paper, we analyze whether the information embedded in different parts of the limit order book affect short- and long-run price dynamics in line with the predictions of recent theoretical models. To this end, we reconstruct the first 20 levels of the historical limit order book every millisecond for all 30 stocks in the DAX30 index based on data from the Xetra electronic trading system. We then distinguish not only between the slopes of the bid and ask sides as has been done in the previous literature, but also between the slopes of the lower and higher levels of the book, which had not been considered previously. We then develop several hypotheses on potential asymmetries in the effects of different slope measures on future price dynamics based on recent theoretical literature. To test these hypotheses, we include the four slope measures as state variables in a linear vector autoregressive system that also includes midquote return and trade direction, similar to the Hasbrouck (1991) model. This empirical framework allows us to test the predictions of recent theoretical models for the immediate, short- and long-run effects of different slope measures on future price dynamics. To the best of our knowledge, this is the first paper to analyze the asymmetries in the short- and long-run effects of slope measures based on different levels.

We find that not only do slopes of different sides based on same levels affect short- and long-run price dynamics quite differently, but more importantly so do slopes of the same side based on different levels, mostly in line with the predictions of recent theoretical models. More specifically, the bid-side slopes have significantly negative short- and long-run effects on prices, which is (in absolute value) significantly different from the significantly positive effects of the corresponding ask-side slopes, regardless of the levels used to measure them. This is in line with our predictions based on Kalay and Wohl (2009), Foucault et al. (2005) and Rosu (2009). However, the empirical evidence in support of our predictions regarding the relative magnitudes of these effects is rather weak. To be more precise, we argue based on these models that the effect of the bid-side slope on prices should be greater in magnitude than that of the ask-side slope based on the same levels. We find this to be true for approximately half of the stocks in our sample, and significantly so for a handful of stocks. Slope measures based on higher levels have significantly stronger effects on prices than do slope measures based on lower levels of the same side. However, slope measures of the same side

¹⁰We also considered trading at the best bid and ask prices as well as at higher value of κ with a latency of 500 milliseconds. Our results remain qualitatively similar.

based on different levels have immediate effects of the same sign, in line with the predictions of Foucault et al. (2005) and Rosu (2009) but in contrast to those of Goettler et al. (2009). In a battery of robustness checks, we show that these results are robust to using alternative definitions and sample periods.


We also provide empirical evidence that these asymmetries are important not only statistically but also economically. We do this in a simple high frequency day trading exercise by comparing the performances of strategies that ignore the information embedded in different types of asymmetries to that of an unrestricted strategy that uses this information. In this framework, we show that the unrestricted strategy provides daily profits that are, on average, 25 basis points higher than each restricted strategy, suggesting that ignoring these asymmetries costs a trader approximately 25 basis points in daily profits. These results are robust to trading under alternative sets of assumptions such as trading at different signal thresholds, the best bid and ask prices instead of midquote prices and slower speeds.

References

- Beltran-Lopez, H., Giot, P., Grammig, J., 2009. Commonalities in the order book. *Financial Markets and Portfolio Management* 23, 209–242.
- Biais, B., Hillion, P., Spatt, C., 1995. An empirical analysis of the limit order book and the order flow in the Paris bourse. *Journal of Finance* 50, 1655–1689.
- Bilodeau, Y., 2013. Xetraparser. [computer software], HEC Montreal .
- Brogaard, J., Hendershott, T., Riordan, R., 2016. Price Discovery without Trading: Evidence from Limit Orders. Available at SSRN: <https://ssrn.com/abstract=2655927> .
- Cao, C., Hansch, O., Wang, X., 2009. The information content of an open limit-order book. *Journal of Futures Markets* 29, 16–41.
- Chakravarty, S., Holden, C. W., 1995. An integrated model of market and limit orders. *Journal of Financial Intermediation* 4, 213 – 241.
- Dufour, A., Engle, R. F., 2000. Time and the price impact of a trade. *Journal of Finance* 55, 2467–2498.
- Foucault, T., Kadan, O., Kandel, E., 2005. Limit order book as a market for liquidity. *Review of Financial Studies* 18, 1171–1217.
- Glosten, L. R., 1994. Is the electronic open limit order book inevitable? *Journal of Finance* 49, 1127–1161.
- Glosten, L. R., Milgrom, P. R., 1985. Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *Journal of Financial Economics* 14, 71–100.
- Goettler, R. L., Parlour, C. A., Rajan, U., 2009. Informed traders and limit order markets. *Journal of Financial Economics* 93, 67 – 87.
- Granger, C. W. J., 1969. Investigating causal relations by econometric models and cross-spectral methods. *Econometrica* 37, 424–438.
- Griffiths, M. D., Smith, B. F., Turnbull, D. A. S., White, R. W., 2000. The costs and determinants of order aggressiveness. *Journal of Financial Economics* 56, 65 – 88.
- Hamilton, J. D., 1994. *Time Series Analysis*. Princeton University Press .
- Hasbrouck, J., 1991. Measuring the information content of stock trades. *Journal of Finance* 46, 179–207.
- Hellwig, M. F., 1980. On the aggregation of information in competitive markets. *Journal of Economic Theory* 22, 477 – 498.

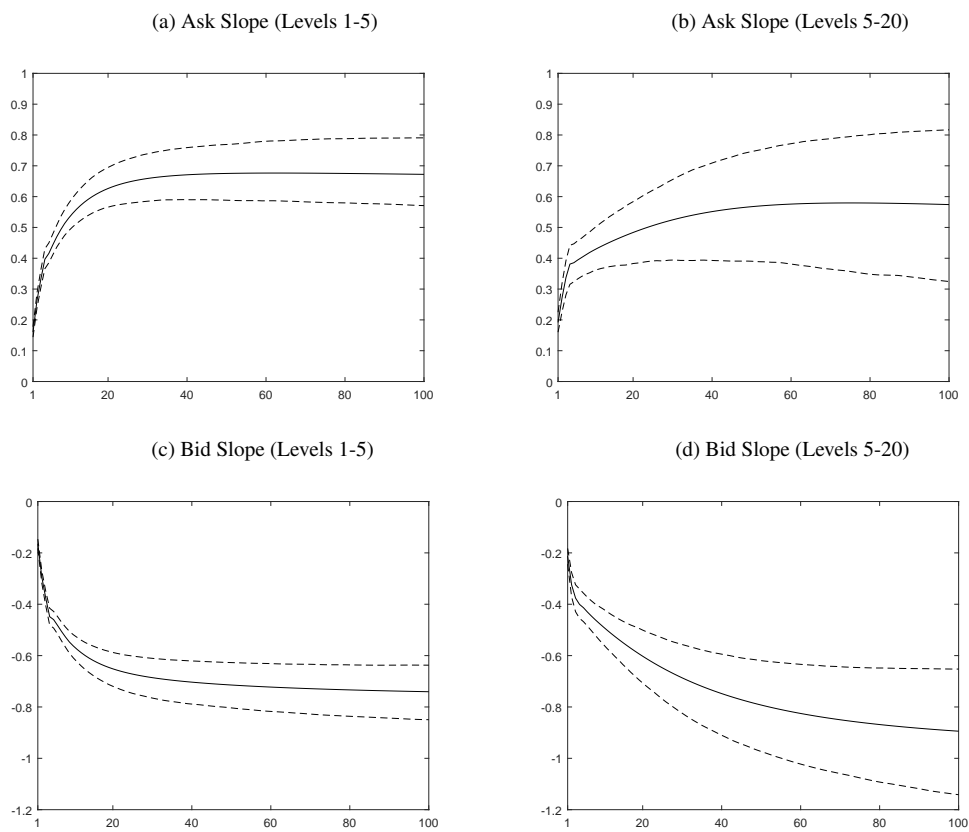
- Kalay, A., Wohl, A., 2009. Detecting liquidity traders. *Journal of Financial and Quantitative Analysis* 44, 29–54.
- Kaniel, R., Liu, H., 2006. So what orders do informed traders use? *Journal of Business* 79, 1867–1913.
- Kozhan, R., Salmon, M., 2012. The information content of a limit order book: The case of an FX market. *Journal of Financial Markets* 15, 1–28.
- Kozhan, R., Tham, W. W., 2012. Execution Risk in High-Frequency Arbitrage. *Management Science* 58, 2131 – 2149.
- Kumar, P., Seppi, D. J., 1994. Limit and market orders with optimizing traders. Working paper Carnegie Mellon University .
- Kyle, A. S., 1985. Continuous auctions and insider trading. *Econometrica* 53, 1315–1335.
- Newey, W. K., West, K. D., 1987. A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica* 55, 703–708.
- Rock, K., 1996. The specialist's order book and price anomalies. Working Paper, Harvard University .
- Rosu, I., 2009. A dynamic model of the limit order book. *Review of Financial Studies* 22, 4601.

