The Conduits of Price Discovery: A Machine Learning Approach*

Amy Kwan† **Richard Philip**‡ **Andriy Shkilko**§

December 30, 2021

Abstract. When examining information flow into prices, empirical studies tend to focus on order submissions. Meanwhile, theory suggests that market conditions should play important roles in augmenting and regulating information flows. Using a machine learning technique known as reinforcement learning, we show that price discovery is notably affected by such theory-based conditions as the state of the limit order book, price history, bid-ask spread, and order arrival frequency. The state of the book and price history stand out as conditions whose importance rivals that of order submissions. The conditions have a substantial influence on trading by sophisticated market participants.

Key words: price discovery, order submission strategies, machine learning

JEL: G14; G15

*We thank Vincent Bogousslavsky, Jonathan Brogaard, Michael Brolley, James Brugler, David Cimon, Rama Cont, Thierry Foucault, Michael Goldstein, Roberto Gomez-Cram, Marco Grotteria, Björn Hagströmer, Jangwook Lee, Giulia Livieri, Ron Masulis, Dmitriy Muravyev, Andreas Park, Brian Roseman, Gideon Saar, Elvira Sojli, Konstantin Sokolov, Wing-Wah Tham, Baozhong Yang, Bart Yueshen, and seminar/session participants at the Central Bank Conference on the Microstructure of Financial Markets, Conference on Markets and Economies with Information Frictions, European Finance Association, Northern Finance Association, Paris December Finance Conference, Microstructure Online Seminar Asia Pacific, University of Miami Conference on Machine Learning and Business, Women in Microstructure Meeting, Babson College, Commodity Futures Trading Commission, Macquarie University, Stockholm University, University of New South Wales, and University of Sydney for valuable comments. Shkilko acknowledges support of the Social Sciences and Humanities Council of Canada and the Canada Research Chairs program. The most recent version of the paper may be downloaded from: https://bit.ly/3cR8bIF.

†University of Sydney, Australia, e-mail: amy.kwan@sydney.edu.au

‡University of Sydney, Australia, e-mail: richard.philip@sydney.edu.au

§Wilfrid Laurier University, Canada, e-mail: ashkilko@wlu.ca

1. Introduction

Among the functions of financial markets, incorporation of relevant information into prices is particularly important. To date, the finance literature has made substantial progress in understanding this function, often referred to as *price discovery*. Starting with the early theoretical models that focus on information flowing into prices through trades, the literature has progressed to models that recognize a non-trivial role of limit orders. Empirical research has followed a similar pattern, with early studies focusing on trade-driven price discovery and more recent work emphasizing limit orders.

In addition to modeling direct price impacts of orders and trades, the theoretical literature has made significant inroads into understanding the effects of market conditions that prevail when orders arrive to exchanges. Goettler, Parlour, and Rajan (2009) and Roşu (2020) show that the current state of the limit order book should play an important role in price discovery, affecting the informativeness of both market and limit orders. Bhattacharya and Saar (2021) further highlight the importance of such conditions as the bid-ask spread and order arrival frequency. Riccó, Rindi, and Seppi (2021) predict that price discovery should also be conditional on prior order history and should vary through time. Given the progress of theory, introducing conditional price discovery analyses to empirical research seems sensible. Yet such analyses are challenging when researchers use conventional methodologies such as vector autoregressions (VARs). First, multiple conditioning variables often make VAR estimation difficult or even impossible. Second, traditional VARs may not be successful in capturing important non-linear relations that permeate price discovery.

We circumvent these challenges using a machine learning technique commonly known as Reinforcement Learning (RL). Easley, Lopez de Prado, O'Hara, and Zhang (2020) argue that machine learning is well-suited to shed light on market structure issues that traditional methodologies struggle to address, and that it shines in complex environments characterized by multiple variables. Philip (2020) shows that RL estimates of permanent price impacts match VAR estimates in simple environments, but substantially outperform VAR estimates when multiple conditioning variables are included in the model. To put it plainly, RL is a flexible technique that is able to shed light on relationships among multiple theory-based variables and to do so with relative ease . Section 4 describes RL estimation in an accessible yet detailed manner.

The data strongly support theoretical predictions in that price discovery occurs through both market and limit orders and is dependent, to a large degree, on the current state of the limit order book, price history, bid-ask spread, order arrival frequency, and order submission timing.[1] Notably, consistent with Riccó, Rindi, and Seppi (2021), the information content of arriving orders may be opposite their direction and aggressiveness, and this relation depends crucially on the state of the limit order book. For instance, large *buy* orders, both market and limit, typically have *negative* permanent price impacts when the order book imbalance is negative and vice versa.

Not only does RL allow a researcher to condition price discovery on multiple variables, but it also facilitates a ranking of these variables by relative importance. In particular, the data show that the current state of the limit order book is at least as important a determinant of information flow into prices as individual order submissions. Price history follows as a close third. The size of orders, the prevailing bid-ask spread, and order submission frequency also serve as important conditions affecting information flow. To our knowledge, these results represent the first attempt to determine the pecking order of the price discovery conduits. In addition, the data show that it truly takes a village of conduits to direct the flow of information into prices. All theoretically-predicted conduits are important, and even the smallest one, the time effects, is responsible for more than 6% of information flow.

In addition to ranking the conditions that affect information flow into prices, our analyses contribute to the line of research into trading behavior of informed agents. A number of theoret-

[1]The distinction between limit order submissions and the current state of the book may appear subtle to some readers. To clarify, for each newly submitted limit order, the state of the book reflects the accumulation of previously submitted, and not yet executed or cancelled, limit orders.

ical models seek to understand this behavior, often delivering contrasting predictions. Namely, depending on factors such as the value of information, its longevity, or the number of informed agents, information may flow into prices either quickly or slowly. Recent empirical research reports results supportive of both possibilities. Shkilko (2020) and Akey, Grégoire, and Martineau (2021) show that some informed agents move prices quickly, while Collin-Dufresne and Fos (2015) and Kacperczyk and Pagnotta (2019) find that others move them slowly if at all. One notable characteristic of the latter group is its ability to time the market.

Collin-Dufresne and Fos (2016) model this ability by augmenting Kyle (1985) and allowing noise trading to be stochastic. They show that an informed agent willing to exploit stochastic noise to disguise her trading may stay undetected by the market. Our data corroborate such behavior; the most informed agents tend to time periods of returns opposite the direction of their own trading. Bhattacharya and Saar (2021) adopt similar reasoning while modeling a fully dynamic limit order market. Their informed traders hide by mimicking the actions of the uninformed. To do so, they trade through market orders when the spread is narrow and through limit orders when the spread is wide. The data strongly support these predictions; price discovery through market orders increases when the spread is narrow, and price discovery through limit orders increases when the spread is wide. Further corroborating Bhattacharya and Saar (2021), price discovery depends on the frequency of order arrivals.

Theoretical models allow information flow into prices to be time variant, that is, the price impacts of market and limit orders may depend on the time of day. The literature on intraday price discovery patterns has mainly focused on market orders, with early work finding that information flows into prices in the morning and in the afternoon. Recent evidence suggests that this pattern has changed rather notably due to the proliferation of indexers that tend to rebalance portfolios closer to the end of the day (Jiang and Yao (2021)). Our analyses support this evidence; the most informed market orders tend to be submitted during the morning hours, followed by the decline in informedness during the day and a further decline by market close. Notably, limit order informativeness, also high in the morning, declines more rapidly from open to close consistent with Riccó, Rindi, and Seppi (2021), who suggest that as the probability of execution declines through the day, informed trading should shift away from limit orders. To our knowledge, the latter result is new to the literature.

Finally, in an empirical application of our approach, we seek to shed new light on the behavior of a prominent group of market participants, high-frequency traders (HFTs), who are known to be sophisticated monitors of market conditions. Brogaard, Hendershott, and Riordan (2014) show that HFTs derive some of their trading signals from the limit order book, and we ask if the other price discovery conduits discussed by the theoretical literature may also affect HFT trading strategies. The data confirm that the state of the order book is a vital source of information for HFT firms, and show that the remaining conduits also play important roles. Together the state of the book, price history, the bid-ask spread, order arrival frequency, and time effects are related to nearly one-half of information flow generated by the HFTs.

