arXiv:1708.02715v1 [q-fin.TR] 9 Aug 2017

Order Flows and Limit Order Book Resiliency on the Meso-Scale

Kyle Bechler* Mike Ludkovski †

June 29, 2017 ‡

Abstract

We investigate the behavior of limit order books on the meso-scale motivated by order execution scheduling algorithms. To do so we carry out empirical analysis of the order flows from market and limit order submissions, aggregated from tick-by-tick data via volume-based bucketing, as well as various LOB depth and shape metrics. We document a nonlinear relationship between trade imbalance and price change, which however can be converted into a linear link by considering a weighted average of market and limit order flows. We also document a hockey-stick dependence between trade imbalance and one-sided limit order flows, highlighting numerous asymmetric effects between the active and passive sides of the LOB. To address the phenomenological features of price formation, book resilience, and scarce liquidity we apply a variety of statistical models to test for predictive power of different predictors. We show that on the meso-scale the limit order flows (as well as the relative addition/cancellation rates) carry the most predictive power. Another finding is that the deeper LOB shape, rather than just the book imbalance, is more relevant on this timescale. The empirical results are based on analysis of six large-tick assets from Nasdaq.

Keywords: limit order books, order flows, scarce liquidity, price impact

1 Introduction: the Meso-Scale

With the proliferation of the electronic trading, research on the fine time-scale of financial markets, exemplified by the microstructure of Limit Order Books (LOBs), continues to grow in scope and importance. The microstructure time-scale (on the order of milliseconds) works with the discrete events that correspond to the order messages submitted to the exchanges and hence captures the fundamental price formation and market behavior. It is often contrasted with the classical "diffusion-limit" framework (on the order of hours and days) that works with continuous-time processes and allows for infinitesimal trading. However, in between these two time-scales, there lies a third: the *meso-scale*, capturing the asset behavior on the order of minutes.

The meso-scale is of key importance for the market participants engaged in market-making and order execution. Consider the proverbial trader who executes a large order (an action originally driven by a macro-scale analysis) by dividing it into smaller "child" orders executed over

*UCSB, Email: kyle.bechler@gmail.com

†Corresponding Author: Department of Statistics and Applied Probability, University of California, Santa Barbara CA, 93106-3110 USA. Email: ludkovski@pstat.ucsb.edu

‡We thank Sebastian Jaimungal for many useful discussions and for help with the datasets. We are also grateful for feedback from the participants at the 2015 Stevanovich Conference on Market Microstructure and High Frequency Data (Chicago) and the 2016 Microstructure: Confronting Many Viewpoints Conference (Paris), where earlier versions of this work were presented.

multiple "slices". She is concerned with execution costs comprised of instantaneous slippage and longer-term price impact. The first component is addressed via the *routing* layer that optimizes individual trades, for example to avoid "walking the book". The second component is addressed by the *scheduling* layer that determines the pace of execution and usually works over 2–10 minute intervals. This scheduling revolves around the expected evolution of the LOB through time, requiring aggregation of the information from the static LOB snapshots and lifting it to the meso-scale.

The dynamics of the book on the meso-scale, and the associated consumption and provision of liquidity drive the concept of resilience, also known as the transient and permanent price impacts, that underlies much of the literature on algorithmic trading and order execution. On the tick-by-tick scale, the effect of an order is mechanical and can be easily described using observable LOB characteristics, namely the resting limit volume at various levels, summarized via LOB depth and shape. In particular, for vast proportion of executions direct price impact is zero, as only a portion of the respective top queue is consumed. However, on a longer time-scale (as implemented in order splitting algorithms that typically unfold over several hours), the enormous volume of the order messages obscures the link between price evolution and individual orders, and renders static liquidity measures irrelevant. Consequently, proper calibration of the HFT models ought to rely on meso-scopic/dynamic, rather than microscopic/static LOB metrics.

To measure liquidity on the mesoscale it is necessary to consider aggregated metrics about types/volumes of orders submitted. A natural quantity of interest therefore becomes the *order flows* which can be viewed as the dynamic version of instantaneous LOB characteristics. Statistical analysis of the order flows is the main focus of this article. We analyze the mesoscopic links between market and limit order flows, and between order flows, prices and static LOB metrics. Our main goal is a phenomenological description of these relationships, so as to a provide a data-driven perspective for LOB behavior. For empirical work we utilize a cross-section of 6 liquid, large-tick Nasdaq tickers spread across 2 calendar years.

Based on our analysis, we report several key findings. First, we document an S-shaped non-linear relationship between mid-price change ΔP and market order flow VM. In tandem, we also record a strong linear relationship between ΔP and Net Liquidity flow, which is a weighted average of VM and limit order flow VL. These empirical results offer an important correction to the commonly assumed framework of linear price impact: to obtain linearity one must take into account touch limit orders, weighted appropriately. Ignoring limit flow dramatically weakens the explanatory power of the models, in fact our findings suggest that limit flows are *at least, if not more* statistically significant than the market trades.

Second, we use statistical tools to identify key predictors for price formation on the meso-scale. We find that the explanatory power of the top-level LOB queue depths is rather weak. This is primarily due to the rapid fluctuations in those quantities, which do not allow for easy time-aggregation. Instead, liquidity measures that look deeper into the book are more relevant for quantifying liquidity. At the same time, we find that the LOB metrics offer an important *modulation* to the effect of order flows and do add a non-negligible predictive power. Third, we disentangle price trend (defined in terms of overall bid/ask-side pressure captured by market order flow VM) from liquidity (primarily captured by limit order flows VL). This allows us to pinpoint the meaning of *scarce liquidity*. We carry out further statistical analysis in that direction, which is of interest for adaptive order scheduling that can respond to varying expected price impact.

The meso-scale is intrinsically defined via aggregating the raw LOB data into slices, and the question of aggregation method looms large. The micro-scale is naturally discrete-event-based. On the macro-scale, the high degree of averaging lends itself naturally to a (calendar) time-based limit. Here we propose that *volume-based* aggregation is most appropriate for the meso-scale.

Thus, our analysis is based on working with LOB buckets, with each bucket containing the same amount of executed (market) volume. As we show, this yields a desired stabilizing effect on the LOB data and removes some known artifacts. It also quantifies what we mean by the meso-scale; informally it can be characterized as dealing with "dozens of trades"; indeed our buckets contain 10–100 market executions, accompanied by 100–1000 limit order events. At that intermediate time-scale, the aforementioned event-based analysis is less tractable for capturing the dynamic quantities that come to the fore; while time-based analysis must contend with huge swings in market activity that are caused by extreme clustering of LOB events.

As far as we are aware, this is the first academic study of volume-aggregated intra-day LOB data. Several previous studies used fixed-interval time-aggregation, such as 15-min slices in [21] and 5-min slices in [10], [20] and [5]. The closest study to ours is by Cont et al. [12] who also investigated the relationship between order flows, price change and book depth, and reached some similar conclusions. However, that study was limited to working only with top-level book information, and most crucially *postulated* the expected (linear) link between the different variables. In contrast, we employ non-parametric tools to obtain data-driven findings about variable importance, empirical non-linearities, et cetera. Moreover, we also investigate the asymmetric features of the order flows which are obscured by the total netting carried out in [12].

Our investigation offers a statistical counterpoint to the research on optimal execution which has put forth a plethora of different *theoretical* price impact models. From the empirical point of view, there is already a growing literature on tick-by-tick data, see [12, 14, 17, 16, 34, 35]. These are complemented by LOB *simulators* that aim to capture short-term behavior and price impacts. The most popular framework views the LOB as a collection of interacting queues [22, 29, 28, 30]. For example, Huang et al [24] model limit order, cancellation, and market order arrival rates as a function of the queue sizes. Unfortunately such models are much less tractable on the meso-scale, so few analytic results are so far available. An alternative way to model the mesoscopic price impact is to infer the impulse response function, known as the propagator, of an *individual* order/trade. The aggregate effect on the mid-price is then viewed as a linear superposition of individual impacts, see [16, 7]. However, such event-by-event analysis largely excludes capturing the statistical effects of observable covariates which "dress" the evolution of the respective dynamical system describing the LOB. Yet another micro-scale analysis which also leads to important meso-scale implications is the emerging strand on long-memory effects of micro-structural phenomena [32], as well as time-series approaches to the functional evolution of the entire LOB shape [20].

The remainder of the paper is structured as follows. In Section 2 we discuss LOB evolution in the context of static and dynamic measures of liquidity and price impact, as well our volume-bucketing method. Sections 3-4 contain our main statistical analysis regarding the meaning of price trend, liquidity, and scarce liquidity. Section 5 takes a different tack, addressing the time-series properties of order flows. Finally, Section 6 discusses ramifications from the previous sections and offers an outlook for further work.

2 Limit Order Books: Price Impact, Order Flows and Liquidity

2.1 Limit Order Book

Electronic trading marketplaces match liquidity providers and consumers via the limit order book (LOB). Participants asynchronously submit trading messages which are aggregated by the exchange into the LOB. The two base classes of trades are market orders and limit orders. Market orders, henceforth called *trades*, indicate actual transactions taking place and are denoted as

$\mathfrak{M} := \{(T_i^M, O_i^M)\}$, where T_i^M are the execution times, and O_i^M are corresponding execution volumes. O_i^M's are signed, with positive indicating a buy order and negative a sell order. Limit orders, henceforth *orders*, are $\mathfrak{L} := \{(T_i^L, O_i^L, S_i^L)\}$ where T_i^L are the message time stamps, O_i^L are (signed) order volumes and S_i^L is the limit order price. Positive values of O_i^L indicate submission of a new limit order, while negative O_i^L indicates a cancellation. Limit orders reside in the LOB until executed by an incoming market order, or being cancelled by the participant who initially placed the order. We refer to [19, 11] for more general introduction on the functioning of LOBs, and other types of orders possible.

At any given moment t, the LOB lists all the resting limit orders, represented as a vector $(p_i^j(t), v_i^j(t))$, listing the volume $v_i(t) \geq 0$ of resting limit orders at price $p_i(t)$. The superscript $j \in A, B$ denotes the Ask and Bid sides, respectively. The subscript $i = 1, 2, \ldots$ indexes the price level which is discretized in terms of the tick size Δp, typically 1 cent for US equities. Indexing is done separately for each side of the LOB, starting from the best-bid and best-ask levels and moving consecutively. Thus, $p_1^A(t)$ is the best-ask (at the touch) price, $p_2^A(t) = p_1^A(t) + \Delta p$ is the second price level, etc. Some queues can be empty, $v_i^j(t) = 0$. However the top queues $v_1^j(t) > 0$ are always strictly positive, and after netting executable limit orders, the best-ask price is at least one tick above the best-bid $p_1^A(t) \geq p_1^B(t) + \Delta p$. Based on the bid/ask prices, we have the midprice P and the spread S:

$$P(t) := \frac{p_1^A(t) + p_1^B(t)}{2}, \qquad Spr(t) := p_1^A(t) - p_1^B(t). \tag{1}$$

2.2 Static LOB Snapshots and Price Impact

A starting point for understanding LOB liquidity is the information contained in a static snapshot of the LOB. The most basic measure of liquidity is the spread $Spr(t)$ which dictates the costs of a round-trip trade that is fundamental to market-making. However, for liquid US equities the spread may not be useful, as it is not sensitive to changing market conditions. In particular for "large tick" assets, the bid-ask spread is almost always equal to one tick. This is the case for all the assets in our sample, where $Spr(t)$ is essentially always 1 or 2 ticks – and the 2-tick spreads are short-lived: after opening up the spread quickly closes as an aggressive limit order is posted inside the spread.