Figure 1: Two Snapshots of the Limit Order Book for ALV on 1 June 2011



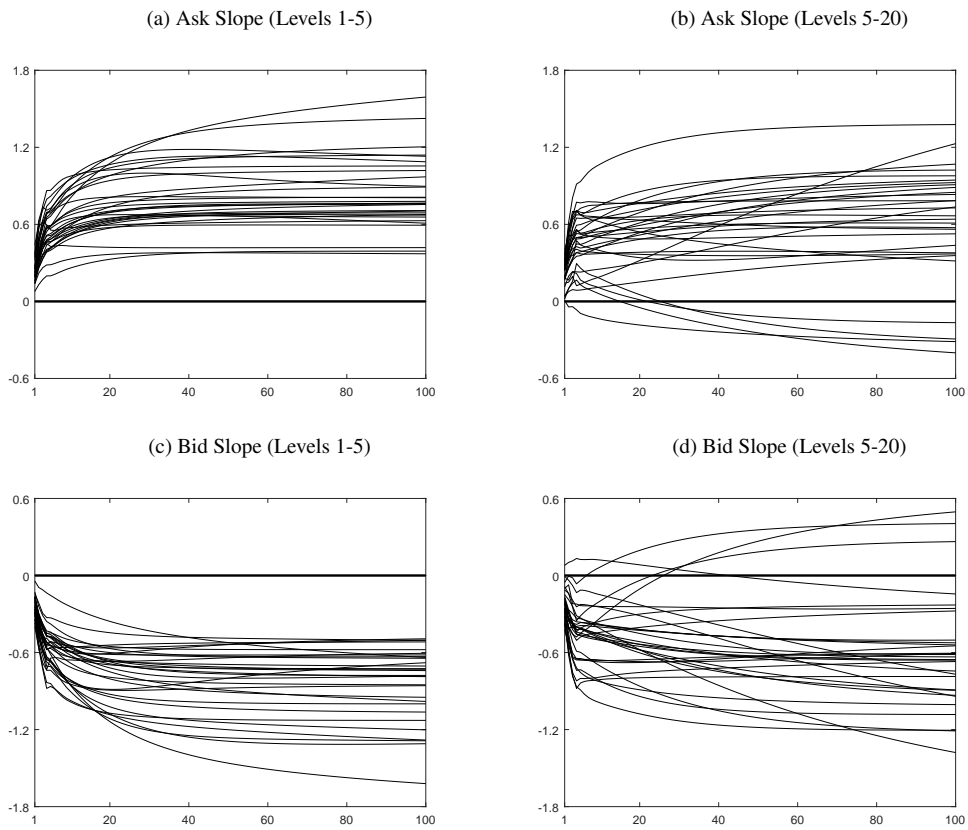
Note: Panel (a) and (b) present the price and (cumulative) depth of the first 20 levels of the limit order book for ALV on 1 June 2011 at 10:30:00.000 and 11:30:00.000, respectively. The diamonds and squares represent different levels of the ask and bid sides, respectively. The annotation displays the corresponding ask- and bid-side slope measures between the first and fifth levels and between the fifth and twentieth levels.

Figure 2: Impulse Response Function of Returns to Slope Measures for ALV



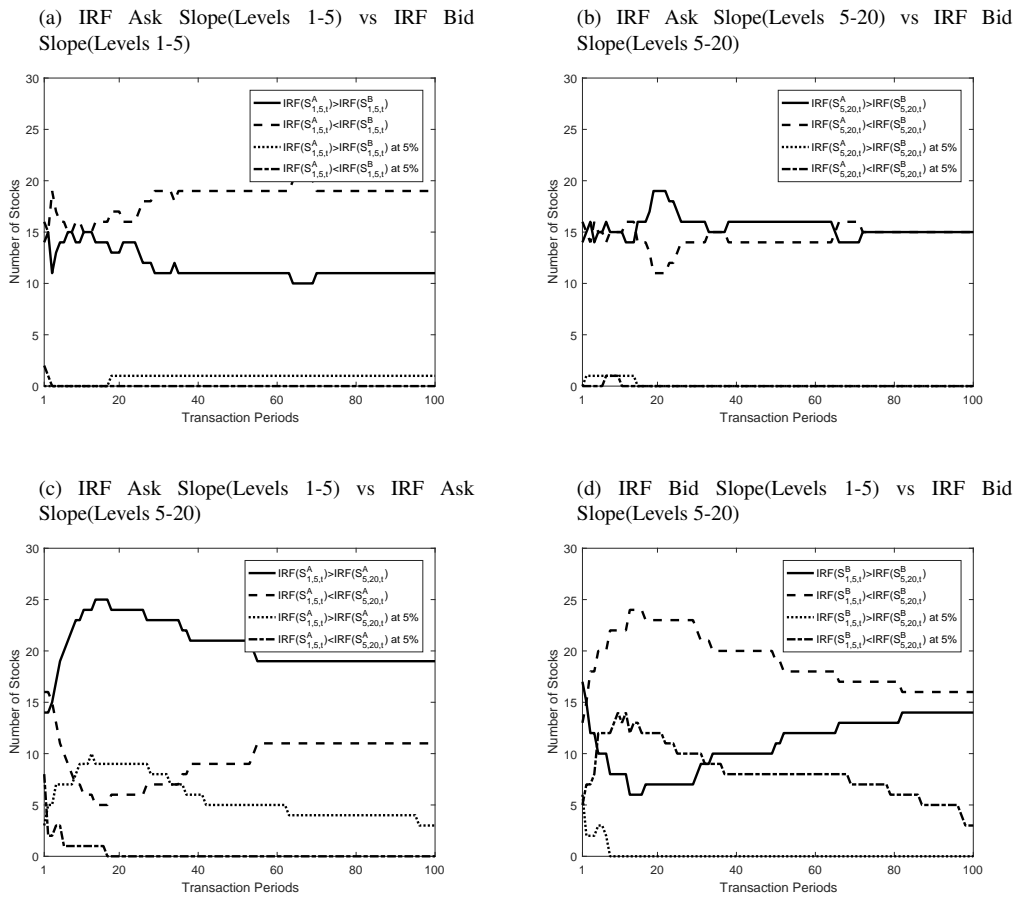
Note: This figure presents the impulse response functions of the midquote return to a one-standard deviation positive shock to slope variables for ALV in June 2011 and the corresponding 95% confidence bands (dashed lines). The horizontal axis represents the transaction periods, that is, the number of transactions since the initial shock, and the vertical axis represents the response of the midquote return in basis points. The impulse response function and its confidence bands are computed based on the simulation approach discussed in Hamilton (1994).

Figure 3: Impulse Response Functions of Returns to Slope Measures



Note: This figure presents the impulse response functions of the midquote return to a one-standard deviation positive shock to slope variables for all 30 stocks in the DAX30 index. The horizontal axis represents the transaction periods, that is, the number of transactions since the initial shock, and the vertical axis represents the response of the midquote return in basis points. The impulse response functions are computed based on the simulation approach discussed in Hamilton (1994).

Figure 4: Impulse Response Function of the Returns to Slope Measures



Note: This figure presents the number of stocks (out of 30 stocks in the DAX30 index) for which the impulse response function of returns to one of the slope measures is (significantly) greater than that to the other slope measure as a function of transaction periods. The horizontal axis represents the transaction periods, that is, the number of transactions since the initial shock, and the vertical axis represents the number of stocks. The solid lines present the number of stocks for which the impulse response function to the first slope variable in the corresponding title is greater than that to the second slope variable and the dashed lines present the number of stocks for which the opposite holds, without taking statistical significance into account. The dotted line presents the number of stocks for which the impulse response function to the first slope variable in the corresponding title is significantly (at the 5% level) greater than that to the second slope variable, while the dashed-dotted lines present the number of stocks for which the opposite holds significantly.