Our contribution to the finance literature is three-fold. First, we holistically examine theoretical predictions that the state of the limit order book, prior order history, bid-ask spread, order arrival frequency, and time effects are important conditions affecting price discovery. Second, we describe a machine learning technique for price discovery analyses that is considerably more flexible than conventional methodologies, especially when it comes to accounting for multiple conditioning variables and non-linearities. In addition to its flexibility, RL is fast. For relatively simple models with only a few conditioning variables, Philip (2020) reports that RL estimates compute up to 245 times faster than the VAR estimates. As the models become more complex, the time savings increase exponentially. Finally, we provide the first ranking of theory-based price discovery conduits and show that these conduits have substantial influence on trading strategies of sophisticated market participants.

While our approach may be a useful alternative to traditional price discovery analyses, we should point out its limitations. First, we acknowledge that permanent price impacts are notoriously difficult to estimate with precision. As such, and similarly to the estimates derived from conventional VARs, the price impacts obtained using RL are empirical estimates and should be viewed accordingly. This said, and as we show shortly, RL is rather successful at capturing the permanent component of price impacts and distinguishing it from the temporary component.

Second, our analysis focuses exclusively on the price discovery conduits proposed by theory models. It is therefore possible that we do not examine all existing conduits. While this may be perceived as a limitation of our approach, we note that focusing on theory-based variables helps us avoid the common 'black-box' critique of ML methods. This critique would have applied if we had asked the machine to examine a large set of variables and select the ones that had a relation to price discovery no matter whether the economics of such a relation were traceable. Instead, we ask the machine to examine a set of theoretically motivated economic links. In the next section, we further detail the literature that has proposed and examined such links.

2. Prior literature

The price discovery process is examined by a number of theory models. The early models of dealer markets assume that this process is mainly driven by market orders (e.g., Glosten and Milgrom (1985); Kyle (1985)). More recent models recognize a significant role of informed limit orders either in a dealer or a limit order market setting (e.g., Kumar and Seppi (1994); Kaniel and Liu (2006); Goettler, Parlour, and Rajan (2009), Brolley and Malinova (2017), Roşu (2020), Riccó, Rindi, and Seppi (2021)), Bhattacharya and Saar (2021). As we mention previously, the latest models emphasize market conditions at the time of order arrival as possible determinants of price discovery.

Empirical research generally follows a similar historical pattern, with early studies focusing on trade-driven price discovery (Hasbrouck (1991a), Hasbrouck (1991b)), and more recent work emphasizing the role of limit orders (Bloomfield, O'Hara, and Saar (2005), Cao, Hansch, and Wang (2009), Cont, Kukanov, and Stoikov (2013), O'Hara (2015), Fleming, Mizrach, and Nguyen (2018), and Brogaard, Hendershott, and Riordan (2019)). Our work contributes to this literature by formalising the role of market conditions that affect price discovery, adding new conditions suggested by recent theory, and ranking the channels through which information flows into prices.

Within the above-mentioned empirical literature, our study is closely related, and complementary, to recent work of Brogaard, Hendershott, and Riordan (2019), who focus on order types and show that although individual market orders are the most informed, limit order submissions dominate price discovery as a group because their quantities far exceed those of market orders. We complement the findings of this study in two ways. First, while Brogaard, Hendershott, and Riordan (2019) focus on order arrivals, we extend the list of information conduits beyond arrivals and show where the arrivals fit in the pecking order of all conduits.

Second, our approach accounts for non-linear relations and variable interactions that permeate price discovery. Among these, a particularly well-known relation is that between price impact and order size. Hasbrouck (1991a) points out that ignoring this relation may lead to incorrect price discovery inferences. Hausman, Lo, and MacKinlay (1992) and Muravyev (2016) report that price impact functions are concave in order size for equities and options, while Philip (2020) shows that conventional VARs often struggle to capture this concavity. In addition, recent theory models (i.e., Bhattacharya and Saar (2021), Riccó, Rindi, and Seppi (2021)) emphasize the importance of interactions between the state of the order book, the bid-ask spread, and other conduits. In summary, our approach complements previous research by painting the most comprehensive picture of price discovery conduits to date, allowing for conditional order submissions, and accounting for non-linearities and variable interactions.

3. Data

We use the complete limit order book data from the Australian Securities Exchange (ASX), Australia's primary equity market. The data come from the SIRCA database and cover January and February 2016. The sample includes 20 largest stocks listed on the ASX. The SIRCA database contains every trade, order submission, amendment, and cancellation, allowing us to reconstruct the order book. A uniquely valuable feature of this database is that it contains broker identifiers. We use these in a subsequent section to understand how different market participants (i.e., proprietary high-speed traders and retail traders) use information derived from the conduits.

We prepare the data for the analysis as follows. First, we reconstruct the limit order book. To account for order splitting, we consolidate consecutive trades reported with the same time stamp, executed in the same direction, at the same price, and initiated by the same broker into one trade. Second, we determine the prevailing bid and ask immediately prior to every order and trade and use them to compute the price impacts. Third, we only include orders and trades that occur during continuous trading hours, effectively omitting opening and closing auctions.

Table 1 contains summary statistics for the 20 sample stocks, ranging from the least active Insurance Australia Group (IAG) with the volume of approximately 72,000 trades to the Commonwealth Bank of Australia (CBA) with over 386,000 trades. The sample covers a range of stock prices, from AUD 3.11 to 103.89, a range of bid-ask spreads, from AUD 0.01 to 0.024, and a range of average trade sizes, from 90 to 6,728 shares.

[Table 1]

4. The Reinforcement Learning Model

In this section, we describe the estimation of the RL model based on Philip (2020).

4.1 Model Parameters

Estimating permanent price impacts via RL requires the researcher to select two key parameters: *States* and *Actions*. A state s is the current market condition at the time an action a takes place. For example, in a very simple application of the model, there is only one state – the limit order book is active. As such, this specification is similar to a typical VAR model of price impact. In contrast, a more complex RL setup that we will explore shortly allows for different states of the limit order book, different liquidity states, time variation and so on. For example, one can describe the state of the book as either a negative order imbalance or a positive order imbalance. To add further realism, the imbalances can be subdivided into large and small.

In turn, actions a are the choices available to the market participant. In a simple RL model, the agent may choose between two possible actions, to buy or to sell. A more complicated model allows for an expanded action space, including limit order submissions, cancellations, market order submissions, and also a choice of order size.

Once the states and actions are defined, they are combined to form state-action pairs (s,a). As an example, a model with 3 states (negative book imbalance, no imbalance, positive imbalance) and 2 actions (to buy or to sell) results in $3 \times 2 = 6$ state-action pairs. In the most comprehensive version of our estimation, the one that takes into account all predictions of current theory, we define 90,000 pairs. Upon defining the pairs, each observation in the dataset is classified into one of the state-action pair categories. The number of observations may vary across these categories. For example, there is likely to be more observations in the (negative imbalance, submission of a marketable sell order) category than in the (negative imbalance, submission of a marketable buy order) category.[2] Given our large sample size, this imbalance does not however impair RL estimation.

[2]To conserve space, we shorten action descriptions such as these to (neg imb, submit mkt buy) in the remainder of the manuscript.

4.2 Estimating the RL Model

Once the parameters are defined, we estimate the model as follows:

$$Q(s,a) = R(s,a) + \gamma \sum_{s' \in S} \sum_{a' \in A} T((s,a),(s',\overline{a}))Q(s',\overline{a}), \tag{1}$$

where $Q(s,a)$ is the permanent price impact of action a in state s, $R(s,a)$ is the immediate price impact of action a in state s, γ is a discount factor, and $T((s,a),(s',a'))$ is the probability of transitioning from the current state-action pair (s,a) to a future state-action pair (s',a').

In equation (1), estimating the model requires three inputs: the immediate price impact R, transition probability T, and discount factor γ. We discuss these in turn. The immediate price impact refers to the change in the log quote midpoint as a result of an action. For instance, the immediate price impact of a marketable order submission is the midquote change from immediately prior to the submission to a submission or cancellation of the next order in the same stock, no matter whether the next order is marketable or limit.

Since we need an immediate impact for each state-action pair, the above-mentioned example with 6 state-action pairs requires 6 immediate impacts. To compute these, we first calculate the immediate impact for every observation in the data. Then, for each state-action pair we compute the average immediate impact across all observations that belong to the pair. For example, to calculate the immediate impact for the state-action pair (neg imb, submit mkt sell), we take all transactions initiated by marketable orders to sell at the time of a negative imbalance and compute the average change in the log midpoint that follows them.

In turn, the transition probability is the probability of observing the future state-action pair (s',a') after observing the current state-action pair (s,a). For example, to calculate the probability of transitioning from the state-action pair (neg imb, submit mkt sell) to the state-action pair (pos imb, submit mkt buy), we compute the number of times the first pair is followed by the second and express this number as a fraction of all (neg imb, submit mkt sell) pairs.