Another commonly used measure is the volume-at-the-touch $v_1^j(t)$, i.e. queue lengths at the best-offer and best-bid price levels. Over very short time horizons, v_1 indicates the immediate liquidity available for incoming market orders, and hence drives the direction and likelihood of a mid-price change. Recall that P changes by a half-tick when v_1^j hits zero, whereupon v_1^j is reset to the posted volume at the new top LOB level. A related quantity that measures the relative "strength" of each side of the book is the Book Imbalance

$$BI := \frac{v_1^A - v_1^B}{v_1^A + v_1^B} \in (-1, 1). \tag{2}$$

The book imbalance BI has been shown to be predictive of the next order side and resulting price move [14, 12, 9]. Consequently, at the slicing/routing level of the execution process, BI is often used as an indicator of when to employ limit orders and when to cross the spread and place a market order [33, 28]. In a similar vein, [24, 30, 26] document a link between BI and order arrival rates.

The top queues offer only a "tip of the iceberg" summary for the overall LOB shape. The direct procedure to describe LOB shape employs units of shares and measures cumulative *depth*

over the first k levels

$$D_k^j := \sum_{i=1}^{k} v_i^j, \qquad j = A, B, \quad k = 1, 2, \ldots. \tag{3}$$

Thus, $D_1^j = v_1^j$ is the top-level depth, $D_2^j = v_1^j + v_2^j$ is the total volume posted at the top two levels and so on. A complementary inverse procedure measures *execution cost* in units of ticks. Execution cost quantifies the slippage from immediately making a market trade, converting depth (shares) into monetary terms. Fix a quantity N of shares and let $\bar{i}^j(N) := \max\{k : D_k^j(t) < N\}$. Then

$$PI_N^j(t) := \frac{1}{N}\left\{ \sum_{i=1}^{\bar{i}^j(N)} v_i^j(t)\left(p_i^j(t) - P(t)\right) + (N - D^j_{\bar{i}^j(N)}(t))\left(p^j_{\bar{i}^j(N)+1}(t) - P(t)\right)\right\}. \tag{4}$$

Thus, PI_N^j blends information about the shape of the LOB by computing weighted average cost (relative to the mid-price) per share of immediately executing N shares in the j-th direction. The map $N \mapsto PI_N^j(t)$ then summarizes the shape of the book in terms of executing different order sizes. Intuitively, $PI \propto 1/D$ is inversely proportional to depth.

Remark 1. We stress that $PI_N(t)$ is a theoretical construct, based on the counterfactual that an immediate market order is placed exactly at t. In practice the vast majority of market trades have zero impact, consuming less than the standing liquidity at the first level. Moreover, latency delays make it difficult to realize the "observed" PI in real-life trades.

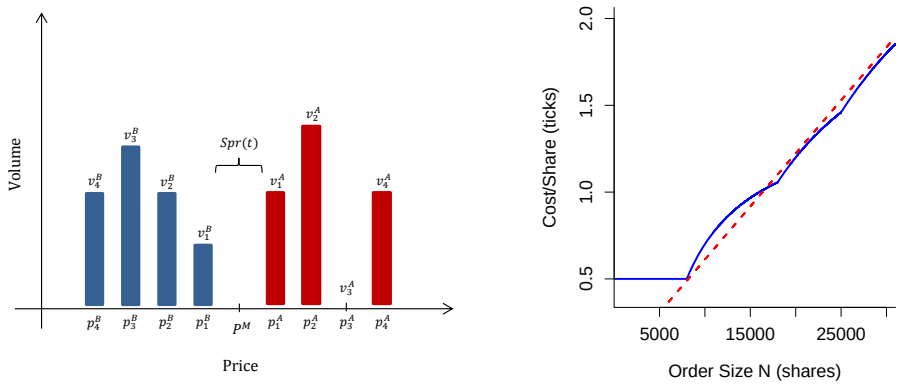

Figure 1: *Left:* Stylized limit order book. In this illustration the spread is two ticks $Spr(t) = 2\Delta p$ and book imbalance is positive $v_1^A > v_1^B \Rightarrow BI(t) > 0$. *Right:* Measuring execution cost in a hypothetical limit order book. The underlying queue lengths are taken to be $(v_1, v_2, v_3, v_4) = (8000, 10000, 7000, 15000)$. We plot $n \mapsto PI_n$ (solid) and the fitted line $n \mapsto \hat{S} \cdot n$ based on the impact slope $\hat{S}$ (dashed).

The right panel of Figure 1 illustrates the above computations for a hypothetical book with $v_1 = 8000, v_2 = 10000, v_3 = 7000, v_4 = 15000$. The solid curve shows $N \mapsto PI_N$. Note that (4) implies that for $N \le v_1$, $PI_N = Spr/2$ is half the spread; for larger order sizes, PI_N has a nonlinear shape due to the N-dependent averaging weights. For example, at $N = 30,000$, we obtain $PI_N = 1.8$ ticks, which averages the cost of 0.5 ticks for the first 8000 shares, 1.5 ticks for the next 10K shares, 2.5 ticks for the 7K thereafter and 3.5 ticks for the last 5K shares.

One can observe from the plot that $N \mapsto PI_N$ is asymptotically linear, which would be consistent with the book having an asymptotic "rectangular" shape with $\bar{v}$ volume per level. (Namely, $\bar{v} \simeq N/(2PI_N)$ which for the computation above yields $\bar{v} = 8333$ taking $N = 30K$.) Motivated by this idea, [10] have proposed to compute an LOB slope S^j which is obtained from a linear regression model

$$PI_n^j = S^j \cdot n + \epsilon, \quad n = 1, 2, \ldots, N, \tag{5}$$

estimated using the standard least-squares fit. The slope coefficient S^j can be interpreted as a linearized price impact (on the j-th side of the book) per share and allows to imply a single liquidity metric from the collection $(PI_n^j)_n$. In Figure 1 the estimated impact slope is $\hat{S} = 0.0611$ ticks/1000 shares, see the fitted dashed line $\hat{S} \cdot n$. $\hat{S}$ can be inverted to map into "averaged" book depth $\bar{v} = 0.5/\hat{S} = 8176$. An advantage of the metric S is the ability to average v_k's over multiple levels and incorporate the spread Spr, which makes it less sensitive to short-term fluctuations that dramatically lower the statistical usefulness of depth measures D_k^j.

2.3 Order Flows and LOB Evolution

To visualize the meso-scopic behavior of the book, Figure 2 shows an event-by-event summary of an LOB for a representative 90-second period. The main plot at the bottom shows the top-level queue lengths v_1^j driven by the market orders (in orange) and limit orders at-the-touch. At the top of Figure 2 we also track the contemporaneous effect on the mid-price, with dotted vertical lines marking a change in bid- or ask-price.

The Figure indicates that following a market execution the LOB response varies widely, at times bouncing back through fresh posted volume, while at other times falling through and fading. We observe several "regimes": episodes of strong resilience when market orders are counteracted with added limit orders (so that $v_1^j(t)$ stays roughly constant over time), and other periods where market orders are accompanied primarily by limit order cancellations, creating a strong negative trend in $v_1^j(t)$, known as an LOB fade. This happens on the left of Figure 2, where the arriving market flow is heavily tilted towards sell orders $TI < 0$ and net limit flow to the bid-side touch is clearly negative. The result is a rapid drop in mid-price, creating *scarce* liquidity for sellers and increasing their respective execution costs. A milder form of this is when few new limit orders are added at the touch, so that the book is not replenished fast enough and "retreats" along with executed trades. Our goal below is to find meso-scopic predictors that explain whether market orders are likely to be met with strong resilience via new limit orders, or minimal liquidity provision or even net cancellations.

2.4 Volume Bucketing

To move from the micro- to meso-scale we divide the trading day into *buckets* that are used for aggregation. Buckets are defined in terms of their start- and end- time-stamps $\{\tau_k\}$, with the k-th bucket consisting of all messages entered during $(\tau_k, \tau_{k+1}]$ (special treatment possibly accorded to the first and last message, see below). Slice times τ_{k+1} are taken to be the time-stamp of the final trade included in bucket k. The resulting bucketed quantities are:

$$VM_k^{A,B} := \sum_{i:\tau_k < T_i^M \le \tau_{k+1}} |O_i^M| 1_{\{O_i^M \gtrless 0\}}; \tag{6}$$

$$VL_k^{j,(1)} := \sum_{i:\tau_k < T_i^L \le \tau_{k+1}} O_i^L 1_{\{S_i^L = p_1^j(T_i^L)\}}, \qquad j = A, B. \tag{7}$$

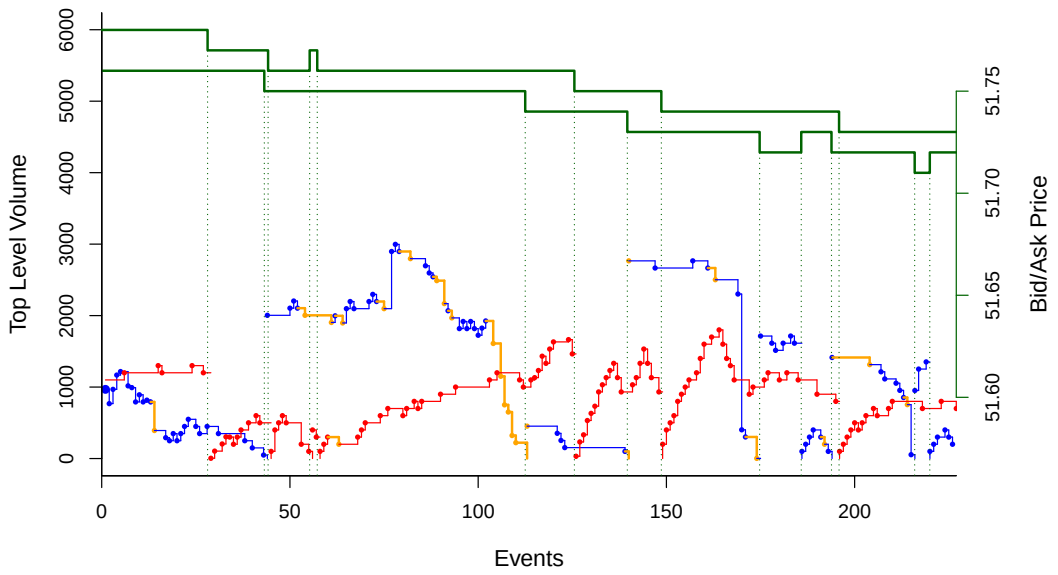

Figure 2: Best bid/ask queue levels $v_1^j(t)$ for TEVA (bottom curves) along with bid/ask price $p_1^j(t)$ (top). Event-by-event data taken from a 90 second window beginning at 2:30pm on 2/18/2011. Limit orders are in red (ask-side) and blue (bid), market executions in orange. All orders are plotted event-by-event (in particular without re-combining split market orders), so that the mechanics and sequence of order arrivals are more clear.

Thus, VM^j are the respective aggregated volumes of executed market buys and sells, and $VL^{j,(1)}$ are the volumes of limit orders at the touch. Note that $VL^{j,(1)}$ tracks the mid-price throughout the bucket, aggregating the respective top-level (at the time of message submission) limit orders. Because the tickers we consider are highly liquid in the sequel we focus exclusively on top-level limit orders $VL \equiv VL^{(1)}$. Recall that market orders capture the consumption of liquidity by aggressive orders and decrease top-level depth v_1^j, while limit orders can either augment v_1^j (liquidity provision) or decrease it (order cancellations). Consequently, $VM^j \geq 0$ is non-negative by construction, whereas VL^j can be of either sign depending on the net limit additions/cancellations.

Given a slice (τ_k, τ_{k+1}) we define the corresponding price change as

$$\Delta P_k = P(\tau_{k+1}) - P(\tau_k). \tag{8}$$

Based on the buckets we also define the discrete snapshots of LOB statistics through indexing by k; for example the book imbalance $BI_k \equiv BI(\tau_k)$ or the depth $D_k \equiv D(\tau_k)$. As analogue to the static BI, we also define the normalized *trade imbalance*

$$TI_k := \frac{VM_k^B - VM_k^A}{V_k}, \tag{9}$$

where $V_k = VM_k^A + VM_k^B$ is the total executed volume in the bucket. By definition, $TI_k \in \{-1, 1\}$; if $TI_k = 1$ then all executed trades in the bucket were 'Buys', and if $TI_k = -1$, all executed trades were Sells. Trade imbalance captures the aggregate supply and demand for the asset and is often used as the basis for measuring price impact [6, 10].