Table 1: Summary Statistics on Stocks in DAX 30 Index in June 2011

Ticker	Company Name	Avg. Market Capitalization (in billion Euros)	Avg. Daily Trading Volume (in million shares)	Avg. Daily Turnover (in percentage)	Avg. Daily Return (in percentage)	Std. Dev. of Daily Return (in percentage)
ADS	ADIDAS AG	10.752	1.046	0.500%	0.196%	1.705%
ALV	ALLIANZ SE	42.349	2.239	0.493%	0.007%	1.506 %
BAS	BASF SE	58.590	3.473	0.378%	0.228 %	1.729%
BAYN	BAYER AG	46.219	3.188	0.386 %	-0.122%	1.853%
BEI	BEIERSDORF AG	10.149	0.458	0.202%	-0.130%	0.815%
BMW	BAYER MOTOREN WERKE AG	38.454	2.741	0.455%	0.512%	1.293%
CBK	COMMERZBANK	15.654	58.946	1.153%	-0.312 %	2.624%
DAI	DAIMLER AG	51.318	5.052	0.474%	0.252 %	1.194%
DB1	DEUTSCHE BOERSE AG	10.249	1.165	0.598%	-0.202 %	1.055%
DBK	DEUTSCHE BANK AG	37.290	6.268	0.674 %	-0.074%	1.497%
DPW	DEUTSCHE POST AG	15.338	3.933	0.325 %	0.057%	1.039%
DTE	DEUTSCHE TELEKOM	44.743	13.803	0.319 %	0.209%	1.291 %
EOAN	E.ON SE	36.486	10.709	0.562%	-0.035 %	1.544%
FME	FRESENIUS MEDI. CARE AG&CO	15.059	0.646	0.217%	0.108 %	0.909%
FRE	FRESENIUS SE & CO KGAA	11.577	0.383	0.236 %	-0.062%	1.312 %
HEI	HEIDELBERGCEMENT AG	8.565	0.862	0.460%	-0.424 %	1.433%
HEN3	HENKEL AG & CO KGAA	8.497	0.709	0.398 %	-0.121%	1.151 %
IFX	INFINEON TECHNOLOGIES AG	8.210	9.927	0.913%	-0.166 %	1.570%
LHA	DEUTSCHE LUFTHANSA AG	6.599	3.514	0.767 %	-0.035%	1.139 %
LIN	LINDE AG	19.663	0.470	0.278%	0.128 %	1.171%
MAN	MAN SE	13.300	1.100	0.780 %	-0.233%	0.440%
MEO	METRO AG	14.091	1.095	0.338 %	-0.470%	1.283%
MRK	MERCK KGAA	4.869	0.356	0.551 %	-0.087%	1.073%
MUV2	MUNICH RE CO	19.440	0.811	0.430 %	-0.049%	1.160 %
RWE	RWE AG	19.946	3.175	0.607%	-0.265 %	1.638%
SAP	SAP SE	51.508	3.309	0.270 %	-0.152%	1.177%
SDF	K&S AG	10.226	1.091	0.570 %	-0.198%	1.401%
SIE	SIEMENS AG	83.571	3.401	0.372 %	0.082%	1.413%
TKA	THYSSENKRUPP AG	17.547	2.755	0.535 %	0.374%	2.024%
VOW3	VOLKSWAGEN AG	22.153	1.086	0.638%	0.649%	1.253%
Mean		25.080	4.924	0.496%	-0.011%	1.356%
Median		16.600	2.490	0.467%	-0.056%	1.292%
Min		4.869	0.356	0.202%	-0.470%	0.440%
Max		83.571	58.946	1.153%	0.649%	2.624%

Note: This table presents the average market capitalization (in billion Euros), the average daily trading volume (in million shares), the average daily turnover (in percentage) defined as the trading volume to the number of shares outstanding, the average daily (log) return and its standard deviation (in percentage) in June 2011. All data are from the Compustat Global Security Daily files and based on the primary issues.

Table 2: Summary Statistics

Ticker	Ask Slope (Levels 1-5)		Ask Slope (Levels 5-20)		Bid Slope (Levels 1-5)		Bid Slope (Levels 5-20)	
	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.
ADS	-11.620	0.592	-11.116	0.454	-11.578	0.612	-11.056	0.475
ALV	-11.076	0.631	-10.940	0.431	-11.074	0.636	-10.962	0.424
BAS	-12.245	0.590	-11.854	0.493	-12.233	0.600	-11.828	0.455
BAYN	-12.293	0.582	-11.799	0.448	-12.304	0.551	-11.824	0.435
BEI	-11.634	0.651	-11.258	0.525	-11.631	0.653	-11.148	0.535
BMW	-11.732	0.570	-11.428	0.440	-11.710	0.543	-11.337	0.421
CBK	-17.086	0.597	-16.798	0.476	-17.127	0.613	-16.920	0.523
DAI	-12.425	0.611	-12.189	0.469	-12.404	0.614	-12.167	0.450
DB1	-11.785	0.586	-11.147	0.414	-11.766	0.557	-11.061	0.404
DBK	-12.707	0.647	-12.441	0.521	-12.635	0.640	-12.443	0.469
DPW	-15.038	0.391	-13.946	0.321	-15.025	0.397	-13.978	0.341
DTE	-15.989	0.461	-14.951	0.351	-15.981	0.459	-14.970	0.409
EOAN	-14.625	0.482	-13.959	0.427	-14.574	0.459	-13.934	0.394
FME	-11.772	0.586	-11.143	0.499	-11.712	0.590	-11.061	0.471
FRE	-10.849	0.632	-10.359	0.451	-10.866	0.629	-10.304	0.457
HEI	-11.418	0.678	-11.440	0.483	-11.314	0.669	-11.337	0.447
HEN3	-11.529	0.675	-11.224	0.531	-11.518	0.678	-11.147	0.533
IFX	-15.201	0.662	-15.149	0.491	-15.198	0.657	-15.131	0.430
LHA	-14.376	0.502	-13.692	0.393	-14.379	0.498	-13.641	0.356
LIN	-10.621	0.428	-9.845	0.302	-10.607	0.425	-9.809	0.271
MAN	-11.300	0.609	-10.950	0.482	-11.591	0.807	-11.452	0.607
MEO	-11.663	0.650	-11.621	0.490	-11.663	0.657	-11.571	0.468
MRK	-10.770	0.614	-10.261	0.439	-10.728	0.608	-10.178	0.434
MUV2	-11.121	0.421	-10.112	0.296	-11.100	0.397	-10.348	0.382
RWE	-12.571	0.643	-12.446	0.553	-12.487	0.625	-12.412	0.538
SAP	-12.742	0.592	-12.472	0.503	-12.719	0.576	-12.482	0.482
SDF	-11.703	0.553	-11.150	0.443	-11.713	0.538	-11.103	0.451
SIE	-11.551	0.617	-11.289	0.489	-11.549	0.601	-11.285	0.473
TKA	-12.558	0.637	-12.398	0.467	-12.547	0.613	-12.346	0.435
VOW3	-10.396	0.509	-10.044	0.442	-10.339	0.511	-9.858	0.377

Note: This table presents the summary statistics for the (log) slope measures in June 2011.

Table 3: Estimation Results for ALV

$\alpha_{z,\tau}$	$\tau = 1$	$\tau = 2$	$\tau = 3$	$\tau = 4$	$\tau = 5$	Sum	F-stat
Ask Slope (Levels 1-5)	0.2591*** (0.0165)	-0.0322* (0.0199)	-0.0159 (0.0209)	0.0018 (0.0194)	-0.0483*** (0.0159)	0.1644*** (0.0087)	408.95***
Ask Slope (Levels 5-20)	0.4588*** (0.0439)	-0.2238*** (0.0480)	-0.0502 (0.0518)	-0.0419 (0.0468)	-0.0921*** (0.0368)	0.0508*** (0.0115)	21.24***
Bid Slope (Levels 1-5)	-0.2680*** (0.0181)	0.0105 (0.0209)	0.0270* (0.0191)	-0.0002 (0.0181)	0.0725*** (0.0148)	-0.1581*** (0.0085)	387.30***
Bid Slope (Levels 5-20)	-0.5170*** (0.0439)	0.2090*** (0.0489)	0.1161** (0.0543)	0.0923** (0.0478)	0.0323 (0.0385)	-0.0673*** (0.0118)	36.73***

Note: This table presents the estimates of the coefficients on lagged slope measures ($\alpha_{z,\tau}$) in the return equation 5a for ALV in June 2011. The coefficients of the return equation 5a are estimated via OLS with Newey-West standard errors, which are presented in parentheses under the coefficient estimates. $\sum_{\tau=1}^5 \alpha_{z,\tau}$ is the sum of the coefficient estimates on lagged slope measures and provides a first raw approximation to the long-run effect of slope measures on price dynamics. F-stat is the F-statistic testing the null hypothesis that the coefficient estimates on lagged slope measures are jointly equal to zero. ***, ** and * denote either coefficient estimates that are significantly different from zero or test statistics that are statistically significant at the 1%, 5% and 10%, respectively.