Finally, γ is a discount factor representing the time value of money that lies between 0 and 1. More precisely, $0 <= \gamma < 1$, as γ must be less than 1 to ensure convergence. Because the price impacts of trades are computed over very short time horizons, γ should be very close to 1. We set $\gamma = 0.999$.

Intuitively, equation (1) shows that the permanent price impact of action a is the immediate impact caused by the action and the expected permanent price impact of all subsequent actions. The latter is captured by the $\sum_{s' \in S} \sum_{a' \in A} T((s,a),(s',a'))Q(s',a')$ term and depends on the probability of observing subsequent actions a' while in s' and the permanent price impact of these subsequent actions.

We note that both the right-hand side (RHS) and the left-hand side (LHS) of equation (1) contain a $Q(s,a)$ term. To obtain the estimates of $Q(s,a)$, we use an iterative learning rule known as *Q-learning*. Specifically, for the first iteration, we initialize all $Q(s,a)$ estimates on the RHS to zero and solve for $Q(s,a)$ on the LHS. Thus, in the first iteration, the price impact of state-action pair (s,a) is equal to the immediate price impact, as equation (1) collapses to $Q(s,a) = R(s,a)$. In the second iteration, we substitute the updated LHS values of $Q(s,a)$ from the first iteration into the RHS values of $Q(s,a)$ and re-estimate the LHS. This process continues until the LHS and RHS converge, which implies that $Q(s,a)$ values no longer change with subsequent iterations. Additional estimation details for Q-learning may be found in Philip (2020).

The readers, who are familiar with the RL approach, may recognize that it is often used for optimization, or selecting optimal actions in a given state. Such optimization is not the focus of this study. Rather, we use another unique feature of RL, that is its ability to estimate the permanent impact of an action when there are multiple possible subsequent states of the world. It is also important to point out the advantage of RL in our setting over other popular ML methods (e.g., neural networks or random forests). RL differs from these methods because it is able to estimate Y (the permanent price impact in our case), whereas the alternative techniques typically require that Y is observed.

4.3 RL and VAR: A Comparison

When examining price discovery, the finance literature usually relies on vector autoregression (VAR) techniques to estimate permanent price impacts. As such, it is important to show that the price impact estimates from the RL model are comparable to the VAR estimates in cases when both models can be estimated. To do so, we estimate a standard VAR model as follows:

$$\begin{aligned} r_t &= \sum_{i=1}^{10} \alpha_i r_{t-i} + \sum_{i=0}^{10} \beta_i x_{t-i} + \varepsilon_{1,t} \\ x_t &= \sum_{i=1}^{10} \delta_i r_{t-i} + \sum_{i=1}^{10} \phi_i x_{t-i} + \varepsilon_{2,t}, \end{aligned} \tag{2}$$

where r_t is the midpoint return for a trade at time t, and x_t is a trade indicator variable equal to 1 for buys and -1 for sells. To ensure a fair comparison, we estimate the RL model using parameters similar to the VAR; with one state (limit order book is active) and two actions (buy or sell). These two actions are analogous to the $+/-1$ trade sign indicator in the VAR model.

We report the resulting permanent price impact estimates in Table 2, with estimates based on the VAR model in Panel A and those based on the RL model in Panel B. The average price impact estimated via the traditional VAR model ranges from 0.94 basis points (bps) for CBA to 5.20 bps for STO. Using CBA as an example, these estimates suggest that a buy (sell) transaction permanently moves the price by an average of 0.94 (-0.94) bps.

[Table 2]

Panel B contains the corresponding price impact estimates from the RL model. In contrast to VAR, which produces the same price impact magnitudes for buys and sells, the RL model produces separate estimates. As should be expected, these estimates have rather similar magnitudes. For example, the first stock in the sample (AMC) has the price impact of -2.49 bps for sells and 2.53 bps for buys. Observing across all stocks and all buy and sell estimates, the magnitude of price impacts ranges from 0.97 (CBA) to 6.80 (STO). We note that both the RL and VAR models produce the lowest price impact estimates for CBA and the highest for STO.

To further compare between RL and VAR estimates, we report the bounds of the 95% confidence interval for the VAR estimates. The results confirm that the two models produce comparable results: 37 out of 40 RL estimates fall within the 95% confidence interval of the VAR estimates. Importantly, correlation between the VAR estimates and the average of the RL buy and sell estimates is over 0.99. As such, the results corroborate the notion that VAR and RL produce comparable results when the state-action space is relatively simple. In the following sections, we explore more complex state-action spaces, adding conditioning variables and allowing for non-linearities. To reiterate, Philip (2020) argues that such spaces are either difficult or impossible to estimate using conventional VARs.

One of the useful features of the VAR is its ability to generate impulse response plots that capture the dynamics of price reactions and help gauge the speed with which information flows into prices. RL is able to generate similar plots as Figure 1 illustrates. As we mentioned earlier, when solving equation (1), the price impact estimate *Q(s,a)* updates to a new value after each iteration until convergence. By construction, after the initial iteration the estimate contains only the immediate price impact, while upon convergence it contains both the immediate and all the subsequent impacts. Thus, as we observe the progression of the estimates between the first iteration and convergence, we are able to see how the price impact evolves from the immediate (and possibly temporary) impact to the permanent impact. The resulting function is visually and conceptually similar to an impulse response function. Panel A of Figure 1 shows that for a typical buy market order the two functions are very similar and therefore likely capture the same economic relations.

[Figure 1]

Another useful aspect of the graphical representation of RL convergence is its ability to distinguish between permanent and temporary price impacts. Panel B of Figure 1 plots an RL convergence function for a buy order submitted when the limit order book imbalance is negative. As we mention in the Introduction and show in detail in the next section, such orders are generally uninformed and have negative permanent price impacts. Panel B shows that this is indeed the case. Notably, the initial price impact of such buy orders is positive but quickly reverses, and at the end the permanent price impact is negative.

4.4 Non-linearities and Expanding the Action Space

Machine learning techniques such as RL are uniquely suited to capture nonlinear relations. One such potentially important relation is between price impacts and trade/order sizes. Philip (2020) shows that ignoring the nonlinear effects of trade sizes often leads to incorrect inferences. In the previous section, we estimate the RL model using only two actions, a buy trade and a sell trade irrespective of size. In this section, we enhance the model to incorporate size effects.

In addition, the previous section focuses solely on trades. Meanwhile, Brogaard, Hendershott, and Riordan (2019) show that limit orders carry significant amounts of information into prices and therefore should be examined along with trades in price discovery analyses. To do the same, we expand the RL analysis to include both trades and limit orders. For consistency of exposition and to conserve space, we refer to trades as *market orders*, even though we recognize that the majority of trades are triggered by marketable orders. When it comes to limit orders, we examine all top of the book order submissions and cancellations.

We begin by estimating the RL model using one state (the order book is active) and 6 order categories: sell market orders, buy market orders, sell limit order submissions, buy limit order submissions, sell limit order cancellations, and buy limit order cancellations. In addition, for each order category we use 5 order size quintiles, resulting in $6 \times 5 = 30$ unique actions. As such, this model has $1 \times 30 = 30$ state-action pairs. To facilitate a fair comparison, we use the same size quintile cutoffs for each action, while allowing the cutoffs to vary across stocks.

Figure 2 plots the permanent price impact estimates for market orders (Panel A), limit order submissions (Panel B), and limit order cancellations (Panel C). Each point in the figure represents the average permanent price impact for a particular size quintile across the sample stocks. The x-axis plots the average value for each size quintile, with negative values representing sell actions and positive values representing buy actions.

[Figure 2]

Panel A contains an S-shaped pattern indicating that price impacts increase in trade size, albeit at a decreasing rate. This finding is consistent with the economic intuition that larger trades should have greater price impacts than smaller trades. Perhaps equally importantly, the data show that the relation between the price impact and trade size is concave. That is, increasing the trade size from 901 shares to 1,000 shares is associated with a smaller increase in price impact than increasing the trade size from 1 share to 100 shares.

Panel B contains a similarly-shaped pattern for limit order submissions, although its magnitude is smaller compared to market orders, and non-linearity is generally absent. The former result is in line with Brogaard, Hendershott, and Riordan (2019), who report that individual limit orders have relatively small price impacts in the VAR setting. Finally, for limit order cancellations reported in Panel C, we expectedly observe a reversal in the S-shape, consistent with the notion that a sell (buy) order cancellation signals a future price increase (decrease). We note that a 'hump' observed for the second order cancellation size category should be interpreted with some caution, as statistical testing (discussed shortly) suggests that the price impact estimates for this category are relatively noisy.