We propose to bucket in terms of *executed market volume*, namely keeping $V_k \equiv V$ fixed across buckets. Technically, this requires that market orders are sometimes split in two as one bucket "fills" up and the next begins. In our main analysis we take $V \simeq \{0.25\%, 1\%, 2\%\}$ of average

daily volume (ADV, based on the considered time period of 10:00am–3:45pm) for the respective ticker. This means that *on average* there are about 400 (resp. 100, 50) buckets per day; actual number of buckets can be 30%–500% of the average due to fluctuating levels of market activity.

Note that slicing is done in terms of executed volume; volume of limit orders $VL_k^A + VL_k^B$ remains random across buckets. The reason for this is that even though market orders account for only $2-4\%$ of total trades, they indicate actual transactions taking place and hence ultimately drive traders' P&L. Thus, due to their intrinsic nature of "putting money on the table", they are typically viewed as influential by other participants and carry the most information for defining the "business-time" for the ticker.

Most related studies have used time-based aggregation. For instance, [12] used 10-second slices, [5, 10, 20] used 5-min slices. Nevertheless, using volume buckets has several attractive properties for our purposes. First, slicing by trade volume rather than time yields consistency in information across buckets. Under clock-time bucketing, one can easily appreciate the difficulty of comparing buckets with dramatically different activity levels (in the extreme, a time-based bucket could be entirely empty). Second, working in volume-time also reduces the intra-day seasonality effects, such as volume and volatility clustering, and improves statistical properties (such as Gaussianity) of the (VL_k, VM_k) time-series. Third, comparison of limit order activity across buckets is naturally normalized when measuring in (trade) volume time, i.e. the stable quantity is not VL_k per se, but rather the ratio VL_k/V_k (or VL_k/VM_k). For example, it is intuitive that limit additions to the touch depend more on quantity of market trades than passing minutes. Fourth, our method avoids the difficult challenge of estimating trade volume effects that are present in trade-by-trade aggregation.

2.5 Datasets

For our statistical analysis we use Nasdaq ITCH TotalView data which contains "Level-2" information on order book events. Specifically, we have access to direction and size of market executions, as well as limit order additions, modifications and cancellations for $i \leq 30$ levels (ticks) into the current book. To avoid synchronization issues, we work only with Nasdaq-based messages, leaving out information from other exchanges. Furthermore, since erratic LOB behavior is often seen near the market open and close, we consider only activity between 10:00am and 3:45pm. Other pre-processing included removing all executions against hidden orders (less than 10% of executed volume) and aggregating executed trades. Market orders are often matched against several smaller limit orders; we re-created the size of the original market order by aggregating orders that were consecutive, in the same direction, and with identical time stamps.

To give a broad slice of market activity, we use a total of six different tickers, covering two distinct time-periods. Thus, we used data on MSFT, TEVA and BBBY from the first 100 days of 2011, and data on INTC, ORCL and NTAP from the last 100 days of 2013. All assets are categorized as large cap stocks, with market caps ranging from about 12B USD for BBBY, and up to 240B for MSFT. Tables 1-2 provide summary info for each stock. We note that all stocks are highly liquid, with several thousand trades per day, and hundreds of thousands of limit orders. Moreover, the analyzed LOBs are quite deep, with average depth of 3–10 average event size (AES). As a result, most of the economically relevant information can be gleaned from the top-level events. These still correspond to 50K–100K touch limit orders per day, and respective nominal order volume $\sum_k |VL_k| \simeq$ 5M–100M shares.

To illustrate how volume aggregation works, we describe it in the case of ORCL. We use $V = 20000$ which is about 1% of the average daily volume of 1.84M shares traded during 10:00am–3:45pm. In total, we obtain 9229 aggregated buckets over the 100 days of data. Since the average

	MSFT		TEVA		BBBY	
	mean	(stdev)	mean	(stdev)	mean	(stdev)
Mean Price (USD) P	26.29	(1.25)	50.34	(2.73)	50.51	(3.64)
Mean Abs Price Change (ticks)	1.80	(1.46)	3.33	(3.07)	3.67	(3.27)
Ave Daily Vol (shares) VM	$10.0M$	$(4.61M)$	$1.21M$	$(1.06M)$	$0.64M$	$(0.30M)$
Num of Daily MO's (events)	5340	(2173)	5455	(2484)	4281	(1483)
Ave Event Size (shares)	1643	(3612)	293	(596)	192	(321)
Ave Limit volume (shares) VL	$116.4M$	$(42.6M)$	$12.04M$	$(5.90M)$	$6.60M$	$(2.58M)$
Num of Daily LO's (events)	$292K$	$(92K)$	$110K$	$(45K)$	$77K$	$(31K)$
Spread (ticks) Spr	1.01	(.067)	1.13	(.362)	1.29	(.599)
Mean depth at Level 1 D_1	$25.5K$	$(27.9K)$	$1.33K$	$(2.62K)$	$0.63K$	$(1.09K)$
Mean depth at Levels 1+2 D_2	$59.2K$	$(29.9K)$	$2.80K$	$(7.4K)$	$1.24K$	$(0.97K)$
Ave Slope (ticks/000's shares) S	.0274	(.00976)	0.691	(0.305)	1.37	(0.663)

Table 1: Summary statistics (mean/standard deviation) for MSFT (Microsoft), BBBY (Bed Bath & Beyond) and TEVA (Teva Pharmaceuticals) for the 100 trading days from 1/2/2011 to 5/25/2011. Average price change is per 1%-ADV volume buckets.

trade size for ORCL is about 800 shares, a typical slice has about 25 trades, and about 500 limit orders. Figure 3 shows a histogram of net limit order-flow VL_k (using combined Bid and Ask data) for ORCL. As expected, average limit flow at the touch is typically positive, i.e. there are more orders entered than cancelled, providing liquidity against executed trades. For ORCL at that aggregation frequency about 5.5% of the buckets had negative VL^j; this number fluctuates in the 2–10% range across the 6 assets. Also observe that typically the nominal limit order flow VL is larger than executed volume V, hinting at the high rate of "churn" in the limit orders, even after netting additions and cancellations. The middle panel of Figure 3 shows the distribution of the ORCL bucket durations, $\tau_{k+1} - \tau_k$; the median duration is about 170 seconds. Note the strong skew with many buckets that last for 5 or even 10 minutes; the longest bucket covered 29 calendar minutes (during a particularly slow lunchtime on day 19). Notably, the correlation between VL_k and $\tau_{k+1} - \tau_k$) is statistically 0, implying minimal relationship between liquidity provided and physical time elapsing within a volume slice. Appendix A provides a summary of the bucketing employed for each asset.

3 Price Trend and Liquidity State

As discussed, at the meso-scale the most important quantities are the order flows, i.e. VM and VL from (6)-(7). In this section we document their relationship to the fundamental financial output of the LOB: the mid-price change ΔP_k.

Based on the price formation mechanism originating in the LOB, the changes in the mid-price are driven by the respective liquidity consumption and provision. A simple narrative ties ΔP_k to the market demand for the asset, exemplified by the trade imbalance TI_k. Self-evidently, if there are more market buys than sells, the price is expected to rise, and if $TI_k < 0$ we expect $\Delta P_k < 0$. A basic model assumes a *linear* relationship between ΔP and TI. A theoretical justification for this is provided in the Obizhaeva and Wang [31] framework and uses a combination of permanent and transient price impact to link executed volume to price move. The basic premise is that limit

	INTC		ORCL		NTAP	
	mean	(stdev)	mean	(stdev)	mean	(stdev)
Mean Price (USD) P	23.66	(0.98)	33.84	(1.28)	41.26	(1.35)
Mean Abs Price Change (ticks)	1.31	(1.10)	2.11	(1.79)	2.71	(2.27)
Ave Daily Vol (shares) VM	$3.58M$	$(1.18M)$	$1.84M$	$(0.80M)$	$0.71M$	$(0.31M)$
Num of Daily MO's (events)	2058	(609)	2880	(1192)	2171	(935)
Ave Event Size (shares)	1740	(3819)	837	(1257)	252	(287)
Ave Limit Volume (shares) VL	$99.62M$	$(33.42M)$	$40.30M$	$(15.96M)$	$15.80M$	$(5.26M)$
Num of Daily LO's (events)	$263K$	$(77K)$	$190K$	$(68K)$	$136K$	$(46K)$
Spread (ticks) Spr	1.01	(0.13)	1.03	(0.17)	1.08	(0.29)
Mean depth at Level 1 D_1	$13.9K$	$(10.0K)$	$3.9K$	$(2.8K)$	$1.24K$	$(1.13K)$
Mean depth at Levels 1+2 D_2	$31.9K$	$(17.9K)$	$8.5K$	$(4.8K)$	$2.69K$	$(1.28K)$
Ave Slope (ticks/000's shares) S	.0367	(.0158)	0.136	(.0659)	0.432	(0.201)

Table 2: Summary statistics (mean/standard deviation) for INTC (Intel), ORCL (Oracle), and NTAP (Net Appliances) for the 100 trading days from 8/9/2013 to 12/31/2013.

orders refill the book over time, but take time to be entered. Therefore, price changes are driven by "too many too fast" trade executions, and quantified by the LOB resilience (e.g. the rate of the respective exponential recovery), which yields the sensitivity of ΔP to TI.

The left panel of Figure 4 shows a scatterplot of ΔP against TI over buckets of 1% ADV (12,000 shares) for TEVA. While a positive relationship can be seen as expected, the link is clearly *non-linear*. In fact, there is an S-shape behavior, whereby for strongly one-sided markets, TI close to ± 1, the expected price change becomes constant or even declines. This puzzling result in fact indicates intelligent order-placing: traders take into account the state of the LOB when executing orders and hence induce dependency between order flow and price moves. In other words, 12,000 shares executed consecutively in the same direction is an infrequent occurrence for TEVA, and is usually precipitated by a deep/resilient book at that moment. Another important feature of Figure 4 is the dispersion in the observed price change ΔP_k conditional on concurrent trade imbalance TI_k. For example, conditional on $TI_k = 0.5$, we observe that $\Delta P_k | TI_k = 0.5$ ranges from -0.08 to 0.16 for buckets across our sample period, which is financially material. This indicates that on its own trade imbalance has low predictive power, i.e. price formation is driven by much more than the net executed volume.

To statistically capture the relationship between price change over each volume slice $\Delta P_k = P_{k+1} - P_k$ and the concurrent trade imbalance TI_k we fit a nonparametric model of the form

$$\Delta P = g(TI) + \epsilon, \tag{10}$$

where g is the link function to be estimated and ϵ_k are i.i.d. regression residuals. Specifically to account for the observed non-linear dependence we use a generalized additive model (GAM), taking $g(\cdot)$ in (10) to be a penalized regression spline that is fit via cross-validation. The analysis is done using the `R` package `mgcv` [36] and is carried out for each stock on volume buckets from the entire sample period, excluding the rare buckets where a single large market order filled the entire bucket (leading to a bucket with zero duration). An example of the resulting fit $TI \mapsto g(TI)$ is shown in the left panel of Figure 4 illustrating the aforementioned S-shape.