Table 4: Estimation Results

	Ask Slope (Levels 1-5)			Ask Slope (Levels 5-20)			Bid Slope (Levels 1-5)			Bid Slope (Levels 5-20)		
	$\alpha_{z,1}$	$\alpha_{z,\tau}$	F-stat	$\alpha_{z,1}$	$\alpha_{z,\tau}$	F-stat	$\alpha_{z,1}$	$\alpha_{z,\tau}$	F-stat	$\alpha_{z,1}$	$\alpha_{z,\tau}$	F-stat
ADS	0.464***	0.263***	214.49***	0.749***	0.006	0.10	-0.620***	-0.222***	162.51***	-0.841***	-0.011	0.35
ALV	0.259***	0.164***	408.95***	0.459***	0.051***	21.24***	-0.268***	-0.158***	387.30***	-0.517***	-0.067***	36.73***
BAS	0.283***	0.140***	205.12***	0.576***	0.038***	15.71***	-0.289***	-0.309***	211.30***	-0.503***	-0.024**	5.21
BAYN	0.403***	0.123***	127.99***	0.899***	-0.001	0.01	-0.307***	-0.129***	125.11***	-0.584***	0.020	2.66
BEI	0.341***	0.207***	61.26***	0.748***	0.152***	32.20***	-0.402***	-0.251***	90.37***	-0.735***	-0.105***	15.96***
BMW	0.365***	0.166***	284.65***	0.418***	0.062***	32.86***	-0.403***	-0.202***	384.46***	-0.437***	-0.030***	6.95
CBK	0.331***	0.190***	343.70***	0.250***	0.000	0.00	-0.342***	-0.197***	378.23***	-0.161***	-0.005	0.27
DAI	0.329***	0.145***	377.62***	0.555***	0.041***	21.10***	-0.313***	-0.151***	399.50***	-0.486***	-0.038***	16.94***
DBI	0.337***	0.197***	143.46***	0.414***	-0.008	0.15	-0.530***	-0.216***	164.62***	-0.712***	0.060***	8.62
DBK	0.292***	0.140***	390.42***	0.476***	0.019**	5.99	-0.284***	-0.133***	349.36***	-0.415***	-0.044***	27.63***
DPW	1.040***	0.337***	57.54***	0.142	-0.022	0.38	-0.713***	-0.388***	74.25***	-0.326**	0.074*	4.44
DTE	0.580***	0.204***	51.17***	-0.036	-0.086**	13.10**	-0.510***	-0.209**	52.06***	0.282***	0.025	1.72
EOAN	0.455***	0.172***	137.59***	0.624***	-0.001	0.00	-0.465***	-0.161***	111.26***	-0.704***	-0.019	1.85
FME	0.475***	0.290***	188.66***	0.403***	0.060***	8.76	-0.434***	-0.227***	116.65***	-0.611***	-0.065***	9.68*
FRE	0.557***	0.303***	168.96***	0.776***	0.138***	28.52***	-0.568***	-0.300***	161.68***	-0.822***	-0.136***	26.45***
HEI	0.402***	0.192***	153.98***	0.643***	0.084***	17.48***	-0.391***	-0.175***	125.33***	-0.755***	-0.115***	29.42***
HEN3	0.355	0.239***	182.02***	0.493***	0.087***	20.50***	-0.398***	-0.227***	169.80***	-0.606***	-0.132**	48.08***
IFX	0.397***	0.198***	239.14***	0.532***	0.081***	26.81***	-0.442***	-0.230***	327.93***	-0.625***	-0.109***	36.48***
LHA	0.669***	0.303***	82.75***	0.813***	-0.010	0.12	-0.706***	-0.286***	75.43***	-0.828***	-0.069*	4.01
LIN	0.923***	0.130**	6.01	0.608***	-0.055	1.18	-0.826***	-0.163***	9.80*	-0.225	0.140**	5.58
MAN	0.143***	0.073***	98.90***	0.090**	0.014	2.79	-0.066***	-0.062***	121.82***	-0.021	-0.023***	10.75*
MEQ	0.346***	0.209***	200.40***	0.644***	0.126***	49.33***	-0.362***	-0.225***	237.11***	-0.733***	-0.131***	50.98***
MRK	0.583***	0.327***	191.72***	0.647***	0.111***	18.42***	-0.471***	-0.349***	215.68***	-0.492***	-0.086***	11.29***
MUV2	1.059***	0.324***	45.88***	0.141	-0.111**	5.91	-1.091***	-0.306***	35.03***	-0.337**	-0.056*	3.13
RWE	0.327***	0.121***	165.02***	0.512***	0.021*	4.09	-0.282***	-0.116***	149.24***	-0.479***	-0.026**	6.17
SAP	0.239***	0.086***	90.58***	0.484***	0.029**	9.09	-0.251***	-0.106***	127.69***	-0.469***	-0.020*	4.45
SDF	0.467***	0.189***	99.85***	0.866***	0.066***	11.74**	-0.548***	-0.153***	63.27***	-0.969***	-0.049**	6.31
SIE	0.243***	0.130***	448.01***	0.387***	0.033***	21.65***	-0.228***	-0.133***	433.89***	-0.360***	-0.030***	16.93***
TKA	0.458***	0.205***	321.77***	0.591***	0.035**	6.45	-0.509***	-0.217***	344.32***	-0.516***	-0.015	0.97
VOW3	0.332***	0.237***	106.24***	0.342***	0.009	0.20	-0.439***	-0.207***	91.08***	-0.284**	0.010	0.15
Bonferroni p-value	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
nb significant	30	30	29	27	19	14	30	30	29	28	19	11
Mean	0.448	0.200	0.448	0.508	0.032	0.449	-0.449	-0.201	0.449	-0.509	-0.036	0.448
Median	0.381	0.194	0.381	0.522	0.031	0.419	-0.419	-0.205	0.419	-0.510	-0.030	0.381
Min	0.143	0.073	0.143	-0.036	-0.111	-1.091	-1.091	-0.388	-1.091	-0.969	-0.136	0.143
Max	1.059	0.337	1.059	0.899	0.152	-0.066	-0.066	-0.062	-0.066	0.282	0.140	1.059

Note: This table presents the coefficient estimates on the first lag of a given slope measure in the return equation 5a ($\alpha_{z,1}$), the sum of coefficient estimates on all lags of the same slope measure ($\sum_{\tau=1}^5 \alpha_{z,\tau}$) and F-statistic testing the null hypothesis that the coefficient estimates on all lags of the same slope measure are jointly equal to zero. Bonferroni p-value is the p-value based on the Bonferroni correction and calculated as $\min(1, \min(p_1, \dots, p_{30}) \times 30)$ where p_i is the p-value of a given statistic from the estimation of the return equation 5a for the i^{th} stock. A Bonferroni p-value less than 0.05 suggests that a given null hypothesis is rejected jointly across all stocks at least 5% significance level. nb significant is the number of stocks (out of 30) for which the results are statistically significant at the 5% level. ***, **, * and * denote either coefficient estimates that are significantly different from zero or test statistics that are statistically significant at the 1%, 5% and 10%, respectively.

Table 5: Results for the Hypotheses on the Immediate Effect of Slope Measures on Price Dynamics

	H1a	H2a	H3a	H4a
ADS	20.57***	1.42	24.39***	15.59***
ALV	0.29	1.79	42.41***	66.27***
BAS	0.05	2.10	64.92***	38.35***
BAYN	10.18***	26.03***	117.33***	37.64***
BEI	3.11*	0.03	58.47***	39.13***
BMW	2.75*	0.14	2.04	0.78
CBK	0.10	1.82	2.74*	17.11***
DAI	1.09	3.28*	68.18***	42.36***
DB1	23.53***	12.89***	1.66	9.86***
DBK	0.27	2.38	43.43***	23.57***
DPW	12.66***	1.86	50.50***	10.08***
DTE	0.81	5.24**	38.83***	70.46***
EOAN	0.05	1.23	8.99***	18.18***
FME	1.14	7.87***	1.79	9.65***
FRE	0.08	0.32	12.32***	16.55***
HEI	0.22	3.82*	31.19***	71.83***
HEN3	2.91*	4.38**	12.37***	27.43***
IFX	3.79*	2.41	10.06***	17.58***
LHA	0.28	0.01	1.81	1.27
LIN	0.78	4.71**	3.73*	13.69***
MAN	17.75***	2.71*	3.31*	2.16
MEO	0.39	2.30	46.47***	73.45***
MRK	9.33***	4.16**	1.38	0.14
MUV2	0.07	1.21	28.93***	23.89***
RWE	4.80**	0.50	30.53***	30.91***
SAP	0.33	0.12	64.27***	48.22***
SDF	4.35**	1.80	56.20***	57.92***
SIE	1.09	0.64	36.92***	30.56***
TKA	4.32**	1.54	8.91***	0.03
VOW3	2.48	0.21	0.01	2.56
Bonferroni p-value	0.000	0.000	0.000	0.000
nb significant	9	7	21	24