We examine these relations in a more formal setting in Table 3, which contains average permanent price impacts for each action. For instance, the price impact of a large market order to sell (*Sell Q5*) is -3.54 bps, compared to -0.83 bps and 0.13 bps for, respectively, a sell limit order submission and sell limit order cancellation in the same size quintile. We also test for differences in price impacts between market and limit order submissions (column 4) and for differences in the absolute values of price impacts between limit order submissions and cancellations (column 5). The data confirm that market orders have larger price impacts than limit order submissions or cancellations, and that limit order submissions have larger price impacts than cancellations in the majority (80%) of cases.

[Table 3]

The results suggest that limit order submissions are often more informative than cancellations. Yet submissions and cancellations are opposite actions, and their contributions to price discovery may be expected to perfectly offset each other. Why is this not the case? One possibility is that many limit order cancellations occur after the book moves away from these orders. For instance, an informed limit order to buy at the prevailing bid of $10.00 submitted in anticipation of a price increase may be cancelled after a period of time not because the submitter now anticipates a price decline, but because the price has moved up, and execution of the original order is no longer likely.

Even though individual limit orders have smaller price impacts than individual market orders, the number of limit order submissions is so large that total price discovery coming from them may surpass that coming from other order types. To examine this possibility, we investigate the contribution to total price discovery attributable to market and limit orders after adjusting for the frequency of each order type. Consistent with the earlier literature, Panel B of Table 3 reports that market orders, limit order submissions, and limit order cancellations comprise 16.65%, 53.01%, and 30.34% of all actions. Taking these proportions into account, RL further shows that the three actions contribute 53.8%, 35.7%, and 10.5% to overall price discovery.

At first glance, it may appear that these results contradict Brogaard, Hendershott, and Riordan (2019), who report that limit order submissions and cancellations lead price discovery on the Toronto Stock Exchange (TSX). We however believe that our results should be viewed as complementary to theirs, with the differences most likely attributable to the unique characteristics of the two markets. First, Australian trading is less fragmented than trading in Canada, and therefore cross-venue market making and arbitrage activities, which typically involve large quantities of limit order submissions and cancellations (e.g., van Kervel (2015); Shkilko and Sokolov (2020)), are less prevalent on the ASX. Specifically, Panel B shows that limit order submissions and cancellations make up 83% of all ASX activity. In contrast, they make up more than 95% of all activity on the TSX. Consistently, there is less price discovery through limit orders on the ASX relative to the TSX.

Second, Brogaard, Hendershott, and Riordan (2019) show that it is mainly price discovery by high-frequency traders (HFT) that occurs via limit orders, while non-HFT price discovery is dominated by market orders. Our sample includes both trader categories, and as such delivers results for an average trader. Finally, it is also possible that our results are unique to the RL technique. To shed light on this possibility, we estimate a similarly structured VAR. Our conclusions are unchanged; Panel B of Table 3 reports that the RL and VAR results are statistically indistinguishable from each other.

5. Conditioning Price Discovery

In this section, we expand the set of price discovery conduits beyond order types and sizes. To do so, we follow the theoretical predictions of Goettler, Parlour, and Rajan (2009), Roşu (2020), Riccó, Rindi, and Seppi (2021), Bhattacharya and Saar (2021) and extend the RL model to include five conditioning variables: the current state of the limit order book, price history, bid-ask spread, order arrival frequency, and order timing. For these analyses, we estimate the RL model with the same 30 actions described earlier while changing the states to reflect changes in the conditioning variables. For instance, we distinguish the state with a large negative limit order book imbalance from that with a large positive imbalance. At the end of this analysis, we rank the contribution of each conditioning variable, effectively answering the question: Which information conduits are most important in determining future price movements?

5.1 Current State of the Limit Order Book

To gain insight into the price discovery effects originating from the limit order book, we estimate the RL model using 10 states (book imbalance deciles at the time of each action) and 30 actions described earlier resulting in a model with $10 \times 30 = 300$ state-action pairs. To define the states, we calculate the limit order book depth imbalance (*DI*) at the time of each market order submission, limit order submission, and limit order cancellation as follows:

$$DI = \frac{\sum_{i=1}^{5} DEP_{Bid,i} - \sum_{i=1}^{5} DEP_{Ask,i}}{\sum_{i=1}^{5} DEP_{Bid,i} + \sum_{i=1}^{5} DEP_{Ask,i}}, \tag{3}$$

where $DEP_{Bid,i}$ is the total bid depth, and $DEP_{Ask,i}$ is the total ask depth at price level i, for a total of ten price levels (five ask and five bid levels). We then divide *DI*s into ten deciles corresponding to the ten price levels and assign each order submission and cancellation to one of the deciles, from the most negative imbalance to the most positive one. Figure 3 plots the price impact estimates for the most negative (*D1*, dotted line with cross markers) and the most positive (*D10*, dashed line with circle markers) deciles. To facilitate comparison, the solid lines replicate the average case from Figure 2.

[Figure 3]

We begin the discussion of Figure 3 by noting that the difference in price impacts between the leftmost and the rightmost markers of any of the *S*-functions is smaller than the difference between the *D1* and *D10* functions. The former difference represents the price discovery contribution of order direction and size, and the latter – of the book imbalance. For example, when a large positive imbalance exists (*D10*), a large market buy order has a price impact of 9.73 bps, while a large market sell order has a price impact of 1.97 bps. The difference of 7.76 bps between these figures represents the price discovery contribution of order size while holding the book imbalance constant. In contrast, when we hold the size constant and vary the imbalance, the contribution of the latter is larger than that of the former. For example, the price impact of a large market buy order in *D10* is 9.73 bps, while the impact of a similarly sized market buy in *D1* is -1.85 bps, a difference of 11.58 bps.

In addition to highlighting the importance of the limit order book in the price discovery process, Figure 3 reveals a curious regularity. Namely, the information content of arriving orders may be opposite their direction and aggressiveness, and this relation depends crucially on the state of the book. More specifically, large buy orders, both market and limit, have negative permanent price impacts when the order book imbalance is negative. Similarly, large sell orders have positive price impacts when the book imbalance is positive. For instance, a large market order to buy has a -1.85 bps price impact when the book imbalance is significantly negative . This result is consistent with the theoretical predictions of Riccó, Rindi, and Seppi (2021) and, to our knowledge, is the first empirical confirmation of these predictions.

To further highlight these results, Table 4 reports the permanent price impacts across all order types, directions, and sizes as well as for the most positive (*D10*) and the most negative (*D1*) book imbalance deciles. It also reports the differences between *D10* and *D1*, in columns 3, 6, and 9 for, respectively, market order submissions, limit order submissions, and limit order cancellations. All of these differences are statistically significant suggesting that the state of the book plays an important role in price discovery across all order sizes. In the subsequent sections, we report reduced versions of this analysis for the remaining conduits, focusing specifically on the statistical significance of differences between states.

[Table 4]

The effects of the limit order book discussed thus far are conditional on an order submission or cancellation. Is it possible to gauge the unconditional effects? We note that although interesting, this question is largely hypothetical. It is important to recognize that the order book itself cannot move the midquote – our proxy for the efficient price. Only the orders submitted to (or cancelled from) the top of the book can do so. These orders are already the focus of our main tests.

This said, the unconditional effects may be approximated. To do so, we make one simplifying assumption that is relaxed later in a regression setting – that order submissions and cancellations occur with equal probability upon observing the current book state. With this assumption in mind, recall that the average contributions of the book for each action are reported in columns 3, 6, and 9 of Table 4. Averaging these gives the unconditional price impact of 11.3 bps as illustrated by the horizontal grey lines in Figure 3. To clarify, this result suggests that when the limit order book is in state *D1* (*D10*), the future price should be expected to decrease (increase) by 5.6 bps.

Overall, the results in this section highlight the importance of the current state of the order book in influencing prices. Some theoretical models assume that market orders come from informed investors, whose information may be short-lived or highly valuable and therefore may require prompt action (e.g., Kaniel and Liu (2006), Chau and Vayanos (2008)). Our results suggest that this is not always the case. Consistent with Riccó, Rindi, and Seppi (2021), even a large market buy order may have a negative effect on prices if the limit order book imbalance is negative. More generally, two orders of the same size may have notably different price effects depending on the state of the book during their submission or execution.