Figure 5 shows a cross-section of fitted GAM models (10) across assets (left panel) and bucket sizes V (right panel). We find that the S-shape observed in Figure 4 is highly persistent, so

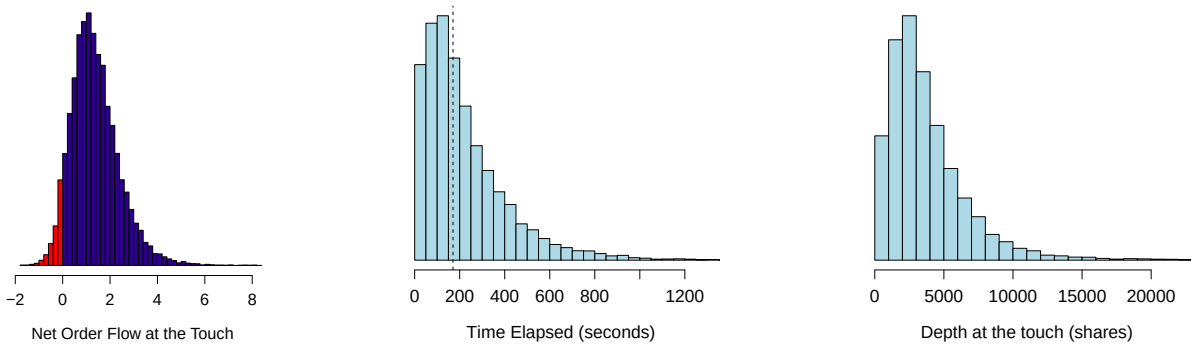

Figure 3: *Left:* Volume-normalized limit order flow at the touch VL_k/V (net of additions and cancellations). About 5.5% of buckets had negative VL indicated in red. *Middle:* Time elapsed during volume buckets $\Delta\tau_k$; the mean duration is indicated by the dashed line. *Right:* one-sided depth D_1^j at the touch (combined Bid/Ask information). Figures drawn for ORCL over 100 trading days with buckets of $V = 20K$ traded volume (a total of 9229 buckets).

that it is a *fundamental stylized feature* of the liquid, large-tick LOBs under consideration. In order to compare the fitted g's, we normalize the y-axis measuring price changes by the respective mean absolute price move. This yields a consistent way of mapping trade imbalance, which is by definition in the range of $\{-1, 1\}$ to the relative price move $\Delta P_k/Ave(|\Delta P_k|)$.

The above findings are in contrast to the widely used linear price impact framework that can be traced back to [31] and has since appeared in numerous optimal execution/market-making models, e.g. [1, 18, 2]. Part of the discrepancy can be attributed to our volume-based bucketing which through normalizing TI to be in $[-1, 1]$ crystalizes the underlying nonlinearity. From a different direction, our finding is consistent with the documented volume-concave impact of *individual* market trades which implies a concave relationship between $|TI|$ and ΔP [16].

Remark 2. One approach to improve the goodness-of-fit of (10) is to fit a separate regression model for each trading day. Indeed, it is often assumed that the LOB behavior is driven by a daily cycle. We found that while daily models do reveal significant day-to-day variation in the shape of the best fit curve (cf. [10]), the resulting R^2 is only marginally better compared to fitting a single model over all days. For example, the R^2 based on individual day fits for MSFT 1% ADV improves from 46.9% in Table 3 to 53.3%. So the significant noise in (10) is not simply the result of inter-day variation, but rather shows that liquidity provision is stochastic on an intra-day basis. Furthermore, checking the residuals $\widehat{\epsilon_k}$'s from (10) fitted to the whole sample period, we find that they are reasonably Gaussian and exhibit minimal temporal autocorrelation.

3.1 Limit Flow

Fundamentally, price is driven by the *net* liquidity supply; looking just at traded volume as was done in (10) only considers liquidity consumption. Thus, a natural next step is to look at both market orders that consume liquidity, and limit orders that provide liquidity. For the latter, we concentrate on orders added/cancelled at the present bid/ask touch which are by far the most relevant given the liquid LOBs under consideration. To capture this idea, we fit the linear model

$$\Delta P = \alpha_0 + \alpha_1 TI + \alpha_2 (VL^B - VL^A) + \epsilon, \tag{11}$$

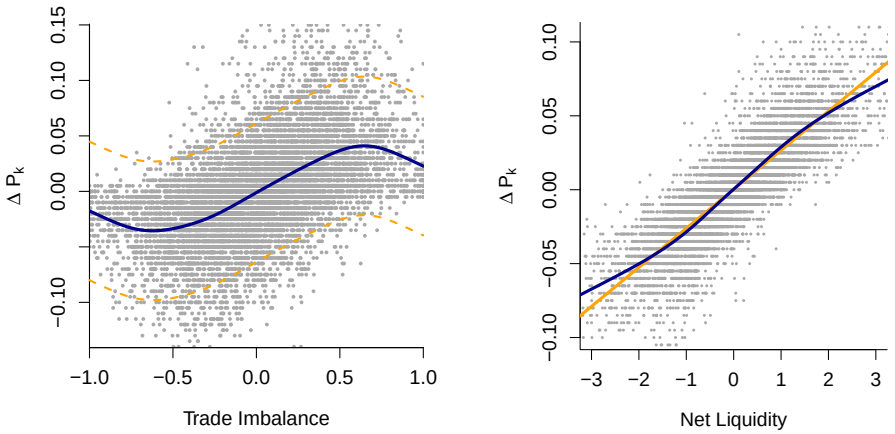

Figure 4: *Left*: Price change ΔP_k plotted against trade imbalance TI_k for TEVA for the first 100 days of 2013 (9954 buckets with $V = 12000$). The fitted curve is based on (10). *Right:* Price change against net liquidity provision for NTAP over buckets of $V = 8000$. We plot the linear fit from (11) (light line) and the non-parametric GAM model $g(\cdot)$ based on $\Delta P = g(NetLiq) + \epsilon$, with $NetLiq$ from (12) (dark curve). Since $\Delta P_k \in \{0.005k : k \in \mathbb{Z}\}$, the y-axes are discretized by half-ticks.

	Only TI (10)			Net Liquidity (12)		
Ticker	0.25%ADV	1%ADV	2%ADV	0.25%ADV	1%ADV	2%ADV
BBBY	0.290	0.270	0.243	0.493	0.596	0.590
INTC	0.303	0.403	0.421	0.371	0.682	0.767
MSFT	0.370	0.469	0.485	0.579	0.785	0.765
NTAP	0.293	0.315	0.302	0.470	0.712	0.751
ORCL	0.320	0.314	0.276	0.449	0.708	0.750
TEVA	0.269	0.276	0.269	0.520	0.607	0.592

Table 3: R^2 scores from the nonparametric fit against TI in (10) compared to R^2 scores from a linear fit against TI and VL in (11).

which connects price change to net executed volume and net limit order flow. The main finding is that (11) yields an *excellent* fit. Moreover, the linearity in (11) is *intrinsic*, i.e. taking into account the limit flows VL completely removes the previously documented S-shape. Let us define

$$NetLiq := TI + \beta \cdot VL, \tag{12}$$

where $\beta = \alpha_2/\alpha_1$ is the ratio between the impact of a market order vis-a-vis that of a touch limit order. The right panel of Figure reffig:SvsTI shows (for a different ticker, to showcase the generality of the phenomenon) that $NetLiq$ captures the fundamental price formation story of liquidity provision, yielding $\Delta P \propto \Delta NetLiq$ (the estimated intercept coefficient α_0 in (11) is statistically indistinguishable from zero). In the Figure we also compare the linear and GAM-based fits of ΔP against $NetLiq$ which are found to be nearly identical, confirming the linearity postulated in (11)).

Incorporating order flows not only removes the nonlinear price response, but also dramatically improves the goodness-of-fit, see Figure 4 again. Table 3 compares the R^2 scores from (10) and

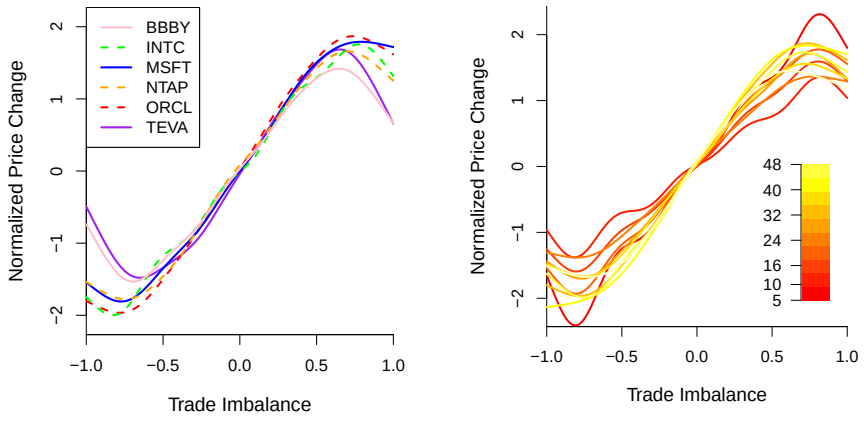

Figure 5: Relationship between normalized price change $\Delta P_k/Ave(|\Delta P_k|)$ and bucket trade imbalance TI_k. *Left:* six tickers across the two sample periods; *Right:* for ORCL across 10 different bucket sizes V.

(11). We observe a major improvement from about $R^2 \simeq 0.4$ when using only TI as a predictor for ΔP, to $R^2 \simeq 0.7$ when using both TI and VL^j. This suggests that the predictive power of VL^j is *at least as significant* as of the executed volume. Table 3 also shows that goodness-of-fit improves as we consider larger bucket sizes, implying that the other predictors get averaged out over time (i.e. are primarily about the microstructural effects of the LOB).

While mechanically net liquidity is simply $VM^A - VM^B + VL^B - VL^A$, it is commonly assumed that the price impact of limit orders should be less than that of market orders. This is precisely the reason for using separate coefficients α_1, α_2 in (11). The resulting ratio β that is used to construct $NetLiq$ in (12) is another *fundamental stylized feature* of the meso-scale. Table 4 reports the estimated $\hat{\beta} := \widehat{\alpha_2}/\widehat{\alpha_1}$ across assets and bucket sizes, which can be observed to be highly stable across all cases. As expected, the respective price impacts are highly non-equal, namely the relative effect of a limit order is only about 50%–70% of a market order. Again, we observe that β is larger for larger V's: on the macro-scale price formation is mechanically determined by total liquidity supplied/consumed; on the meso-scale there are further effects coming from LOB shape and shorter-term patterns.

Remark 3. We also tested for the relative price impact of limit order additions and cancellations by further decomposing into $VL^j = VL^{j,+} - VL^{j,-}$. Statistically, the price impact of the two latter order types was the same, giving minimal improvement to the fit in (11). However, as we will show below, cancellations are important for capturing periods of scarce liquidity.

The idea of (12) echoes the analysis in Cont et al. [13], who however worked with time-based buckets (of 10 seconds, so generally shorter than ours) and combined all flows into a single predictor, i.e. pre-assigned equal price impact $\alpha_2 = \alpha_1$ or $\beta = 1$ to market and limit orders. Our analysis suggests that β is significantly less than 100%, which seems to be closer to market practice. On the one hand, intuition suggest that executed trades carry more information, i.e. are more relevant for questions of adverse selection or informed trading and hence contribute more to price formation. On the other hand, the nominal volume of limit orders VL is 1–3 times larger than nominal volume of VM; without separate coefficients, ΔP would be unduly affected by limit orders. The above findings are also consistent with the event-by-event analysis done in [16, 22] who found that the impact of market trades (measured by the temporal impulse response

Ticker	$0.25\%ADV$	$1\%ADV$	$2\%ADV$
BBBY	0.39	0.55	0.58
INTC	0.20	0.42	0.55
MSFT	0.41	0.67	0.59
NTAP	0.33	0.64	0.68
ORCL	0.24	0.57	0.68
TEVA	0.49	0.66	0.66

Table 4: Price impact β of touch Limit Orders relative to Market Orders based on (12).

function) for large-tick Nasdaq stocks is 2–4 times larger than that of limit orders.