Note: This table presents the F-statistics for the null hypotheses on the immediate effect of slope measures on price dynamics. $H1a : \alpha_{z,1}(S_{1,5,t}^A) = (1 - 2I_{1a})\alpha_{z,1}(S_{1,5,t}^B)$ and $H2a : \alpha_{z,1}(S_{5,20,t}^A) = (1 - 2I_{2a})\alpha_{z,1}(S_{5,20,t}^B)$ are the null hypotheses on the equality between the immediate effects of the ask- and bid-side slopes based on lower and higher levels, respectively, where $\alpha_{z,i}(\cdot)$ denote the element of $\alpha_{z,i}$ corresponding to the slope measure of interest in parenthesis. The terms multiplying the coefficients on the bid-side slope measures, i.e. $(1 - 2I_{1a})$ and $(1 - 2I_{2a})$, allow us to compare the magnitudes of the coefficient estimates if they are positive for the ask- and negative for the bid-side slope. To be more precise, I_{1a} and I_{2a} are binary variables defined as $I_{1a} = 1_{\{\alpha_{z,1}(S_{1,5,t}^A) > 0, \alpha_{z,1}(S_{1,5,t}^B) < 0\}}$ and $I_{2a} = 1_{\{\alpha_{z,1}(S_{5,20,t}^A) > 0, \alpha_{z,1}(S_{5,20,t}^B) < 0\}}$ where $1_{\{x,y\}}$ is an indicator function that takes the value of one if both conditions x and y are satisfied and zero otherwise. Similarly, $H3a : \alpha_{z,1}(S_{1,5,t}^A) = \alpha_{z,1}(S_{5,20,t}^A)$ and $H4a : \alpha_{z,1}(S_{1,5,t}^B) = \alpha_{z,1}(S_{5,20,t}^B)$ are the null hypotheses on the equality between the immediate effects of the ask- and bid-side slopes based on different levels, respectively. Bonferroni p-value is the p-value based on the Bonferroni correction and calculated as $\min(1, \min(p_1, \dots, p_{30}) \times 30)$ where p_i is the p-value of a given statistic from the estimation of the return equation 5a for the i^{th} stock. A Bonferroni p-value less than 0.05 suggests that a given null hypothesis is rejected jointly across all stocks at 5% significance level. nb significant is the number of stocks (out of 30) for which the results are statistically significant at 5%. ***, ** and * denote statistical significant test statistics at the 1%, 5% and 10% levels, respectively.

Table 6: Results for the Hypotheses on the Overall Effect of Slope Measures on Price Dynamics

	H1b	H2b	H3b	H4b
ADS	247.20***	501.13***	133.91***	85.22***
ALV	65.80***	392.56***	118.34***	133.30***
BAS	175.82***	428.29***	135.88***	122.46***
BAYN	256.81***	280.64***	198.67***	137.70***
BEI	47.52***	331.80***	76.62***	76.00***
BMW	275.08***	251.62***	49.07***	117.47***
CBK	114.84***	50.21***	149.89***	145.30***
DAI	467.68***	668.92***	168.76***	170.31***
DB1	147.73***	242.97***	65.07***	117.73***
DBK	377.44***	482.69***	162.52***	93.23***
DPW	81.98***	32.83***	104.38***	58.16***
DTE	48.70***	11.08**	78.25***	111.59***
EOAN	71.15***	291.17***	86.34***	71.72***
FME	78.59***	173.85***	63.08***	54.78***
FRE	103.48***	307.33***	52.39***	53.84***
HEI	233.30***	449.92***	68.99***	94.58***
HEN3	65.60***	320.10***	57.91***	52.60***
IFX	274.93***	264.78***	48.98***	53.31***
LHA	67.47***	113.66***	47.31***	32.89***
LIN	95.55***	20.61***	15.99***	35.64***
MAN	36.04***	14.30**	30.35***	21.99***
MEO	104.09***	385.66***	78.01***	114.65***
MRK	67.33***	160.64***	43.27***	61.47***
MUV2	88.90***	15.38***	56.83***	39.21***
RWE	312.08***	374.36***	94.98***	89.04***
SAP	192.21***	469.38***	98.12***	121.91***
SDF	203.53***	479.36***	95.41***	92.20***
SIE	261.23***	440.31***	150.55***	152.80***
TKA	433.92***	327.19***	95.17***	97.49***
VOW3	44.30***	36.55***	60.32***	50.82***
Bonferroni p-value	0.000	0.000	0.000	0.000
nb significant	30	30	30	30

Note: This table presents the F-statistics for the null hypotheses on the overall dynamic effect of slope measures on price dynamics. $H1b : \alpha_{z,\tau}(S_{1,5,t}^A) = (1 - 2I_{1b,\tau})\alpha_{z,\tau}(S_{1,5,t}^B)$ and $H2b : \alpha_{z,\tau}(S_{5,20,t}^A) = (1 - 2I_{2b,\tau})\alpha_{z,\tau}(S_{5,20,t}^B)$ for $\tau = 1, 2, \dots, 5$ are the null hypotheses on the equality between the overall effects of the ask- and bid-side slopes based on lower and higher levels, respectively. $\alpha_{z,i}(\cdot)$ denote the element of $\alpha_{z,i}$ corresponding to the slope measure of interest in parenthesis. The terms multiplying the coefficients on the bid-side slope measures, i.e. $(1 - 2I_{1b,\tau})$ and $(1 - 2I_{2b,\tau})$, allow us to compare the magnitudes of the coefficient estimates if they are positive for the ask- and negative for the bid-side slope. $I_{1b,\tau}$ and $I_{2b,\tau}$ are binary variables defined as $I_{1b,\tau} = 1_{\{\alpha_{z,\tau}(S_{1,5,t}^A) > 0, \alpha_{z,\tau}(S_{1,5,t}^B) < 0\}}$ and $I_{2b,\tau} = 1_{\{\alpha_{z,\tau}(S_{5,20,t}^A) > 0, \alpha_{z,\tau}(S_{5,20,t}^B) < 0\}}$ for $\tau = 1, 2, \dots, 5$ where $1_{\{x,y\}}$ is an indicator function that takes the value of one if both conditions x and y are satisfied and zero otherwise. Similarly, $H3b : \alpha_{z,\tau}(S_{1,5,t}^A) = \alpha_{z,\tau}(S_{5,20,t}^A)$ and $H4b : \alpha_{z,\tau}(S_{1,5,t}^B) = \alpha_{z,\tau}(S_{5,20,t}^B)$ for $\tau = 1, 2, \dots, 5$ are the null hypotheses on the equality between the overall effects of the ask- and bid-side slopes based on different levels, respectively. Bonferroni p-value is the p-value based on the Bonferroni correction and calculated as $\min(1, \min(p_1, \dots, p_{30}) \times 30)$ where p_i is the p-value of a given statistic from the estimation of the return equation 5a for the i^{th} stock. A Bonferroni p-value less than 0.05 suggests that a given null hypothesis is rejected jointly across all stocks at 5% significance level. nb significant is the number of stocks (out of 30) for which the results are statistically significant at 5%. ***, ** and * denote statistical significant test statistics at the 1%, 5% and 10% levels, respectively.

Table 7: Results for the Hypotheses on the Long-run Effect of Slope Measures on Price Dynamics

	H1c	H2c	H3c	H4c
ADS	4.87**	0.04	78.50***	55.42***
ALV	0.43	1.50	56.69***	37.28***
BAS	0.11	1.00	47.35***	54.37***
BAYN	0.27	1.37	47.92***	64.33***
BEI	2.35	1.60	1.89	13.27**
BMW	11.56***	5.28**	42.76***	110.55***
CBK	0.31	0.12	126.79***	136.58***
DAI	0.37	0.09	68.12***	78.01***
DB1	1.13	5.50**	52.01***	91.23***
DBK	0.81	4.91**	96.79***	54.28***
DPW	2.27	4.57**	34.02***	51.71***
DTE	0.06	15.14***	48.82***	36.03***
EOAN	0.69	0.90	67.04***	43.40***
FME	10.45***	0.04	59.09***	29.14***
FRE	0.01	0.00	19.74***	17.95***
HEI	0.72	1.57	14.78**	4.50
HEN3	0.39	3.47*	28.21***	11.68**
IFX	4.38**	2.01	27.56***	26.42***
LHA	0.36	1.64	39.77***	17.30***
LIN	0.61	6.25**	5.40	12.68**
MAN	2.32	0.88	26.41***	14.85**
MEO	0.88	0.05	10.84*	14.61**
MRK	1.03	0.57	36.49***	56.40***
MUV2	0.25	1.01	35.45***	14.34**
RWE	0.16	0.20	43.16***	34.32***
SAP	3.71*	0.44	16.17***	36.99***
SDF	3.57*	0.49	18.83***	12.95**
SIE	0.20	0.15	87.15***	94.43***
TKA	0.77	1.31	73.29***	91.80***
VOW3	1.84	0.00	43.01***	35.08***
Bonferroni p-value	0.020	0.003	0.000	0.000
nb significant	4	6	29	30