5.2 Price History

The results in the previous section show that the current state of the limit order book is an important determinant of price formation. Riccó, Rindi, and Seppi (2021) suggest that prior order history should have information beyond the current state of the book. To examine this suggestion, we must choose an appropriate proxy for prior order history, a metric that captures all changes in the order book in a single variable. We suggest that prior price movements ultimately reflect a substantial portion of order history and use them as a proxy.

To investigate the impact of historical price movements on price discovery, we estimate an RL model with 10 states and 30 actions as before. This time however, the states are determined based on midpoint returns (price history) over ten actions prior to the current action. As such, we focus on event time as recommended by Easley, López de Prado, and O'Hara (2012) rather than clock time. We classify each observation into return deciles *Ret1* to *Ret10*, containing observations with the most negative and the most positive returns, respectively.

We plot the price impact estimates for the most negative (*Price Down*) and most positive (*Price Up*) price histories in Figure 4. Return timing appears to play a non-trivial role. For instance, market buy orders that follow large negative returns (circle markers) have greater price impacts than buys that follow large positive returns (cross markers). Similarly, sells following large positive returns have greater price impacts than sells following large negative returns. Table 5 reports that these results are highly statistically significant across all order size categories.

[Figure 4 and Table 5]

These findings add new supportive evidence to the literature that aims to understand trading behavior of informed investors. Collin-Dufresne and Fos (2016) show theoretically that such investors may choose to time their trading to reduce trading costs. In our context, this implies that the informed prefer to trade after prices move in the direction opposite their trading interest. The data support this possibility; the most informed trades in Figure 4 indeed occur after prices move in the opposite direction.

5.3 Bid-Ask Spread

Bhattacharya and Saar (2021) propose a model in which informed traders conceal their presence in the market by mimicking uninformed traders. The uninformed tend to demand liquidity when it is cheap and supply it when it is expensive. In their mimicking efforts, the informed therefore trade through market orders when the spread is narrow and through limit orders when the spread is wide. Consequently, information flows into prices mainly via market orders when the spread is narrow and limit orders when the spread is wide.

To examine this possibility, we augment the RL model to include two states defined by the size of the bid-ask spread: *narrow* and *wide*. The former includes observations when the spread equals to the tick size (AUD 0.01), and the latter – when the spread is greater than the tick size. We opt to use AUD 0.01 for the narrow state because our sample consists of frequently traded securities, for which the spreads are usually close to the minimum tick size as we show in Table 1. As before, we estimate an RL model with 30 actions, resulting in $2 \times 30 = 60$ state-action pairs.

The data fully corroborate the above-mentioned theory predictions. Figure 5 plots the price impact estimates for the *narrow* and *wide* states and shows that market orders arriving when the spread is narrow contribute more to price discovery than market orders arriving when the spread is wide. The opposite is true for limit orders, whose contribution to price discovery is considerably greater when the spread is wide. Bhattacharya and Saar (2021) do not make explicit predictions for limit order cancellations, so we report the cancellation results largely for completeness. Table 5 indicates that the results are statistically significant for all limit order size categories and for all but the two market order categories that represent the largest sizes.

[Figure 5]

5.4 Order Arrival Frequency

In Bhattacharya and Saar's (2021) model, order arrivals intensify when the informed are present. This effect is conditional on the order type the informed choose, which in turn depends on the spread as discussed earlier. More specifically, when the spread is narrow, and the informed choose market orders, more frequent market order arrivals serve as stronger signals. The same applies to limit orders when the spread is wide.

To test these predictions, we augment the RL model from the previous section to include five additional states that capture order arrival frequency. Specifically, we compute the time between the current and previous actions of the same type (e.g., from one market order submission to another) and form five arrival frequency quintiles, from lowest to highest. Thus, the new model contains 10 states (2 spread states and 5 arrival frequency states) and 30 actions from before, resulting in $10 \times 30 = 300$ state-action pairs. For illustrative purposes, in the lowest frequency quintile orders arrive every 12.8 seconds and in the highest quintile – every 0.46 milliseconds.

The results reported in Figure 6 are consistent with theory predictions. When the spread is narrow, market orders arriving frequently have larger price impacts compared to market orders arriving infrequently. A similar result holds for limit orders when the spread is wide. Table 5 shows that these results are statistically significant for all order size groups but one.

[Figure 6]

5.5 Time Effects

So far, our results show that the current state of the limit order book, price history, the bid-ask spread, and order arrival frequency contribute to price discovery. In this section, we ask if the timing of order submission matters as well, as suggested by Riccó, Rindi, and Seppi (2021). To examine this issue, we augment the RL model to include the time of day effects by splitting the trading day (from market opening at 10:00 a.m. to market closing at 4:00 p.m.) into three states:

10:00 - 10:30 (morning), 11:30 - 14:30 (midday), and 15:30 - 16:00 (afternoon). As before, we estimate the RL model with 30 actions, resulting in a final model consisting of $3 \times 30 = 90$ state-action pairs.

Recent research points to an *S*-shaped intraday spread pattern, whereby greater spread values are observed early in the day and lower values are observed near market close (Upson and Van Ness (2017)). This pattern is linked to the proliferation of indexing institutions that tend to rebalance their portfolios closer to the end of the day, thereby reducing order flow informedness (Jiang and Yao (2021)). Consistent with this recent literature, the results in Figure 7 show that trading activity in the morning contributes more to price discovery than trading activity in the midday and afternoon sessions.

[Figure 7]

Viewed across order types, the results suggest that market order informativeness declines by 25% from the morning to the afternoon sessions. Notably, the informativeness of limit order submissions, also high in the morning, declines more significantly from open to close, by 31%. This result is consistent with Riccó, Rindi, and Seppi (2021), who suggest that as the probability of execution declines through the day, informed trading should shift away from limit orders.

5.6 Contribution to Price Discovery

The results so far confirm the existence of several price discovery conduits proposed by the theoretical literature. In this section, we probe deeper to understand the *relative* importance of these conduits in a comprehensive model that includes all of them simultaneously.

For this investigation, we follow Easley, Lopez de Prado, O'Hara, and Zhang (2020) and compute the *mean decreased accuracy* (MDA) metric for each identified conduit. The MDA is commonly used in machine learning, but is relatively new to finance. As we explain in more detail shortly, it measures the decrease in accuracy of the price impact estimate if one of the conduits (e.g., order submissions, state of the book, etc.) is intentionally estimated with error. Importantly, the MDA does not assume a specific functional form and, compared with a regression analysis, automatically accounts for non-linearities and variable interactions.

Estimating the MDA requires two parameters: the true predicted price impact and the randomized predicted price impact. The true predicted price impact $\hat{P}I_{s,a}$ is the estimated price impact for each state-action pair (s,a) based on a comprehensive RL model with 30 actions and 3,000 states, all of which are separately discussed in the earlier sections. As previously, the 30 actions include 6 order categories (sell market orders, buy market orders, sell limit order submissions, buy limit order submissions, sell limit order cancellations, and buy limit order cancellations), with each category containing five size quintiles. The 3,000 states include 10 limit order book state deciles, 10 price history deciles, 2 bid-ask spread states, 5 order arrival frequency quintiles, and 3 time of day states $(10 \times 10 \times 2 \times 5 \times 3 = 3,000)$. We estimate this model to obtain $\hat{P}Is$ for $3,000 \times 30 = 90,000$ state-action pairs.

In turn, to obtain the second parameter, the randomized predicted price impact $\tilde{P}I^{k}_{s,a}$, we assign a random value to a conduit k while holding all other conduits at their true values. For example, to test the relative importance of the state of the book, let us assume that the true price impacts in state-action pairs (*D1, c*) and (*D4, c*) are 1 bp and 2 bps, where D is one of the depth imbalance deciles, and c indicates that all other conduits in both pairs are the same. To randomize the depth imbalance variable, the observation (*D1, c*) is flipped to have the value of (*D4, c*). Thus, the $\tilde{P}I^{k}_{s,a}$ value for (*D1, c*) is assigned as 2 bps instead of 1 bp. Using these assignments, the MDA for the conduit k is estimated as follows:

$$MDA^{k} = \sum_{(s,a)=1}^{(S,A)} \left(\frac{|(\hat{P}I_{s,a} - \tilde{P}I^{k}_{s,a})|}{\hat{P}I_{s,a}} \right) / (S,A), \tag{4}$$

where (S,A) is the total number of state action pairs. We repeat this process 100 times for each variable k and report the mean and standard deviation of the MDA for each variable in Table 6.