Remark 4. We also mention the less-known work of [23] who proposed a different way of defining Net Liquidity by using power-weighing, namely considering $\sum_{i:\tau_k<T_i^L\leq\tau_{k+1}}(O_i^L)^\alpha 1_{\{S_i^L=p_1^j(T_i^L)\}}$ in (7), with the power coefficient $\alpha\in[0,1]$. In particular, he advocated looking at square-root-volume ($\alpha=0.5$) or just order counts ($\alpha=0$) when constructing the $NetLiq$ measure. On the latter point, [5] used the *number* of market trades to estimate permanent price impact.

3.2 Limit Flows vs Market Flows

The previous subsection documented that price change is driven by a weighted average of VM and VL, with the link being essentially linear. This motivates the analysis of the relationship between trade imbalance and limit order flows, shown for a representative ticker in Figure 6 below. We separately plot VL^A and VL^B against TI to highlight the inherent *nonlinear* dependence between the two sides of the book. Indeed, when the market is unbalanced (large $|TI|$), the book behaves asymmetrically: the net limit order activity is stable on the passive side, while liquidity provision on the active side declines. We conjecture that this effect is due to market-makers competing to avoid adverse selection, leading to diminished new limit orders and existing order cancellations. Thus, Figure 6 shows that there is a hockey-stick-response curve $VL^j=g^j(TI)$ on each side of the book. For example, looking at VL^A, there is a "base-line" level of liquidity provision that takes place as long as $TI<0.3$ or so (i.e. less than 2/3 of execution are buys). However, when the market is dominated by Buys, VL^A tends to be lower, i.e. the Ask side is replenished less than usual under upside price pressure. For some stocks (BBBY, TEVA see Figure 9 in Appendix B) we see an upside-down V-shape, whereby VL^A also falls for TI close to -1, i.e. the Ask side "hibernates" under intense sell pressure. Mirror effects take place on the Bid side. These findings suggest that LOB resilience (and hence price impact) is essentially driven by *one-sided flows* rather than aggregate two-sided metrics. For example, when net limit order flow at the best bid VL^B is large, the price is unlikely to move lower even if trade imbalance is negative.

Figure 6 also shows that large upside price moves tend to coincide with (i) positive TI; (ii) smaller than usual VL^A and (iii) larger than usual VL^B; the converse holding for downside price moves. Hence when prices move rapidly, both the market and limit order flows are highly unbalanced, lending credence to weak resilience manifesting itself as a "wrong-way" correlation between VM and VL. Given the very large number of heterogeneous players participating in trading activity, it is difficult to draw any conclusions as to the respective causality. Whether HFT market-makers are predicting one-sided market flow or reacting to it, liquidity provision via VL seems to reflect liquidity consumption in VM.

Remark 5. An extreme example of weak resilience was documented in [25] in the context of the market events of July 12, 2012. On that day four large-cap US stocks exhibited an unusual trading

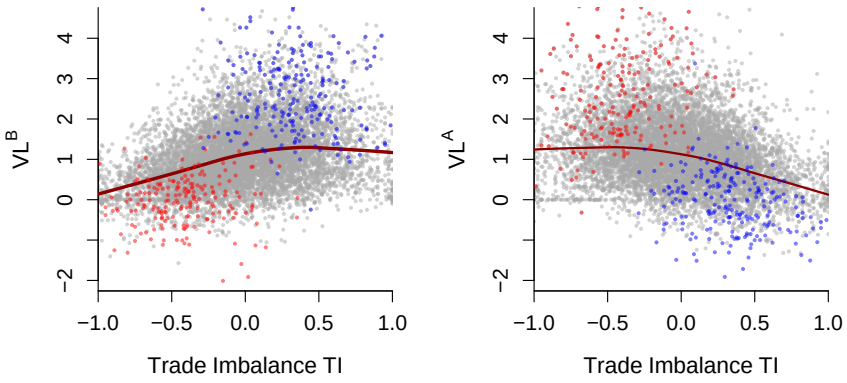

Figure 6: Scatterplot of the net limit order flow at the best bid VL_k^B (left) and best ask VL_k^A (right) against the trade imbalance TI_k across the 8886 buckets of $V = 8000$ for NTAP. The y-axis is normalized via VL^j/V. Red (blue) points indicate volume buckets with price decrease (increase) ΔP_k of more than 0.07. The solid lines indicate the average response $g^j(TI)$ obtained from a spline-based GAM fit to $VL^j = g^j(TI) + \epsilon$.

pattern: heavy buying following by heavy selling in a predictable fashion over 30 minute intervals. The result was a sawtooth pattern in the asset prices, apparently generated by significant price slippage for the aggressive buyers/sellers. Lehalle et al. [25] conclude that this weak LOB resilience was driven by liquidity providers who anticipated one-sided market order flow and were jockeying for queue priority via rapid cancellation of their resting top-level orders after each execution.

4 Liquidity Predictors

The results in the previous section strongly indicate that including concurrent limit order flow VL_k along with trade imbalance TI_k dramatically improves the goodness-of-fit for predicting ΔP_k. Nevertheless, for purposes of optimal execution the role of trade imbalance remains paramount. Indeed, the execution algorithm displaces some of the market orders, so that the order scheduler in effect partially controls the trade imbalance TI_{k+1} in the upcoming bucket. Hence, forecasting the conditional price change $\Delta P_{k+1}|TI_{k+1}$ is crucial in scheduling the next slice.

Unlike the time slicing approach, each bucket in our setup contains an identical amount of traded volume. Thus the difference between observed price change ΔP for two buckets with similar trade imbalance TI must be due to differences in LOB resilience and/or shape. Therefore, the findings in the previous section lead to the natural decomposition of price formation into two components: (1) *price trend*, which is primarily about the supply and demand for the asset captured in trade imbalance TI_k, and (2) *liquidity*, which is seen in the deviation of the observed ΔP_k from the best fit curve between the two variables, and has been already shown to strongly depend on VL_k^j. In this section we seek to more systematically identify the key variables important for explaining LOB liquidity on the meso-scale. This is done by constructing a series of regression models taking in a variety of potential covariates, and quantifying their significance and effect. The following Section 4.1 then concentrates on predicting the periods of scarce liquidity that lead to out-sized price impact.

For the above purpose, we utilized four classes of statistical models: (i) classical linear models (LM) using stepwise selection criteria; (ii) LASSO models fit via cross-validation; (iii) MARS

models; (iv) random forests (RF). Generically, we model

$$\Delta P = \sum_r \phi_r(\mathbf{X}) + \epsilon, \tag{13}$$

where $\mathbf{X}$ is the vector of all the covariates and the basis functions ϕ_r are adaptively picked by the model. We focus our attention on the primary effect of various predictors and therefore take ϕ_r to be low-dimensional in $\mathbf{X}$: they are linear for LM and LASSO, of the hockey-stick form $(X^{(r)} - K_r)_+$ for MARS, and indicator functions of hypercubes for RF. Thus, except for RF no interactions among covariates are a priori considered.

The three model frameworks beyond LM offer complementary perspectives on variable importance and ultimately model selection. LASSO uses L^1-regularization to drop less relevant predictors giving a crisp quantification of variable importance and easily interpretable coefficients for covariate impact. Multivariate adaptive regression splines (MARS) aims to identify nonlinear relationships through using piecewise linear hinge basis functions. It can also search for statistically significant predictors from low-degree interactions among covariates. Finally, RF constructs a black-box nonlinear model by randomly considering subsets of predictors to obtain a piecewise constant relationship. Random Forests yield only indirect measures of variable importance or contribution via model averaging, but makes minimal parametric assumptions about the functional form of covariate impact. Note that our goal is *not* to maximize predictive performance, but to obtain interpretable qualitative analysis. Hence, we do not explore fully non-parametric "machine learning" tools.

To build a statistical description for ΔP we utilized both static measures based on the LOB state at bucket times τ_k, as well as dynamic quantities measured over time. The static predictors tested in our regression models included the measures of LOB depth and shape described in Section 2, specifically price impact PI_N, impact slope S, cumulative depth D_i for $i = 1, 2$ and book imbalance BI. For completeness, we also included time-of-day τ_k. All static quantities were measured at the start of the volume slice i.e. the price change $P_k - P_{k-1}$ from the activity during $[\tau_{k-1}, \tau_k]$ was regressed against static variables measured at time τ_{k-1}. Dynamic quantities included the contemporaneous net limit flows VL^A_k, VL^B_k, as well as lagged quantities: trailing limit flow $VL^j_{k-\ell}$, lagged price change $\Delta P_{k-\ell}$, and lagged trade imbalance $TI_{k-\ell}$ over the previous ℓ volume buckets, where $\ell \in 1, 5, 10, 20$. Also included was the exponentially weighted moving average of recent trade imbalance, first introduced in [4],

$$TIMA^{(\beta)}_{i+1} := e^{-\beta|O^M_i|} TIMA^{(\beta)}_i + (1 - e^{-\beta|O^M_i|}) \operatorname{sign}(O^M_i), \tag{14}$$

where i indexes arrival of market orders of size $|O^M_i|$ and β dictates the exponential weights; we set $\beta = 0.5/V$.

Important Covariates. Our results indicate that the trade imbalance TI and the limit order flows VL^j are by far the most statistically significant predictors for ΔP. Further predictors that were consistently picked by the models include the price impact PI^j, the depth at the top-2 levels D_2, and the proportion of limit order cancellations $PC^j := VL^{j,-}/VL^j$. For the latter, we observe that when cancellations at the bid dominate additions, $PC^B > 0.5$, the price can move lower on very little volume. Weaker predictors (picked by some models but not all) are Book Imbalance BI and Impact Slope S. There are additional covariates that were statistically significant for one or two of the tickers, but these yield negligible improvement to overall model fit. See Table 6 in Appendix B for a summary of variable importance.

The RF models also assigned strong statistical significance to the interactions $TI \times PI$, $VL^j \times PI^j$ (one-sided interaction between limit flow and book impact) and $TI \times S$. Since both PI

and S are in units of ticks/share, the latter two terms are denominated in ticks, and provide an impact-normalized measure of market trend. Indeed, the static LOB metrics capture the shape of the book and hence *modulate* the effect of the major predictors in the flows. The relevance of the above cross-terms was confirmed by a LASSO model where we manually added all possible interactions between TI, VL^A, VL^B and the impact covariates PI and S. As an additional check, no further interactions (considered by building a degree-2 MARS model) were deemed significant. Overall, using the R^2 as a simple goodness-of-fit metric, we find that it starts out at $R^2 \simeq 30\%$ when regressing ΔP against just TI; rises to about $R^2 \simeq 70\%$ once VL^A, VL^B are included; and reaches about $R^2 \simeq 80\%$ after adding in the mentioned LOB shape metrics. Another 1-2% gain in R^2 is possible by including all the covariates listed, including their lagged versions.

Of note, top level depth D_1 was statistically insignificant in the regression tests. Given the amount of activity in each volume bucket it is not surprising that "shallow" LOB measures capturing only the top of the LOB are not all that useful. Instead, the models identify (partly due to the underlying collinearities) a mix of the deeper LOB shape metrics, including the top-2 levels volume D_2, the price impact PI_N (which looks at 3-5 levels deep, $N \simeq Ave(D_4)$), and the impact slope S (a cost-averaged depth over $\simeq 4$ levels). We note that for more liquid tickers (such as MSFT), D_2 is more relevant, while for less liquid ones the models emphasize PI. Another interesting find was the (moderate) significance of trailing moving average of executions $TIMA$ from (14), with a consistently negative relationship to ΔP_k. The inclusion of $TIMA$ appears to be capturing the tendency of stock prices to, more often than not, retrace recent movement. More precisely, when TI leans with the prevailing trend ($TIMA$) the impact on price is less.