Note: This table presents the F-statistics for the null hypotheses on the long-run effects of slope measures on price dynamics. $H1c : \sum_{\tau=1}^5 \alpha_{z,\tau}(S_{1,5,t}^A) = (1 - 2I_{1c}) \sum_{\tau=1}^5 \alpha_{z,\tau}(S_{1,5,t}^B)$ and $H2c : \sum_{\tau=1}^5 \alpha_{z,\tau}(S_{5,20,t}^A) = (1 - 2I_{2c}) \sum_{\tau=1}^5 \alpha_{z,\tau}(S_{5,20,t}^B)$ are the null hypotheses on the equality between the long-run effects of the ask- and bid-side slopes based on lower and higher levels, respectively. $\alpha_{z,i}(\cdot)$ denote the element of $\alpha_{z,i}$ corresponding to the slope measure of interest in parenthesis. The terms multiplying the coefficients on the bid-side slope measures, i.e. $(1 - 2I_{1c})$ and $(1 - 2I_{2c})$, allow us to compare the magnitudes of the sums if they are positive for the ask- and negative for the bid-side slope. I_{1c} and I_{2c} are binary variables defined as $I_{1c} = 1_{\{\sum_{\tau=1}^5 \alpha_{z,\tau}(S_{1,5,t}^A) > 0, \sum_{\tau=1}^5 \alpha_{z,\tau}(S_{1,5,t}^B) < 0\}}$ and $I_{2c} = 1_{\{\sum_{\tau=1}^5 \alpha_{z,\tau}(S_{5,20,t}^A) > 0, \sum_{\tau=1}^5 \alpha_{z,\tau}(S_{5,20,t}^B) < 0\}}$, respectively, where $1_{\{x,y\}}$ is an indicator function that takes the value of one if both conditions x and y are satisfied and zero otherwise. Similarly, $H3c : \sum_{\tau=1}^5 \alpha_{z,\tau}(S_{1,5,t}^A) = \sum_{\tau=1}^5 \alpha_{z,\tau}(S_{5,20,t}^A)$ and $H4c : \sum_{\tau=1}^5 \alpha_{z,\tau}(S_{1,5,t}^B) = \sum_{\tau=1}^5 \alpha_{z,\tau}(S_{5,20,t}^B)$ are the null hypotheses on the equality between the long-run effects of the ask- and bid-side slopes based on different levels, respectively. Bonferroni p-value is the p-value based on the Bonferroni correction and calculated as $\min(1, \min(p_1, \dots, p_{30}) \times 30)$ where p_i is the p-value of a given statistic from the estimation of the return equation 5a for the i^{th} stock. A Bonferroni p-value less than 0.05 suggests that a given null hypothesis is rejected jointly across all stocks at 5% significance level. nb significant is the number of stocks (out of 30) for which the results are statistically significant at 5%. ***, ** and * denote statistical significant test statistics at the 1%, 5% and 10% levels, respectively.

Table 8: Summary Statistics for Hypotheses Test Results using Alternative Definitions and Sample Periods

(a) Sample Period: July 2011

Hypothesis	$S_{1,5}$ vs $S_{5,20}$	$S_{1,10}$ vs $S_{10,20}$	$S_{2,10}$ vs $S_{10,20}$	$S_{2,5}$ vs $S_{5,20}$	$S_{1,2}$ vs $S_{2,20}$
H1a	9 (0.000)	5 (0.015)	5 (0.024)	5 (0.023)	9 (0.012)
H2a	7 (0.000)	7 (0.000)	7 (0.000)	8 (0.000)	11 (0.000)
H3a	21 (0.000)	23 (0.000)	23 (0.000)	25 (0.000)	27 (0.000)
H4a	24 (0.000)	24 (0.000)	23 (0.000)	24 (0.000)	25 (0.000)
H1b	30 (0.000)	26 (0.000)	24 (0.000)	20 (0.000)	28 (0.000)
H2b	30 (0.000)	27 (0.000)	28 (0.000)	28 (0.000)	28 (0.000)
H3b	30 (0.000)	29 (0.000)	29 (0.000)	29 (0.000)	30 (0.000)
H4b	30 (0.000)	29 (0.000)	29 (0.000)	28 (0.000)	29 (0.000)
H1c	4 (0.020)	5 (0.634)	4 (0.323)	4 (0.469)	6 (0.042)
H2c	6 (0.003)	4 (0.119)	4 (0.145)	4 (0.067)	2 (1.000)
H3c	29 (0.000)	27 (0.000)	25 (0.000)	20 (0.000)	18 (0.000)
H4c	30 (0.000)	26 (0.000)	25 (0.000)	20 (0.000)	18 (0.000)

(b) Sample Period: May 2011

Hypothesis	$S_{1,5}$ vs $S_{5,20}$	$S_{1,10}$ vs $S_{10,20}$	$S_{2,10}$ vs $S_{10,20}$	$S_{2,5}$ vs $S_{5,20}$	$S_{1,2}$ vs $S_{2,20}$
H1a	13 (0.000)	5 (0.011)	4 (0.026)	4 (0.085)	16 (0.000)
H2a	9 (0.000)	7 (0.000)	8 (0.000)	9 (0.001)	8 (0.009)
H3a	23 (0.000)	19 (0.000)	20 (0.000)	25 (0.000)	26 (0.000)
H4a	22 (0.000)	23 (0.000)	23 (0.000)	24 (0.000)	26 (0.000)
H1b	30 (0.000)	24 (0.000)	20 (0.000)	18 (0.000)	30 (0.000)
H2b	29 (0.000)	27 (0.000)	27 (0.000)	28 (0.000)	28 (0.000)
H3b	30 (0.000)	29 (0.000)	29 (0.000)	30 (0.000)	30 (0.000)
H4b	30 (0.000)	30 (0.000)	30 (0.000)	29 (0.000)	30 (0.000)
H1c	5 (0.027)	2 (1.000)	0 (1.000)	1 (1.000)	7 (0.006)
H2c	5 (0.000)	3 (0.000)	3 (0.000)	5 (0.000)	8 (0.000)
H3c	30 (0.000)	26 (0.000)	26 (0.000)	19 (0.000)	16 (0.000)
H4c	30 (0.000)	27 (0.000)	25 (0.000)	18 (0.000)	17 (0.000)

(c) Sample Period: June 2010

Hypothesis	$S_{1,5}$ vs $S_{5,20}$	$S_{1,10}$ vs $S_{10,20}$	$S_{2,10}$ vs $S_{10,20}$	$S_{2,5}$ vs $S_{5,20}$	$S_{1,2}$ vs $S_{2,20}$
H1a	7 (0.000)	5 (0.016)	5 (0.027)	4 (0.000)	11 (0.025)
H2a	11 (0.000)	10 (0.000)	10 (0.000)	11 (0.000)	11 (0.000)
H3a	19 (0.000)	18 (0.000)	18 (0.000)	23 (0.000)	21 (0.000)
H4a	22 (0.000)	22 (0.000)	23 (0.000)	27 (0.000)	26 (0.000)
H1b	29 (0.000)	23 (0.000)	21 (0.000)	24 (0.000)	28 (0.000)
H2b	26 (0.000)	27 (0.000)	27 (0.000)	25 (0.000)	27 (0.000)
H3b	27 (0.000)	28 (0.000)	26 (0.000)	27 (0.000)	27 (0.000)
H4b	28 (0.000)	28 (0.000)	28 (0.000)	29 (0.000)	28 (0.000)
H1c	6 (0.232)	6 (0.026)	6 (0.026)	8 (0.016)	4 (0.072)
H2c	3 (0.115)	7 (0.024)	7 (0.034)	7 (0.005)	8 (0.010)
H3c	17 (0.000)	19 (0.000)	20 (0.000)	17 (0.000)	18 (0.000)
H4c	19 (0.000)	24 (0.000)	22 (0.000)	15 (0.000)	17 (0.000)

Note: This table presents the number of stocks for which we reject H1a to H4c as well as the Bonferroni p-values for testing these hypotheses jointly across stocks in parentheses underneath. S_{l_1, l_2} is the slope between levels l_1 and l_2 of either the bid or the ask-sides depending on the hypothesis under consideration. $H1a$ and $H2a$ are the null hypotheses on the equality between the immediate effects of the ask- and bid-side slopes based on lower and higher levels, respectively. $H3a$ and $H4a$ are the null hypotheses on the equality between the immediate effects of the ask- and bid-side slopes based on different levels, respectively. $H1b$ and $H2b$ are the null hypotheses on the equality between the overall effects of the ask- and bid-side slopes based on lower and higher levels, respectively. $H3b$ and $H4b$ are the null hypotheses on the equality between the overall effects of the ask- and bid-side slopes based on different levels, respectively. $H1c$ and $H2c$ are the null hypotheses on the equality between the long-run effects of the ask- and bid-side slopes based on lower and higher levels, respectively. $H3c$ and $H4c$ are the null hypotheses on the equality between the long-run effects of the ask- and bid-side slopes based on different levels, respectively. Bonferroni p-value is the p-value based on the Bonferroni correction and calculated as $\min(1, \min(p_1, \dots, p_{30}) \times 30)$ where p_i is the p-value of a given statistic from the estimation of the return equation 5a for the i^{th} stock. A Bonferroni p-value less than 0.05 suggests that a given null hypothesis is rejected jointly across all stocks at 5% significance level.