[Table 6]

Since the MDA measures deterioration in estimation accuracy, a variable with a high MDA is more important than a variable with a low MDA. Table 6 suggests that the state of the book is the most important variable (MDA = 10.38), followed by order occurrence (9.25), price history (8.79), order size (6.59), arrival frequency (4.57), the bid-ask spread (4.34), and time of day (2.86). To our knowledge, this ranking represents the first attempt at establishing a pecking order of price discovery conduits.

How should one think of the limit order book having a greater MDA than order submissions and cancellations? We suggest that this ranking derives naturally from our earlier results and characterizes the market as a two-part system, in which both orders and market conditions play important price discovery roles. The first part of the system, order submissions and cancellations, physically carries information into prices. The second part, market conditions, augments and regulates this flow. Some components of the second part, e.g., the limit order book, are so crucial that they significantly affect (enhance or diminish), and occasionally overpower information carried by orders. The earlier example of a large market sell order submitted when the book is positively imbalanced illustrates this effect. Such components have high MDAs. Other components, e.g., the bid-ask spread, have smaller yet non-trivial effects and therefore smaller MDAs.

5.7 An Application: Proprietary vs. Retail Traders

Prior literature proposes that sophisticated market participants may monitor price discovery conduits and trade on information obtained from them. For instance, Brogaard, Hendershott, and Riordan (2014) argue that the limit order book is an important source of information for some high-frequency traders (HFTs). Do the other conduits also serve as determinants of their trading strategies? For instance, may HFTs learn from the frequency of order arrivals? Do they time their informed trades to the state of the bid-ask spread or to recent price movements?

To shed light on these possibilities, we compute the permanent price impacts of HFT trades and compare them to price impacts generated by traders, who are less likely to monitor the conduits. To benchmark for the latter group, we use retail investors. Even though we do not expect an average retail investor to monitor information flow into prices in real time, we note that her trades may have non-zero price impacts. This expectation is based on studies showing that retail investors often trade on price-relevant information (e.g., Kaniel, Saar, and Titman (2007), Kaniel, Liu, Saar, and Titman (2012), Barber, Odean, and Zhu (2009), Kelley and Tetlock (2013), Boehmer, Jones, Zhang, and Zhang (2020)).

A unique feature of our data is that all trades come with broker identifiers. Several HFT firms in Australia (e.g., Optiver, Susquehanna, Virtu) have brokerage status, and therefore we can see their trades in the data. We identify nine such HFT firms. Similarly, we are able to observe activities by several brokerages that cater predominantly to retail investors (e.g., ANZ Securities, Bell Potter Securities, E-Trade Australia), for a total of 28 retail brokerages. Using this information, we categorize brokerages into three categories: (i) those predominantly associated with HFT activity, (ii) those predominantly associated with retail activity, and (iii) others. Similarly to Nasdaq HFT data used in prior studies (e.g., Brogaard, Hendershott, and Riordan (2014), O'Hara, Yao, and Ye (2014)), the first category tends to capture the activity of pure HFT firms. HFT activities of larger integrated investment institutions are in the third category, as we are unable to separate them from the non-HFT activities of these institutions.

We begin the comparison between HFT and retail investors using permanent price impacts from a simple RL model. The univariate results in Panel A of Table 7 indicate that HFT trades have substantially larger price impacts than those of retail investors, 3.03 bps vs. 1.8 bps. In Panel B, we formalize this comparison in the following regression setting:

$$\hat{PI}^{HFT}_{i,s,a} = \alpha + \beta \times \hat{PI}^{Retail}_{i,s,a} + \varepsilon_{i,s,a}, \tag{5}$$

where $\hat{PI}_{i,s,a}^{HFT}$ is the estimated price impact of trades originating from the HFT brokerages, and $\hat{PI}_{i,s,a}^{Retail}$ is the estimated price impact of trades originating from the retail brokerages. The β coefficient measures the relative informedness of HFT trades compared to Retail trades, and its statistical significance is benchmarked against 1. If $\beta = 1$, then HFTs and retail traders have the same price impact functions, implying that they are equally informed.

We first estimate a baseline model that compares simple price impact functions for the two account types. This baseline model includes the two actions that are a trader's prerogative, the occurrence of the trade itself and its size. The results of the baseline model are comparable to the univariate results in that HFT trades are about 58% more informative than those of retail investors.

[Table 7]

Next, we begin augmenting the baseline estimate with a set of controls for the previously identified price discovery conduits. We add the conduits one by one, in the order of importance identified in the MDA analysis. For instance, we first assume that both HFTs and retail investors may condition their trading on the state of the limit order book (even though we believe that HFTs are much more likely to do so) and ask if the baseline relation between their price impact functions holds. We then add price history as a determinant of informed order flow timing and so on. The full model at the bottom of Table 7 includes all previously identified price discovery conduits and therefore is comprised of 90,000 state-action pairs that allow for conditioning on the state of the book, price history, the bid-ask spread, order submission frequency, and time effects.

As we allow for more conditioning variables in the HFT decision set, the baseline β coefficient of 1.58 decreases. First, it decreases substantially, to 1.01, after adding the state of the book. In other words, once we take the book away as a catalyst for HFT trades, the informedness of these trades equals that of retail trades. Other variables eventually reduce the β coefficient to 0.87, implying that HFT permanent price impacts are 13% smaller than those of retail investors once the price discovery conduits are accounted for. In absolute terms, this equates to a 1.64 bps (= 1.88 $\times$ 0.87) price impact of HFT trades, a 46% reduction compared to the base case.

The last figure suggests that the previously identified price discovery conduits represent a substantial portion, but not the entirety, of the HFT information and trade timing set. What is in the remainder of this set? We propose that even though HFTs may learn from some of the conduits and condition their trading on the state of others, a lot of their information likely comes from external sources that are not included in our group of conduits.

More specifically, it is well-documented that HFTs perform a dual function when it comes to price discovery. First, they observe information flow into prices through orders and monitor the state of the conduits that affect this flow. Second, they bring additional information from outside sources such as correlated assets (Budish, Cramton, and Shim (2015), Shkilko and Sokolov (2020)) as well as corporate and macroeconomic announcements (Hu, Pan, and Wang (2017), Chordia, Green, and Kottimukkalur (2018)). Given the structure of our tests, such outside information may not yet be reflected in the conduits and as such constitutes an additional and important component of the HFT information set.

6. Conclusion

Recent theory models condition price discovery and trader behavior on market states such as the state of the limit order book, order price history, the bid-ask spread, order arrival frequency, and time of day. In the meantime, conditional empirical analyses have not yet entered the mainstream due to methodological limitations. Specifically, the VAR models, which are often used to understand how information flows into prices, are difficult or impossible to estimate with multiple conditioning variables and when the relations between variables are non-linear.

To circumvent these limitations, we describe and apply a machine learning technique known as Reinforcement Learning (RL), which is particularly useful in applications characterized by non-linearities and multiple conditioning variables. Our findings support theoretical predictions in that price discovery occurs through both market and limit orders and is conditional, to a large degree, on the state of the limit order book and price history. The bid-ask spread and the frequency of order arrivals also play non-trivial roles in determining how information flows into prices.

Our analyses also shed light on the predictions of models of trader behavior. Consistent with these models, the most informed agents tend to time periods of returns opposite the direction of their own trading. They also time liquidity, demanding it when spreads are narrow and supplying it when spreads are wide. In an application of the RL methodology, we show that a large portion of informed trading by high-frequency traders, who are believed to be sophisticated market monitors, is dependent on the state of the conduits.

We view our contribution to the literature as threefold. First, we shed light on predictions of several theory models that have not yet been tested due to limitations of conventional methodologies. Second, we describe a new method for measuring price discovery, namely a machine learning technique that allows a computer to understand what relations, if any, exist between variables discussed by theory. Finally, we report, for the first time in the literature as far as we know, a ranking of price discovery conduits.

References

Akey, P., V. Grégoire, and C. Martineau, 2021, "Price Revelation from Insider Trading: Evidence from Hacked Earnings News," *Journal of Financial Economics*, conditionally accepted.

Barber, B. M., T. Odean, and N. Zhu, 2009, "Do retail trades move markets?," *Review of Financial Studies*, 22(1), 151–186.

Bhattacharya, A., and G. Saar, 2021, "Limit Order Markets under Asymmetric Information," *Working Paper*, University of Chicago and Cornell University.