Returning to a more qualitative description, Figure 8 (Appendix C) shows paired scatterplots across 9229 buckets of ORCL for limit flows VL^j, proportion of cancellations PC^j and impact slopes S^j. Some important observations are:

- Strong negative correlation $\rho \simeq -0.4$ between VL^A and VL^B, meaning that Ask/Bid limit order activity is complementary: at any given moment most of the limit orders go to one side;
- Positive correlation ($\rho \simeq 0.45$) between order flow VL and proportion of cancellations PC on the *other* side: as more orders are added, say, on the Ask side, limit orders are also cancelled on the Bid side;
- Negative correlation between depth D and order flow VL on the *opposite side*, so that the less-active side of the book also tends to be more shallow;
- Positive correlation between book depth on the two sides, D^A and D^B, meaning that the book as a whole fluctuates between being deep and shallow;
- Weak positive correlation ($\rho \simeq 0.2$) between cancellation proportions PC^A, PC^B on the two sides, dovetailing with the previous point.

These complex cross-relationships imply that disentangling the marginal effect of a given predictor is difficult, and a simple additive model, such as MARS or LASSO, should be interpreted with care. For instance, the ask-side depth metric D_2^A is intrinsically correlated with D_2^B, VL^B, S^A, PC^A and so on. Therefore, the counterfactual of changing ask-side depth while keeping all other covariates fixed, captured by a regression coefficient (or partial dependence plot) of D_2^A, is not too meaningful.

4.1 Scarce Liquidity

Intuitively, scarce liquidity describes the situation where large price moves occur under low volume, indicating weak resilience. Within our framework, the latter corresponds to an "abnormally" large ΔP_k relative to the contemporaneous trade imbalance TI. Pinpointing the confluence of factors that coincide with scarce liquidity is of practical importance due to the potentially large ramifications of executing during a period of low resilience in the LOB.

To capture the idea of a "disproportionately large" price move relative to executed volume we remove the effect of trade imbalance to focus on the resulting residuals $\widehat{\epsilon}_k = \Delta P_k - g(TI_k)$ from (10). Namely, we define a scarce liquidity indicator variable $SL_k^j \in \{0, 1\}$ in a binary fashion in terms of $\widehat{\epsilon}_k$, so that SL^j occurs on at most one side of the LOB per volume bucket:

$$SL_k^A := I\{\widehat{\epsilon}_k \geq \bar{M}\}, \qquad SL_k^B := I\{\widehat{\epsilon}_k \leq -\bar{M}\}, \tag{15}$$

where we empirically choose the scarcity threshold $\bar{M}$ in terms of the standard deviation of the residuals $\hat{\epsilon}$, specifically $\bar{M} = 1.5StDev(\widehat{\epsilon}_{1:K})$. Thus, scarce liquidity occurs on each side of the LOB in approximately 6% of volume buckets. Our data-driven definition of liquidity is similar to Amihud [3]. The pairs plot in Figure 8 in the Appendix color-codes buckets with scarce liquidity, highlighting the relationship of key predictors to SL^j.

To quantitatively analyze the occurrence of scarce liquidity we ran a logistic regression model for SL^A and SL^B with the covariates described in the previous section. Compared to the regression models for ΔP of the previous section, logistic regression is geared towards capturing one-sided LOB effects and also focuses attention on the buckets with more extreme price moves. It can therefore better identify the LOB regimes that cause substantial price moves, compared to (13) that concentrates on fitting the "typical" price formation factors. Recall that logistic regression assumes that $SL_k^j \sim Bernoulli(\pi_k^j)$, where the log-odds of π_k^j are specified via

$$\text{logit}\, \pi^j = \sum_r \phi_r^{SL,j}(\mathbf{X}). \tag{16}$$

This analysis reveals that occurrences of scarce liquidity are driven by a high asymmetry in the limit flows, as well as reduced LOB depth and/or high cancellation rates. For example, the major predictors of SL^A are VL^B (+ve correlation), TI (-), VL^A (-) and $D_2^{A,B}$ (both -). Several of these influences need to come together to explain SL^A. Financially, we observe that scarce liquidity tends to occur on a passive and shallow side of the book. For instance, if the Bid side is passive (high VL^A and low VL^B), and moreover is less deep than normal (low D_1^B), then prices are likely to fall through on the downside, triggering $SL^B = 1$. The logistic regression models also show that the proportion of limit order cancellations has a high predictive power, so that large price moves are associated with a "LOB fade". As in the previous section, the limit order flows were the most statistically significant predictors for SL^j via (16). Unlike the previous section, price impact PI had a heightened significance. Specifically, same-side price impact (PI_N^j when testing π^j) was of approximately equal importance as VL^j for all six tickers. The physical timing τ_k, was also significant, as was bucket duration $\Delta\tau_k$. The fitted models imply that scarce liquidity is more likely in the morning, or when trading is slow (large $\Delta\tau_k$).

Table 7 in Appendix B reports goodness-of-fit of the above logistic models using a hold-out test set. The Table shows the relative frequency of SL^j against the predicted log-odds $\widehat{\pi}^j$. We observe that for vast majority of cases the log-odds are low and in those buckets there is indeed almost no cases of abnormally large price moves. On the other hand, when the model predicts scarce liquidity, it is correct in about 70–80% of the time. We note that due to the low frequency

of SL, the model predicts less than half of all scarce liquidity occurrences, i.e. there are more false negatives than true positives.

Table 6 in Appendix B lists the significant predictors for the three regressions we ran (for ΔP, SL^A, and SL^B). In the Table we used the `importance` function in the `randomForest` package [27] to roughly classify given predictors into Most Important (**), Somewhat Important (*) and Unimportant (-). Since the RF model is data-driven and able to capture nonlinear effects the resulting variable importance is more robust compared to other model classes we tried, such as Lasso or MARS which only capture low-degree interactions and have trouble with collinearity. We also manually checked the predictors selected by LASSO and MARS vis-a-vis those of RF to confirm covariate importance.

Table 6 matches the intuition that same-side covariates (especially the impact-dollar metrics PI^j and S^j) ought to be more significant for explaining SL^j. We also note that the logistic models are more complex (assigning predictive power to more covariates), and moreover include additional "two-sided" predictors, such as bucket duration or price volatility, that would not contribute to ΔP, but make sense for forecasting large price moves. The observed asymmetric effects imply the need to build separate models for different aspects of the limit-flow/price relationship. The listed significant predictors are remarkably consistent across assets. Moreover, they are complementary/symmmetric on the Ask and Bid side (the respective regressions were run completely independently), indicating that these are *generic stylized facts* for liquid, large-tick LOBs.

5 Time Series Perspective

Previous sections carried out a statistical investigation of LOB behavior using the regression paradigm that treats observations as i.i.d. samples from the "true population". Of course, in practice LOB data is sequentially ordered, so time series methods are also appropriate. However, the observed complexity (such as multiple nonlinear relationships) makes constructing multivariate time-series models a major challenge and beyond the scope of this paper. For our purposes we therefore simply present some stylized observations.

A basic test is for the presence of auto-correlation which would imply that there are some persistent factors in the observations, and hence serial dependence among buckets that are close in time. We first consider the net order flow variables: TI, VL^A, VL^B. We find that there is minimal autocorrelation at the shorter scales (0.25% ADV buckets), but some persistence is observed on the longer scales and for less-liquid assets. Positive autocorrelation suggests that periods of strong liquidity provision are followed by more of the same, and conversely, low VL tends to persist for multiple buckets. For the deeper LOB's, we believe the persistence is still there but a better tool than the ACF diagnostic is needed. Remarkably, the persistence carries through for long lags, indicating presence of long-term memory on the order of hours. Figure 7 illustrates the non-zero autocorrelation in VL^j for NTAP and TEVA across 2% ADV buckets. These results echo previous literature (e.g. [35]) that documents long-term persistence in order flows, possibly even on the scale of days. Overall, we note that the computed auto-correlation levels are not too strong (relative to say the very high R^2 reported previously), indicating that standard ARIMA models might not fit well.

Turning attention to scarce liquidity, as remarked there is essentially no auto-correlation in the residuals of (10) and hence minimal ACF levels in the one-sided SL^j indicators. In other words, the ACF test does not detect any obvious clustering in one-sided scarce liquidity. At the same time, we do find cross-correlation, so that an abnormally large positive price move is likely to be followed by an abnormally large negative ΔP_k. This hints at the "whiplash" property of the

LOB, where reversions in price change or in TI seem to trigger reduced liquidity provision. The right panel of Figure 7 shows the ACF plot for the combined indicator $SL = SL^A + SL^B$ (volume buckets with scarce liquidity in either direction), which is positively auto-correlated, i.e. exhibits some temporal clustering.

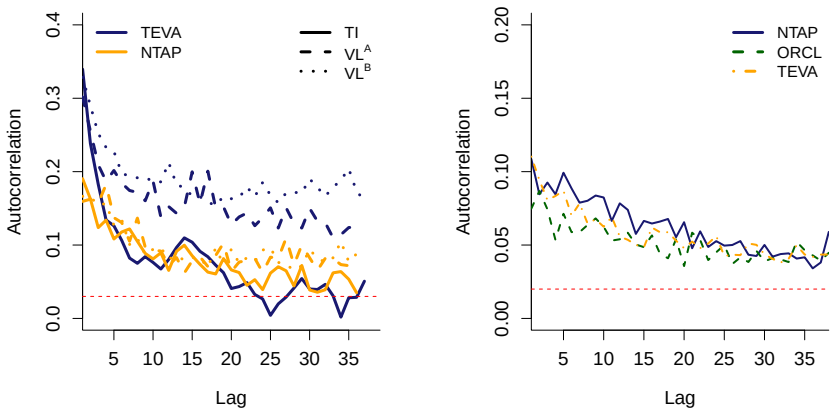

Figure 7: Marginal autocorrelation functions (ACF) for order flows and scarce liquidity. *Left:* the marginal ACF for the trade imbalance TI, as well as one-sided net limit flows at the touch VL^A and VL^B. Data is based on 2% ADV bucketing for NTAP and TEVA. Right: marginal ACF for the scarce liquidity indicator $SL = SL^A + SL^B$. In both panels dashed lines indicate the significance threshold for the presence of autocorrelation.

5.1 Co-Movement of Ask/Bid Order Flows

Recall Figure 6 which presented an aggregated relationship between VM and VL. To capture the respective co-movement of flows, we consider the correlation ρ^j between one-sided limit flows VL^j and VM^j, $j = A, B$. Lower levels of ρ^j are indicative of increased cancellations and hence less liquidity, while positive values are characteristic of strong resilience. *Negative* correlation suggests that liquidity gets removed/added in tandem with market orders, which would be characteristic of book "fading", cf. the event-by-event plot in Figure 2. Rapidly declining ρ^j could then be an indication of the cumulative effect of persistent and aggressive buying/selling that leads market makers to adjust their reference prices to account for added inventory risk.

A major challenge is to operationalize ρ^j given the asynchronous nature of the entered orders. Using the aggregated bucketed quantities VL_k, VM_k is not convenient given the relatively large buckets we consider—to compute statistical correlation one needs time-series of at least a few dozen terms, while the episodes observed in Figure 2 are much more fleeting. One attempt we made was to define the correlation on either side of the LOB as

$$\rho^j_t := \mathrm{corr}(VM^j_{[t-s,t]}, VL^j_{[t-s,t]}), \tag{17}$$

computing empirical correlation over a sliding window of s physical time using smaller time-based buckets. Specifically we tried $s = 1.5$ hours with time buckets of 30 seconds. However the resulting (ρ^j_t) were not very stable in time and the statistical significance of the resulting predictor was too low to be included in any of the models of Section 4. Another challenge is normalizing ρ across assets. For tickers with deeper queues, such as MSFT, the short-term behavior of limit flows is less tied to price formation, so that VL^j and VM^j are less coupled, leading to lower respective correlations $|\rho^j|$.