Table 9: Relative Performances of the Restricted Trading Strategies

Ticker	Restrictions on the Immediate Effect				Restrictions on the Overall Short-Run Effect			
	H1a	H2a	H3a	H4a	H1b	H2b	H3b	H4b
ADS	0.200%	0.247%**	0.274%**	0.107%	0.218%	0.309%***	0.284%**	0.241%
ALV	0.287%*	0.374%***	0.099%	0.272%***	0.277%*	0.323%***	0.266%*	0.367%***
BAS	0.263%**	0.228%	0.118%	0.212%**	0.224%**	0.281%**	0.158%	0.180%
BAYN	0.149%	0.177%	0.328%**	0.215%**	0.230%*	0.254%*	0.252%**	0.206%*
BEI	-0.044%	0.056%	0.050%	-0.032%	-0.083%	-0.011%	0.023%	0.077%
BMW	0.389%***	0.156%	0.306%**	0.156%	0.379%***	0.194%	0.414%***	0.357%***
CBK	0.590%*	0.214%	0.289%*	0.292%	0.480%	0.177%	0.561%**	0.801%***
DAI	0.424%*	0.407%**	0.398%**	0.252%	0.418%*	0.411%**	0.535%***	0.352%
DB1	0.167%	0.089%	0.200%**	0.099%	0.241%*	-0.016%	0.259%***	0.172%
DBK	0.756%***	0.114%	0.284%***	0.469%**	0.753%***	0.234%	0.252%	0.430%
DPW	-0.031%	0.007%	0.111%*	-0.049%	0.041%	0.021%	0.128%	-0.019%
DTE	-0.034%	-0.056%	-0.057%	0.041%	-0.171%	-0.010%	-0.049%	0.034%
EOAN	0.051%	0.095%	0.133%	0.066%	0.033%	0.148%	0.204%	0.061%
FME	0.193%**	0.218%***	0.123%	0.303%***	0.169%**	0.233%***	0.285%**	0.320%***
FRE	0.399%***	0.397%**	0.302%***	0.312%**	0.356%***	0.348%***	0.245%	0.438%***
HEI	0.059%	0.405%*	0.208%	0.155%	0.107%	0.432%**	0.401%**	0.300%
HEN3	0.383%***	0.476%***	0.426%***	0.582%***	0.422%***	0.530%***	0.570%***	0.615%***
IFX	0.102%	0.096%	0.114%	0.186%	0.061%	0.032%	0.350%	0.241%
LHA	0.114%	-0.077%	-0.100%	-0.020%	0.069%	0.009%	-0.121%	-0.020%
LIN	-0.025%	0.001%	-0.004%	-0.035%	-0.065%	-0.079%	-0.035%	-0.130%
MAN	0.112%**	0.002%	0.048%	0.023%	0.117%**	0.015%	0.136%**	0.102%*
MEO	0.687%***	0.398%***	0.393%**	0.473%***	0.686%***	0.424%***	0.452%**	0.704%***
MRK	0.192%*	0.131%	0.119%*	0.252%**	0.202%	0.119%	0.228%**	0.238%**
MUV2	-0.129%	-0.093%	-0.080%	-0.187%	-0.220%*	-0.114%	0.070%	0.167%
RWE	0.222%	0.092%	0.126%	0.118%	0.140%	0.133%	0.146%	0.107%
SAP	-0.105%	0.176%	0.123%	0.022%	-0.084%	0.219%	0.055%	0.190%
SDF	0.090%	-0.021%	0.127%	0.162%	0.007%	0.022%	0.261%	-0.008%
SIE	0.610%***	0.336%***	0.263%*	0.503%***	0.816%***	0.373%**	0.582%***	0.651%***
TKA	0.479%**	0.306%**	0.231%	0.361%***	0.495%***	0.264%*	0.556%**	0.379%**
VOW3	-0.124%	-0.089%*	-0.100%	-0.238%	-0.184%	-0.037%	-0.073%	-0.306%
Mean	0.214%	0.162%	0.162%	0.169%	0.205%	0.175%	0.246%	0.242%
Median	0.180%	0.144%	0.126%	0.159%	0.186%	0.185%	0.252%	0.222%
nb pos	23	25	25	24	24	24	26	25
nb pos & signif	10	9	9	11	10	10	14	10
Bonf. p-val	0.000	0.004	0.015	0.000	0.000	0.002	0.014	0.001

Note: This table presents the differences between the average daily cumulative returns of the unrestricted and restricted trading strategies. A positive number implies that the unrestricted strategy provides, on average, higher daily cumulative returns than the trading strategy that uses the restricted forecasting model implied by hypotheses in the corresponding column heading. Mean and median are calculated over all stocks in our sample. nb pos and nb pos & signif present the number of stocks out of 30 in our sample for which the return difference is positive and significantly positive, respectively. Bonf. p-val is the p-value based on the Bonferroni correction for stocks with positive return differences and calculated as $\min(1, \min(p_1, \dots, p_k) \times k)$ where the minimum is calculated only over stocks for which the return difference is positive and, thus, k is the number of stock for which the return difference is positive. p_i is the p-value of the hypothesis testing whether the mean of daily cumulative returns of the unrestricted and a given restricted trading strategy is equal. ***, ** and * denote that the return differences are significantly different from zero at the 1%, 5% and 10% levels, respectively.

Table 10: Average Relative Performances of the Restricted Trading Strategies at Different Thresholds

κ	Restrictions on the Immediate Effect				Restrictions on the Overall Short-Run Effect			
	H1a	H2a	H3a	H4a	H1b	H2b	H3b	H4b
0	0.214%	0.162%	0.162%	0.169%	0.205%	0.175%	0.246%	0.242%
0.1	0.447%	0.261%	0.238%	0.259%	0.438%	0.264%	0.375%	0.412%
0.2	0.389%	0.107%	0.118%	0.085%	0.350%	0.102%	0.245%	0.294%
0.3	0.850%	0.405%	0.462%	0.416%	0.909%	0.386%	0.740%	0.620%
0.4	1.232%	0.582%	0.604%	0.583%	1.269%	0.616%	0.910%	0.953%
0.5	0.807%	0.390%	0.425%	0.450%	0.859%	0.397%	0.696%	0.688%

Note: This table presents the difference between the daily cumulative returns of the unrestricted and the restricted trading strategy that uses the restricted forecasting model implied by hypotheses in the corresponding column heading, averaged over all stocks and trading days in June 2011, for different values of the threshold parameter, κ . We trade less frequently as higher values of κ are considered, which filters out potentially weak signals. A positive number implies that the unrestricted strategy provides, on average, higher daily cumulative returns than the trading strategy that uses the restricted forecasting model implied by hypotheses in the corresponding column heading.

Table 11: Relative Performances of the Restricted Strategies when Trading at the Best Bid and Ask Prices

Ticker	Restrictions on the Immediate Effect				Restrictions on the Overall Short-Run Effect			
	H1a	H2a	H3a	H4a	H1b	H2b	H3b	H4b
ADS	0.536%**	0.190%	0.253%**	0.084%	0.535%**	0.226%	0.268%**	-0.013%
ALV	0.563%**	0.437%**	0.172%	0.341%**	0.568%**	0.364%**	0.375%**	0.485%**
BAS	0.326%**	0.017%	-0.137%	0.164%	0.315%**	0.017%	-0.255%**	-0.091%
BAYN	0.332%**	0.558%**	0.537%**	0.390%**	0.258%	0.387%**	0.272%**	0.108%
BEI	-0.101%	0.094%	0.013%	-0.009%	-0.115%	0.052%	-0.278%	0.017%
BMW	0.184%	0.080%	0.171%*	0.175%	0.215%	0.116%	0.070%	0.246%*
CBK	0.240%**	-0.032%	0.069%	0.050%	0.236%**	-0.040%	0.265%**	0.316%**
DAI	0.551%**	0.435%**	0.316%**	0.300%**	0.560%**	0.406%**	0.412%**	0.477%**
DB1	0.473%**	0.291%**	0.216%*	0.194%	0.517%**	-0.001%	0.275%	0.169%
DBK	0.974%**	0.218%	0.352%**	0.378%**	0.974%**	0.194%	0.527%**	0.531%**
DPW	-0.463%**	-0.061%	0.049%	-0.329%**	-0.492%**	-0.037%	-0.036%	-0.480%**
DTE	-0.394%**	-0.171%**	-0.123%	0.165%**	-0.603%**	-0.150%*	-0.052%	0.015%
EOAN	-0.289%*	-0.097%	-0.283%**	-0.161%	-0.401%**	-0.153%	-0.250%	-0.184%
FME	0.060%	0.162%	0.247%*	0.139%	0.033%	0.084%	0.272%	-0.080%
FRE	1.149%**	0.656%**	0.646%**	0.743%**	1.136%**	0.624%**	0.952%**	0.869%**
HEI	0.418%**	0.480%**	0.314%**	0.449%**	0.349%**	0.460%**	0.454%**	0.202%
HEN3	1.114%**	0.632%**	0.582%**	0.847%**	1.323%**	0.697%**	0.931%**	1.054%**
IFX	1.093%**	0.252%	0.210%	0.487%**	1.039%**	0.193%	0.476%**	0.527%**
LHA	-0.169%	-0.291%**	-0.237%*	-0.410%**	-0.295%*	-0.357%**	-0.125%	-0.605%**
LIN	-0.142%	-0.135%	-0.267%*	-0.352%**	-0.301%	-0.262%**	-0.481%**	-0.314%*
MAN	0.274%**	0.019%	0.199%**	0.037%*	0.314%**	0.046%	0.214%**	0.355%**
MEO	1.043%**	0.450%**	0.444%**	0.692%**	1.031%**	0.452%*	0.720%**	0.968%**
MRK	0.429%*	0.036%	0.120%	0.315%**	0.488%**	-0.071%	0.309%	0.264%
MUV2	-1.442%**	-0.085%	-0.109%	-0.880%**	-1.647%**	-0.132%	-0.219%	-0.663%**
RWE	0.475%**	0.145%	0.297%*	0.103%	0.410%**	0.197%	0.205%	0.081%
SAP	-0.117%	0.134%	0.116%	0.074%	-0.066%	0.186%	-0.058%	0.000%
SDF	-0.011%	-0.183%	-0.026%	-0.202%	-0.168%	-0.239%	-0.033%	-0.334%*
SIE	0.977%**	0.185%	0.119%	0.393%**	1.086%**	0.213%*	0.385%**	0.508%**
TKA	0.810%**	0.196%*	0.194%	0.204%	0.860%**	0.104%	0.308%	0.218%
VOW3	0.010%	-0.195%**	-0.240%**	-0.210%**	-0.091%	-0.402%**	-0.027%	-0.404%**
Mean	0.297%	0.147%	0.140%	0.139%	0.269%	0.106%	0.196%	0.141%
Median	0.329%	0.139%	0.172%	0.165%	0.314%	0.094%	0.267%	0.138%
nb positive	21	21	22	22	20	19	19	19
nb positive & signif	17	7	8	12	17	6	12	10
Bonferroni p-val	0.000	0.001	0.024	0.000	0.000	0.000	0.001	0.002