Bloomfield, R., M. O'Hara, and G. Saar, 2005, "The "make or take" decision in an electronic market: Evidence on the evolution of liquidity," *Journal of Financial Economics*, 75(1), 165–199.

Boehmer, E., C. M. Jones, X. Zhang, and X. Zhang, 2020, "Tracking retail investor activity," *Journal of Finance*, forthcoming.

Brogaard, J., T. Hendershott, and R. Riordan, 2014, "High frequency trading and price discovery," *The Review of Financial Studies*, 27(8), 2267–2306.

Brogaard, J., T. Hendershott, and R. Riordan, 2019, "Price Discovery without Trading: Evidence from Limit Orders," *Journal of Finance*, 74(4), 1621–1658.

Brolley, M., and K. Malinova, 2017, "Informed Trading in a Low-Latency Limit Order Market," *Working Paper*, Wilfrid Laurier University.

Budish, E., P. Cramton, and J. Shim, 2015, "The high-frequency trading arms race: Frequent batch auctions as a market design response," *Quarterly Journal of Economics*, 130(4), 1547–1621.

Cao, C., O. Hansch, and X. Wang, 2009, "The information content of an open limit-order book," *Journal of Futures Markets*, 29(1), 16–41.

Chau, M., and D. Vayanos, 2008, “Strong-Form Efficiency with Monopolistic Insiders,” *The Review of Financial Studies*, 21(5), 2275–2306.

Chordia, T., T. C. Green, and B. Kottimukkalur, 2018, “Rent seeking by low-latency traders: Evidence from trading on macroeconomic announcements,” *Review of Financial Studies*, 31(12), 4650–4687.

Collin-Dufresne, P., and V. Fos, 2015, “Do Prices Reveal the Presence of Informed Trading?,” *The Journal of Finance*, 70(4), 1555–1582.

Collin-Dufresne, P., and V. Fos, 2016, “Insider Trading, Stochastic Liquidity, and Equilibrium Prices,” *Econometrica*, 84(4), 1441–1475.

Cont, R., A. Kukanov, and S. Stoikov, 2013, “The Price Impact of Order Book Events,” *Journal of Financial Econometrics*, 12(1), 47–88.

Easley, D., M. Lopez de Prado, M. O’Hara, and Z. Zhang, 2020, “Microstructure in the Machine Age,” *Review of Financial Studies*, forthcoming.

Easley, D., M. M. López de Prado, and M. O’Hara, 2012, “Flow Toxicity and Liquidity in a High-frequency World,” *Review of Financial Studies*, 25(5), 1457–1493.

Fleming, M. J., B. Mizrach, and G. Nguyen, 2018, “The microstructure of a U.S. Treasury ECN: The BrokerTec platform,” *Journal of Financial Markets*, 40, 2 – 22.

Glosten, L. R., and P. R. Milgrom, 1985, “Bid, ask and transaction prices in a specialist market with heterogeneously informed traders,” *Journal of Financial Economics*, 14(1), 71 – 100.

Goettler, R. L., C. A. Parlour, and U. Rajan, 2009, “Informed traders and limit order markets,” *Journal of Financial Economics*, 93(1), 67–87.

Hasbrouck, J., 1991a, "Measuring the Information Content of Stock Trades," *The Journal of Finance*, 46(1), 179–207.

——— , 1991b, "The Summary Informativeness of Stock Trades: An Econometric Analysis," *The Review of Financial Studies*, 4(3), 571–595.

Hausman, J. A., A. W. Lo, and A. C. MacKinlay, 1992, "An ordered probit analysis of transaction stock prices," *Journal of Financial Economics*, 31(3), 319–379.

Hu, G. X., J. Pan, and J. Wang, 2017, "Early peek advantage? Efficient price discovery with tiered information disclosure," *Journal of Financial Economics*, 126(2), 399–421.

Jiang, W., and C. Yao, 2021, "U Disappeared: How Index Funds Reshape Intraday Stock Market Dynamics," *Working Paper*, The Chinese University of Hong Kong.

Kacperczyk, M., and E. S. Pagnotta, 2019, "Chasing Private Information," *The Review of Financial Studies*, 32(12), 4997–5047.

Kaniel, R., and H. Liu, 2006, "So What Orders Do Informed Traders Use?," *The Journal of Business*, 79(4), 1867–1914.

Kaniel, R., S. Liu, G. Saar, and S. Titman, 2012, "Individual investor trading and return patterns around earnings announcements," *Journal of Finance*, 67(2), 639–680.

Kaniel, R., G. Saar, and S. Titman, 2007, "Individual investor trading and stock returns," *Journal of Finance*, 62(3), 1139–1168.

Kelley, E. K., and P. C. Tetlock, 2013, "How wise are crowds? Insights from retail orders and stock returns," *Journal of Finance*, 68(3), 1229–1265.

Kumar, P., and D. J. Seppi, 1994, "Information and Index Arbitrage," *The Journal of Business*, 67(4), 481–509.

Kyle, A. S., 1985, "Continuous Auctions and Insider Trading," *Econometrica*, 53(6), 1315–1335.

Muravyev, D., 2016, "Order flow and expected option returns," *The Journal of Finance*, 71(2), 673–708.

O'Hara, M., 2015, "High frequency market microstructure," *Journal of Financial Economics*, 116(2), 257–270.

O'Hara, M., C. Yao, and M. Ye, 2014, "What's not there: Odd lots and market data," *The Journal of Finance*, 69(5), 2199–2236.

Philip, R., 2020, "Estimating permanent price impact via machine learning," *Journal of Econometrics*, 215, 414–449.

Riccó, R., B. Rindi, and D. J. Seppi, 2021, "Information, Liquidity, and Dynamic Limit Order Markets," *Working Paper*, Bocconi University.

Roşu, I., 2020, "Liquidity and Information in Limit Order Markets," *Journal of Financial and Quantitative Analysis*, pp. 1792–1839.

Shkilko, A., 2020, "Insider trading under the microscope," *Working Paper*, Wilfrid Laurier University.

Shkilko, A., and K. Sokolov, 2020, "Every Cloud Has a Silver Lining: Fast Trading, Microwave Connectivity and Trading Costs," *Journal of Finance*, forthcoming.

Upson, J., and R. A. Van Ness, 2017, "Multiple markets, algorithmic trading, and market liquidity," *Journal of Financial Markets*, 32(C), 49–68.

van Kervel, V., 2015, "Competition for Order Flow with Fast and Slow Traders," *The Review of Financial Studies*, 28(7), 2094–2127.

Table 1
Summary Statistics

The table reports sample summary statistics. The sample period covers January 4, 2016 to February 29, 2016 for the top 20 stocks listed on the Australian Securities Exchange based on market capitalization on January 4, 2016. We report the total number of trades executed during the sample period (*NTrades*), the average trade price in AUD (*Price*), the average bid-ask spread in AUD (*Spread*), and summary statistics for trade size (*Size*).

	NTrades	Price	Spread	Size (Mean)	Size (Median)
AMC	113,898	13.17	0.011	840	250
AMP	73,224	5.34	0.010	2,759	591
ANZ	248,416	24.00	0.011	932	300
BHP	244,834	15.73	0.011	1,429	476
BXB	107,908	11.06	0.011	919	252
CBA	386,060	76.46	0.018	233	94
CSL	302,654	103.89	0.024	90	42
IAG	72,301	5.23	0.010	2,585	546
MQG	313,652	68.93	0.020	146	59
NAB	251,155	26.39	0.012	692	269
NCM	171,776	14.76	0.011	749	205
ORG	110,871	4.02	0.010	2,724	556
QBE	130,574	10.73	0.011	1,129	294
RIO	267,337	40.83	0.015	282	110
STO	112,709	3.11	0.010	3,576	594
SUN	114,260	11.36	0.011	1,048	283
TLS	104,559	5.43	0.010	6,728	1000
WBC	259,572	30.07	0.012	605	244
WOW	190,625	23.18	0.013	527	200
WPL	242,240	27.04	0.013	368	148

Table 2
VAR and RL Comparison

The table reports price impacts estimated via the VAR and RL models. In Panel A, the VAR is estimated using a +1/-1 trade sign indicator. We report the price impact estimate for each stock and the corresponding 95% confidence interval. In Panel B, the RL model is estimated with one state (market is open) and two actions (to buy or to sell). For each stock, we report the price impact estimate for buy and sell actions. The bold fonts indicate that the magnitude of the RL estimate falls outside the corresponding 95% confidence interval based on the VAR model.