Remark 6. A related mechanism to quantify the co-movement of market and limit orders for the purposes of explaining book liquidity are so-called *toxicity* indicators. Toxicity is linked to adverse selection: in a toxic environment market-makers withdraw due to orders coming from "informed" traders who possess better knowledge of the future asset price. In that direction, we may mention the VPIN metric [15] which can be operationalized as the average of the absolute trade imbalance across the ℓ most recent volume buckets,

$$|\overline{TI}|_{k-\ell:k} := \frac{1}{\ell}\sum_{i=1}^{\ell}|TI_{k-i}|, \tag{18}$$

Large values of $|\overline{TI}|_{k-\ell:k}$, are supposed to lead to less liquidity and more volatility. Another toxicity measure described in [8] is constructed from the event-based correlation between market flow and price change based on the previous 200 market order arrivals:

$$\rho^{Tox}(\tau_k) := \mathrm{corr}(\Delta P_{i-200:i}, O^M_{i-200:i}). \tag{19}$$

We also tried to include these in our regressions, but preliminary results suggested that they had weak predictive power (this is consistent with the feature that all lagged variables were poor predictors relative to contemporaneous covariates). It remains an open question how to construct a relevant volume-bucketed toxicity measure.

6 Conclusion

In this article we have investigated the meso-scopic phenomenology of price formation and liquidity provision/consumption. The documented stylized facts are crucial for algorithmic scheduling of child trades, where the LOB state is used as a dynamic input rather than a static parameter. Our central take-away is that at the minute-scale the distinct order flows, both of market and limit type, are the primary predictors of price moves. We have also observed that top-level depth and the associated book imbalance are insufficient for forecasting price impact; rather deeper book characteristics, such as PI and S, are better suited.

Our results pave the way for further research directions. First, we stress that it is imperative to develop models for limit flows in parallel with the market flow. This is a challenging task as one must consider the book asymmetry (interaction between market and limit orders on the same side of the book), cross-effects (interactions between market buys and sells, and between limit orders on the two sides, e.g. via book imbalance) and interaction between limit Additions and Cancellations. Second, our approach offers a way to disentangle price evolution into price trend (impact of market executions), book resilience(impact of limit flows) and liquidity (residual effects that drive price fluctuations beyond their expected level). This gives a starting point for quantifying LOB dynamics, for example to obtain a dynamic definition of liquidity regimes. In analogue to the developed models for individual order arrivals, we need probabilistic description for the evolution of, say, (PI_k) or SL_k. Third, the unique meso-scopic phenomena raise new questions on connecting the event-based market microstructure, the time-based macro-diffusion and the volume-based bucketing proposed herein in a single multi-scale framework.

First-generation models for execution, liquidation and market-making have been prescriptive in nature, postulating some features of the limit order books and optimizing participant behavior under those assumptions. The phenomenological investigation herein takes the opposite view, attempting to crystalize the data-driven features. For the practical application, the next step would be to develop *predictive* models that can statistically forecast book characteristics going forward,

and therefore inform execution algorithms. Theoretically, we need models that move away from working with static book snapshots (e.g. an abstract "depth" parameter) and towards using more robust definitions of impact. More sophisticated frameworks for capturing the interaction between market makers and observed flows, the so-called "toxicity" effects, are also warranted.

By design, our analysis has been limited to the most liquid, large-tick equities. While these cover many of the most active tickers, including many blue-chip companies such as Dow Jones components, more work is needed on other ticker types. This would include liquid small-tick assets, such as AAPL or GOOG, which might be very active, but where the spread has non-trivial dynamics and book depth has to be understood differently. One challenge for such tickers is that a nontrivial portion of the limit orders are aggressive, i.e. placed inside the spread, which requires to adjust the meaning of VL. Another issue is that when share prices are very high, the discrete nature of order volumes becomes prominent. A different direction for future work is analysis of large-tick but illiquid assets, where the event-by-event scale might still be relevant even on the minutes-scale.

References

[1] A. Alfonsi and A. Schied, *Optimal trade execution and absence of price manipulations in limit order book models*, SIAM Journal on Financial Mathematics, 1 (2010), pp. 490–522.

[2] A. Alfonsi, A. Schied, and A. Slynko, *Order book resilience, price manipulation, and the positive portfolio problem*, SIAM Journal on Financial Mathematics, 3 (2012), pp. 511–533.

[3] Y. Amihud, *Illiquidity and stock returns: cross-section and time-series effects*, Journal of financial markets, 5 (2002), pp. 31–56.

[4] K. Bechler and M. Ludkovski, *Optimal execution with dynamic order flow imbalance*, SIAM Journal on Financial Mathematics, 6 (2015), pp. 1123–1151.

[5] M. Benzaquen, I. Mastromatteo, Z. Eisler, and J.-P. Bouchaud, *Dissecting cross-impact on stock markets: An empirical analysis*, tech. report, Available at SSRN 2836472, 2016.

[6] J.-P. Bouchaud, J. D. Farmer, and F. Lillo, *How markets slowly digest changes in supply and demand*, in Handbook of Financial Markets: Dynamics and Evolution: Dynamics and Evolution, T. Hens IV and K. R. Schenk-Hoppé, eds., Elsevier, 2009, pp. 57–160.

[7] J.-P. Bouchaud, Y. Gefen, M. Potters, and M. Wyart, *Fluctuations and response in financial markets: the subtle nature of random price changes*, Quantitative finance, 4 (2004), pp. 176–190.

[8] R. Carmona and K. Webster, *The self-financing equation in high frequency markets*, tech. report, 2013. arXiv preprint arXiv:1312.2302.

[9] Á. Cartea, R. F. Donnelly, and S. Jaimungal, *Enhancing trading strategies with order book signals*, tech. report, Available at SSRN 2668277, 2015.

[10] Á. Cartea and S. Jaimungal, *Incorporating order-flow into optimal execution*, Mathematics and Financial Economics, 10 (2016), pp. 339–364.

[11] Á. Cartea, S. Jaimungal, and J. Penalva, *Algorithmic and high-frequency trading*, Cambridge University Press, 2015.

[12] R. Cont and A. De Larrard, *Price dynamics in a Markovian limit order market*, SIAM Journal on Financial Mathematics, 4 (2013), pp. 1–25.

[13] R. Cont, A. Kukanov, and S. Stoikov, *The price impact of order book events*, Journal of Financial Econometrics, 12 (2014), pp. 47–88.

[14] R. F. Donnelly, *Ambiguity Aversion in Algorithmic and High Frequency Trading*, PhD thesis, 2014. Available at SSRN 2527808.

[15] D. Easley, M. M. L. de Prado, and M. O'Hara, *Flow toxicity and liquidity in a high-frequency world*, Review of Financial Studies, 25 (2012), pp. 1457–1493.

[16] Z. Eisler, J.-P. Bouchaud, and J. Kockelkoren, *The price impact of order book events: market orders, limit orders and cancellations*, Quantitative Finance, 12 (2012), pp. 1395–1419.

[17] J. D. Farmer, A. Gerig, F. Lillo, and H. Waelbroeck, *How efficiency shapes market impact*, Quantitative Finance, 13 (2013), pp. 1743–1758.

[18] J. Gatheral, A. Schied, and A. Slynko, *Transient linear price impact and Fredholm integral equations*, Mathematical Finance, 22 (2012), pp. 445–474.

[19] M. D. Gould, M. A. Porter, S. Williams, M. McDonald, D. J. Fenn, and S. D. Howison, *Limit order books*, Quantitative Finance, 13 (2013), pp. 1709–1742.

[20] W. K. Härdle, N. Hautsch, and A. Mihoci, *Modelling and forecasting liquidity supply using semiparametric factor dynamics*, Journal of Empirical Finance, 19 (2012), pp. 610–625.

[21] J. Hasbrouck and D. J. Seppi, *Common factors in prices, order flows, and liquidity*, Journal of financial Economics, 59 (2001), pp. 383–411.

[22] N. Hautsch and R. Huang, *The market impact of a limit order*, Journal of Economic Dynamics and Control, 36 (2012), pp. 501–522.

[23] C. Hopman, *Do supply and demand drive stock prices?*, Quantitative Finance, 7 (2007), pp. 37–53.

[24] W. Huang, C.-A. Lehalle, and M. Rosenbaum, *Simulating and analyzing order book data: The queue-reactive model*, Journal of the American Statistical Association, 110 (2015), pp. 107–122.

[25] C.-A. Lehalle and M. Lasnier, *What does the saw-tooth pattern on US markets on 19 july 2012 tell us about the price formation process*, tech. report, CA Chevreux Quantitative Research, 2012.

[26] C.-A. Lehalle and O. Mounjid, *Limit order strategic placement with adverse selection risk and the role of latency*, tech. report, arXiv preprint arXiv:1610.00261, 2016.

[27] A. Liaw and M. Wiener, *Classification and regression by randomforest*, R News, 2 (2002), pp. 18–22.

[28] A. Lipton, U. Pesavento, and M. G. Sotiropoulos, *Trade arrival dynamics and quote imbalance in a limit order book*, tech. report, arXiv preprint arXiv:1312.0514, 2013.

[29] C. C. Moallemi, M. Saglam, and M. Sotiropoulos, *Short-term trading skill: An analysis of investor heterogeneity and execution quality*, tech. report, Available at SSRN 2463952, 2016.

[30] I. Muni Toke, *The order book as a queueing system: average depth and influence of the size of limit orders*, Quantitative Finance, 15 (2015), pp. 795–808.

[31] A. A. Obizhaeva and J. Wang, *Optimal trading strategy and supply/demand dynamics*, Journal of Financial Markets, 16 (2013), pp. 1–32.

[32] E. E. Omar, F. Masaaki, and R. Mathieu, *The microstructural foundations of leverage effect and rough volatility*, tech. report, arXiv preprint arXiv:1609.05177, 2016.

[33] S. Stoikov and R. Waeber, *Optimal asset liquidation using limit order book information*, tech. report, Available at SSRN 2113827, 2012.

[34] B. Toth, Z. Eisler, F. Lillo, J. Kockelkoren, J.-P. Bouchaud, and J. D. Farmer, *How does the market react to your order flow?*, Quantitative Finance, 12 (2012), pp. 1015–1024.

[35] B. Toth, I. Palit, F. Lillo, and J. D. Farmer, *Why is order flow so persistent?*, Journal of Economic Dynamics and Control, 51 (2015), pp. 218–239.

[36] S. N. Wood, *Fast stable restricted maximum likelihood and marginal likelihood estimation of semiparametric generalized linear models*, Journal of the Royal Statistical Society (B), 73 (2011), pp. 3–36.

Appendices

A Volume Bucketing Parameters

	0.25%ADV			1%ADV			2%ADV		
Ticker	V	med($\Delta\tau$)	Trades	V	med($\Delta\tau$)	Trades	V	med($\Delta\tau$)	Trades
BBBY	$1.5K$	32.9	7.8	$6K$	161	31.2	$12K$	334	62.5
INTC	$9K$	21.4	5.2	$36K$	149	20.7	$72K$	324	41.4
MSFT	$25K$	28.0	14.6	$100K$	159	58.2	$200K$	341	116.5
NTAP	$2K$	36.6	7.9	$8K$	188	31.7	$16K$	393	63.5
ORCL	$5K$	31.6	6.0	$20K$	170	23.9	$40K$	357	47.8
TEVA	$3K$	31.1	10.2	$12K$	158	40.9	$24K$	337	81.9

Table 5: Summary statistics for the volume-based bucketing. We list the size V of each bucket (in 000's executed shares), median duration $med(\tau_{k+1} - \tau_k)$ in seconds, and average number of market trades per slice.