Note: This table presents the difference between the daily cumulative returns of the unrestricted and the trading strategy that uses the restricted forecasting model implied by hypotheses in the corresponding column heading, when we trade at the best bid and ask prices instead of midquote prices and a κ of zero. A positive number implies that the unrestricted strategy provides, on average, higher daily cumulative returns than the trading strategy that uses the restricted forecasting model implied by hypotheses in the corresponding column heading. Mean and median are calculated over all stocks in our sample. nb pos and nb pos & signif present the number of stocks out of 30 in our sample for which the return difference is positive and significantly positive, respectively. Bonf. p-val is the p-value based on the Bonferroni correction for stocks with positive return differences and calculated as $\min(1, \min(p_1, \dots, p_k) \times k)$ where the minimum is calculated only over stocks for which the return difference is positive and, thus, k is the number of stock for which the return difference is positive. p_i is the p-value of the hypothesis testing whether the mean of daily cumulative returns of the unrestricted and a given restricted trading strategy is equal. ***, ** and * denote that the return differences are significantly different from zero at the 1%, 5% and 10% levels, respectively.

Table 12: Relative Performances of the Restricted Strategies when Trading at a Slower Speed

Ticker	Restrictions on the Immediate Effect				Restrictions on the Overall Short-Run Effect			
	H1a	H2a	H3a	H4a	H1b	H2b	H3b	H4b
ADS	0.241%*	0.351%*	0.164%*	0.301%*	0.187%*	0.340%*	0.179%*	0.187%*
ALV	0.213%**	0.200%**	0.042%**	0.161%**	0.208%**	0.223%**	0.176%**	0.265%**
BAS	0.166%*	-0.089%*	-0.012%*	0.129%*	0.185%*	-0.071%*	-0.074%*	-0.005%*
BAYN	0.174%**	-0.053%**	0.153%**	0.141%**	0.248%**	0.017%**	0.261%**	0.128%**
BEI	0.007%**	0.059%**	0.048%**	0.036%**	0.014%**	0.045%**	0.001%**	0.187%**
BMW	0.179%*	0.039%*	0.028%*	0.088%*	0.133%*	0.047%*	0.082%*	0.104%*
CBK	0.266%*	0.029%*	0.125%*	0.269%*	0.236%*	0.038%*	0.251%*	0.554%*
DAI	0.339%**	0.096%**	0.291%**	0.093%**	0.329%**	0.182%**	0.298%**	0.061%**
DB1	0.202%**	-0.009%**	0.113%**	0.062%**	0.259%**	-0.077%**	0.100%**	0.093%**
DBK	0.366%**	-0.079%**	-0.003%**	0.210%**	0.352%**	-0.008%**	0.008%**	0.062%**
DPW	0.019%**	-0.009%**	0.048%**	0.046%**	0.011%**	0.072%**	-0.026%**	-0.053%**
DTE	-0.174%**	-0.036%**	-0.176%**	-0.024%**	-0.218%**	-0.051%**	-0.190%**	-0.065%**
EOAN	-0.013%**	-0.027%**	-0.108%**	0.134%**	0.040%**	0.034%**	-0.157%**	0.133%**
FME	0.218%**	0.182%**	0.154%**	0.201%**	0.181%**	0.123%**	0.256%**	0.139%**
FRE	0.472%**	0.224%**	0.233%**	0.256%**	0.414%**	0.218%**	0.352%**	0.404%**
HEI	-0.145%*	0.424%*	0.174%*	0.228%*	-0.070%*	0.468%*	0.271%*	0.194%*
HEN3	0.309%*	0.146%*	0.146%*	0.418%*	0.334%**	0.165%**	0.203%**	0.357%**
IFX	0.094%*	-0.023%*	-0.130%*	0.111%*	0.035%*	-0.043%*	0.275%*	0.162%*
LHA	0.122%**	0.012%**	0.026%**	-0.037%**	0.084%**	0.071%**	-0.044%**	0.050%**
LIN	0.137%**	0.071%**	0.061%**	0.141%**	-0.064%**	0.031%**	0.058%**	-0.136%**
MAN	0.059%**	0.012%**	0.053%**	0.027%**	0.032%**	0.010%**	0.064%**	0.051%**
MEO	0.307%**	0.309%**	0.192%**	0.288%**	0.306%**	0.311%**	0.182%**	0.270%**
MRK	0.239%**	-0.053%**	0.123%**	0.163%**	0.172%**	-0.072%**	0.339%**	0.237%**
MUV2	0.055%**	-0.017%**	-0.012%**	0.003%**	0.047%**	-0.074%**	0.100%**	0.246%**
RWE	0.109%**	0.066%**	0.135%**	0.070%**	0.099%**	0.068%**	0.081%**	0.036%**
SAP	-0.165%**	-0.109%**	-0.118%**	-0.138%**	-0.199%**	-0.063%**	-0.171%**	-0.226%**
SDF	0.036%**	-0.134%**	0.035%**	0.100%**	0.004%**	-0.061%**	0.202%**	0.024%**
SIE	0.233%*	-0.123%*	-0.038%*	0.252%*	0.351%*	-0.100%*	0.127%*	0.293%*
TKA	0.024%*	0.163%*	0.102%*	0.261%*	0.037%*	0.163%*	0.177%*	0.316%*
VOW3	-0.132%**	-0.172%**	-0.140%**	-0.234%**	-0.022%**	-0.200%**	-0.204%**	-0.392%**
Mean	0.132%	0.048%	0.057%	0.125%	0.124%	0.060%	0.106%	0.123%
Median	0.151%	0.012%	0.050%	0.131%	0.116%	0.036%	0.114%	0.130%
nb positive	25	16	21	26	25	19	23	24
nb positive & signf	14	10	14	15	16	13	13	15
Bonferroni p-val	0.250	0.151	0.208	0.247	0.244	0.184	0.228	0.225

Note: This table presents the differences between the average daily cumulative returns of the unrestricted and restricted trading strategies when we trade at a latency of 500 milliseconds and at midquote prices with a κ of zero. A positive number implies that the unrestricted strategy provides, on average, higher daily cumulative returns than the trading strategy that uses the restricted forecasting model implied by hypotheses in the corresponding column heading. Mean and median are calculated over all stocks in our sample. nb pos and nb pos & signf present the number of stocks out of 30 in our sample for which the return difference is positive and significantly positive, respectively. Bonf. p-val is the p-value based on the Bonferroni correction for stocks with positive return differences and calculated as $\min(1, \min(p_1, \dots, p_k) \times k)$ where the minimum is calculated only over stocks for which the return difference is positive and, thus, k is the number of stock for which the return difference is positive. p_i is the p-value of the hypothesis testing whether the mean of daily cumulative returns of the unrestricted and a given restricted trading strategy is equal. ***, ** and * denote that the return differences are significantly different from zero at the 1%, 5% and 10% levels, respectively.