B Predictor Importance

	Price Change ΔP						Ask Scarce Liquidity SL^A						Bid Scarce Liquidity SL^B					
Predictor	BBBY	INTC	MSFT	ORCL	NTAP	TEVA	BBBY	INTC	MSFT	ORCL	NTAP	TEVA	BBBY	INTC	MSFT	ORCL	NTAP	TEVA
TI_k	**	**	**	**	**	**	**	*	**	**	**	*	**	**	**	**	**	**
$TI_{k-3:k-1}$	–	–	–	–	–	–	*	*	*	*	*	*	–	*	*	*	–	*
TI_{k-1}	–	–	–	–	–	–	*	*	*	*	–	*	–	*	*	*	*	–
VL_k^a	*	**	*	**	**	**	**	**	**	**	**	**	**	**	**	**	**	**
VL_k^b	*	**	*	**	**	**	**	**	**	**	**	**	**	**	**	**	**	**
VL_{k-1}^a	–	–	–	–	–	–	*	–	–	–	–	–	*	*	*	–	*	*
VL_{k-1}^b	–	–	–	–	–	–	*	*	*	*	–	*	*	–	–	*	–	–
PC_k^a	*	*	*	*	*	*	**	**	**	**	**	**	*	*	*	**	**	*
PC_k^b	*	*	*	*	*	*	*	*	*	**	**	*	**	**	**	**	**	**
PI_k^a	*	*	*	*	*	*	**	**	**	**	**	**	*	**	*	*	**	*
PI_k^b	*	*	*	*	*	*	*	**	**	**	**	*	**	**	**	**	**	**
S_k^a	*	*	*	*	*	–	**	**	**	**	**	**	**	**	*	*	**	*
S_k^b	*	*	*	*	*	–	*	**	**	**	**	*	**	**	**	**	**	**
D_1^a	–	–	–	–	–	–	*	*	*	*	*	*	–	*	*	*	*	*
D_1^b	–	–	–	–	*	–	*	**	*	*	**	*	*	*	*	*	*	*
D_2^a	*	*	*	*	*	*	**	**	**	**	**	**	**	**	**	**	**	*
D_2^b	*	*	*	*	*	*	**	**	**	**	**	**	**	**	**	**	**	**
$BI_{k-3:k}$	–	–	–	–	*	–	*	*	*	*	**	*	*	*	*	*	**	*
τ_k	*	–	–	*	*	–	**	**	*	**	**	**	*	*	*	**	**	*
$\tau_k - \tau_{k-1}$	–	–	–	–	–	–	*	*	*	**	**	**	*	*	*	*	**	**
$P_k - P_{k-20}$	–	–	–	–	–	–	*	–	*	–	*	–	–	–	–	–	*	–

Table 6: Significant predictors for price change ΔP and scarce liquidity SL^j using a Random Forest model. We use a RF with 500 trees, and the mean decrease in accuracy variable importance metric reported in `randomForest::importance` function. Highly significant (**) predictors are those with relative importance above 0.4; significant (*) predictors have relative importance in $[0.15, 0.4]$; the rest are labeled as not significant (-). Results are based on 1% ADV bucketing. Predictors with subscripts other than k are (averaged) lagged values over the indicated previous buckets.

PLOT OMITTED FROM ARXIV DUE TO SIZE

Figure 8: Pairs plot for main predictors of scarce liquidity. Blue points indicate buckets with Ask-side scarce liquidity $SL^A = 1$, and green are buckets with $SL^B = 1$. Plot for NTAP and $V = 8000$ (1% ADV). Lower-triangle entries indicate the correlation between the predictor pairs over the entire sample period.

		SL^A			SL^B		
		Low	Med	High	Low	Med	High
BBBY		AUC = 0.938			AUC = 0.942		
	$NoSL$	84.73	10.71	0.40	84.28	10.59	0.37
	SL	0.77	2.76	0.63	0.68	3.02	1.05
INTC		AUC = 0.928			AUC = 0.925		
	$NoSL$	86.67	7.98	0.24	86.55	8.61	0.55
	SL	1.34	3.04	0.73	0.97	2.47	0.85
MSFT		AUC = 0.946			AUC = 0.963		
	$NoSL$	86.62	8.33	0.24	87.96	7.47	0.18
	SL	0.77	2.36	1.68	0.49	2.70	1.19
NTAP		AUC = 0.958			AUC = 0.961		
	$NoSL$	86.11	8.74	0.34	86.62	8.12	0.31
	SL	0.72	2.97	1.13	0.51	3.11	1.33
ORCL		AUC = 0.959			AUC = 0.967		
	$NoSL$	86.51	8.54	0.33	87.62	7.72	0.26
	SL	0.56	2.69	1.38	0.43	3.02	0.95
TEVA		AUC = 0.934			AUC = 0.945		
	$NoSL$	85.95	9.45	0.24	86.67	8.37	0.27
	SL	0.93	2.67	0.75	0.78	2.85	1.05

Table 7: Goodness-of-fit for predicting Scarce Liquidity. We list observed percentages of scarce liquidity SL (rows) grouped by predicted model response (columns). Specifically, let $\widehat{\pi}_k^j$ be the predicted probability of SL_k^j as obtained from individual tree votes in a RF logistic regression model (16) with 500 trees. Then $Low = \{k : \widehat{\pi}_k \leq 0.1\}$, $Med = \{\widehat{\pi} \in (0.1, 0.5)\}$, $High = \{\widehat{\pi} \geq 0.5\}$. AUC is the area under the receiver operator curve, computed using a hold-out test set: 2/3 of data was randomly selected for training and the rest for above testing. As a result, the actual frequency of SL^j in the test sets varies across assets in the range 4%–5%. All results are based on 1% ADV buckets.

C Relationship between Market and Limit Order Flows

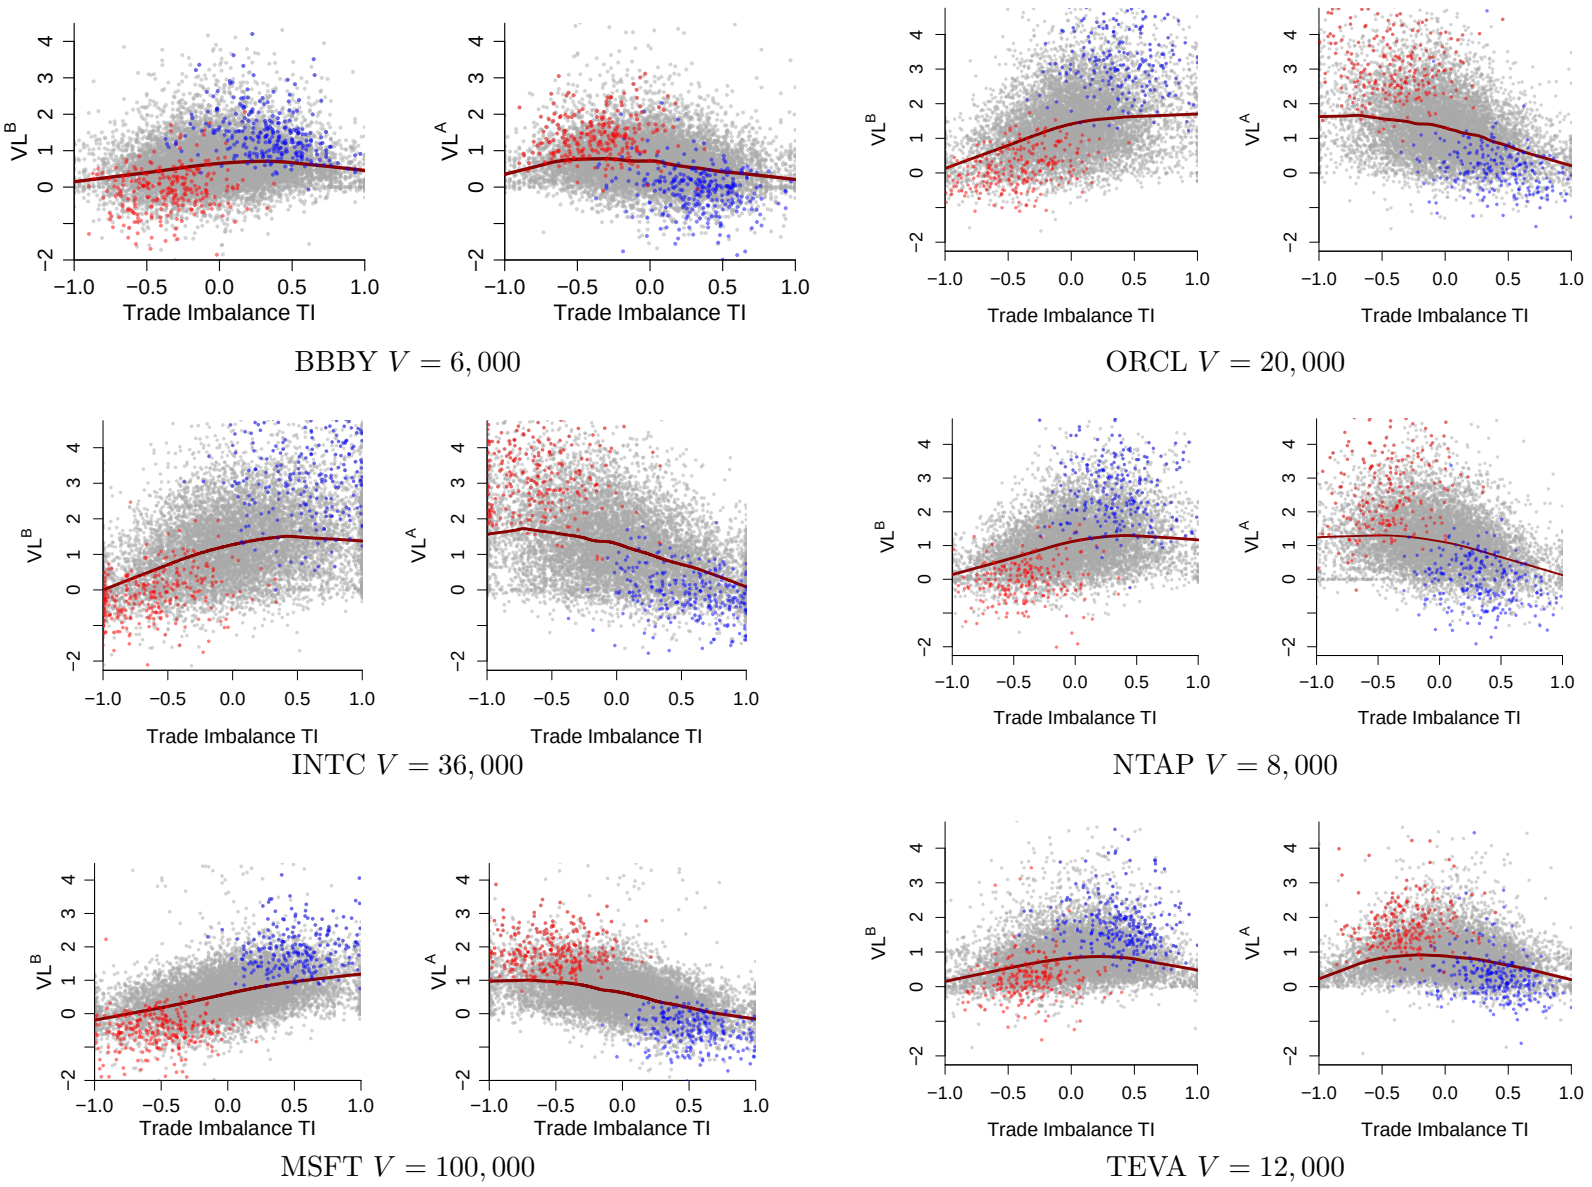

Figure 9: Net limit order flow at the best bid (left panels) and best ask (right panels) plotted against the market order imbalance TI. The y-axis is normalized via VL^j/V. Blue (red) points indicate buckets with large price moves: ΔP_k of more (less) than $\pm 1.645 \cdot StDev(\Delta P)$. The solid line shows the GAM fit for $g^j(TI) = \mathbb{E}[VL^j|TI]$, cf. Figure 6.