	Panel A: VAR Model			Panel B: RL Model	
	Estimate	Lower 5%	Upper 95%	Sell	Buy
AMC	2.29	1.92	2.65	-2.49	2.53
AMP	3.27	2.72	3.82	-3.41	3.69
ANZ	1.40	1.29	1.51	-1.34	1.47
BHP	1.91	1.74	2.08	-1.95	2.03
BXB	2.40	2.11	2.68	-2.61	2.52
CBA	0.94	0.88	1.00	-0.97	0.97
CSL	1.02	0.94	1.09	-1.06	1.03
IAG	3.14	2.66	3.61	-3.26	3.69
MQG	1.12	1.04	1.20	-1.15	1.18
NAB	1.51	1.39	1.63	-1.52	1.60
NCM	2.36	2.09	2.63	-2.40	2.50
ORG	4.55	3.93	5.17	-4.55	**5.37**
QBE	2.62	2.27	2.97	-2.57	**3.08**
RIO	1.62	1.52	1.72	-1.63	1.69
STO	5.20	4.41	6.00	-4.58	**6.80**
SUN	2.20	1.75	2.64	-2.13	2.60
TLS	2.07	1.07	3.08	-2.16	1.94
WBC	1.45	1.36	1.54	-1.45	1.47
WOW	1.77	1.59	1.95	-1.77	1.92
WPL	1.91	1.76	2.06	-1.93	1.90

Table 3
Price Discovery Contribution of Different Order Types and Sizes

Panel A reports average price impacts for market order submissions (*MO*), limit order submissions (*LO*), and limit order cancellations (*CAN*). We estimate the price impact function using the RL model with one state (limit order book is active) and 6 order types: sell market orders, buy market orders, sell limit order submissions, buy limit order submissions, sell limit order cancellations, and buy limit order cancellations. Each order type is further categorized into 5 size quintiles, resulting in $6 \times 5 = 30$ actions. The estimated model has $1 \times 30 = 30$ state-action pairs. We report the average permanent price impact across 20 sample stocks for each state-action pair. We test for differences between *MO* and *LO* (column 4) and between the absolute values of *LO* and *CAN* (column 5) using a t-test. *** and ** indicate significance at the 1% and 5% levels, respectively. Panel B reports the frequency of occurrence of market and limit order submissions and limit order cancellations in all actions (Occurrence, %). It also reports the total contribution of each action to price discovery as estimated by RL and by VAR as well as the results of a t-test for differences between the RL and VAR estimates.

	Order submissions		Cancellations		
	Market (MO)	Limit (LO)	(CAN)	MO - LO	LO - CAN
	[1]	[2]	[3]	[4]	[5]
Panel A: Permanent price impacts					
Sell Q5	-3.54	-0.83	0.13	***	***
Sell Q4	-2.59	-0.46	0.16	***	***
Sell Q3	-2.33	-0.33	0.22	***	**
Sell Q2	-2.07	-0.31	0.54	***	
Sell Q1	-1.59	-0.30	0.08	***	***
Buy Q1	1.53	0.32	-0.13	***	***
Buy Q2	1.97	0.36	-0.53	***	
Buy Q3	2.24	0.36	-0.24	***	**
Buy Q4	2.50	0.49	-0.21	***	***
Buy Q5	3.43	0.82	-0.18	***	***
Panel B: Price discovery contribution					
Occurrence, %	16.65	53.01	30.34		
RL contribution	53.83	35.65	10.51		
VAR contribution	49.71	37.96	12.32		
RL-VAR (t-stat)	(-1.60)	(1.26)	(1.27)		

Table 4
Price Impacts Conditional on the State of the Limit Order Book

The table reports the average price impacts for market orders, limit order submissions and cancellations conditional on the order book imbalance. We estimate the price impact function using the RL model with 10 states based on the book imbalance at the time of the action, which are formed into 10 deciles, and on 6 order types: sell market orders, buy market orders, sell limit order submissions, buy limit order submissions, sell limit order cancellations and buy limit order cancellations. Each order type is further categorized into 5 size quintiles, resulting in $6 \times 5 = 30$ overall actions. The estimated model has $10 \times 30 = 300$ state-action pairs. We compute the average price impact across the sample stocks for each state-action pair and report the average price impact for the most negative (*D1*) and most positive (*D10*) depth imbalance deciles. *** indicates statistical significance at the 1% level between these deciles.

	Market Orders			Limit Orders			Cancellations		
	D10	D1	D10 - D1	D10	D1	D10 - D1	D10	D1	D10 - D1
	[1]	[2]	[3]	[4]	[5]	[6]	[7]	[8]	[9]
Sell Q5	1.97	-9.92	***	4.32	-6.31	***	5.70	-4.86	***
Sell Q4	3.43	-9.21	***	4.90	-5.85	***	5.72	-4.94	***
Sell Q3	3.67	-8.94	***	5.07	-5.70	***	5.76	-4.91	***
Sell Q2	4.08	-9.11	***	5.10	-5.70	***	5.93	-4.85	***
Sell Q1	4.48	-8.55	***	5.19	-5.66	***	5.69	-5.18	***
Buy Q1	8.06	-4.17	***	5.81	-5.08	***	5.36	-5.53	***
Buy Q2	8.44	-3.81	***	5.82	-4.95	***	4.89	-5.78	***
Buy Q3	8.62	-3.26	***	5.81	-4.90	***	5.02	-5.63	***
Buy Q4	9.03	-3.23	***	6.01	-4.73	***	5.09	-5.54	***
Buy Q5	9.73	-1.85	***	6.38	-4.20	***	5.02	-5.54	***

Table 5
Statistical Significance of Additional Conditioning Variables

For each conditioning variable (i.e., price history, bid-ask spread, order arrival frequency, and time of day), the table reports statistical significance of differences in permanent price impacts between the states that are captured by the corresponding figures (Figures 4 through 7). *MO* denotes market orders, *LO* – limit order submissions, and *CAN* – limit order cancellations. For *Price History*, Columns 1 through 3 report on the difference in price impacts between the most negative (*Down)* and positive (*Up*) return deciles. Columns 4 through 6 report on the difference between the narrow and wide bid-ask spreads. Columns 7 and 8 report on the difference between the fast and slow order arrivals for market orders when the spreads are narrow and for limit orders when the spreads are wide. We omit this analysis for limit order cancellations, because Bhattacharya and Saar (2021) do not make predictions for cancellations. Finally, columns 10 through 12 report on the differences between price impacts during the early morning and late afternoon. To obtain the statistics, we first compute the differences for each stock and then compute the average difference across all stocks. Asterisks ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels based on a t-test.

	Price History (Down-Up)			Bid-Ask Spread (Narrow-Wide)			Order Arriv. Freq. (Fast-Slow)			Time of Day (AM-PM)		
	MO	LO	CAN	MO	LO	CAN	MO	LO	CAN	MO	LO	CAN
	[1]	[2]	[3]	[4]	[5]	[6]	[7]	[8]	[9]	[10]	[11]	[12]
Sell Q5	***	***	***		***	**		***	-	***	***	***
Sell Q4	***	***	***	**	***	*	**	***	-	***	***	***
Sell Q3	***	***	***	*	***		***	***	-	***	***	**
Sell Q2	***	***	***	***	***	*	***	***	-	***	***	**
Sell Q1	***	***	***	***	***	**	***	***	-	***	***	***
Buy Q1	***	***	***	**	***	**	***	***	-	***	***	***
Buy Q2	***	***	***	**	***		***	***	-	***	***	
Buy Q3	***	***	***	*	***	***	***	***	-	***	***	
Buy Q4	***	***	***	*	***		***	***	-	***	***	***
Buy Q5	***	***	***		***		**	***	-	***	***	***

Table 6
MDA-based Contribution to Price Discovery

This table reports the Mean Decreased Accuracy (MDA) estimated as follows for each price discovery conduit k:

$$MDA^k = \sum_{(s,a)=1}^{(S,A)} \left(\frac{|(\hat{PI}_{s,a} - \tilde{PI}^k_{s,a})|}{\hat{PI}_{s,a}} \right) / (S,A),$$

where (S,A) is the total number of state-action pairs, $\hat{PI}_{s,a}$ is the estimated price impact for each state-action pair (s,a) based on the ML model with 90,000 state-action pairs, and $\tilde{PI}^k_{s,a}$, is the estimate created by randomly changing the conduit k. For example, to test the relative importance of the state of the book, let us assume that the true price impacts in state-action pairs (*D1, c*) and (*D4, c*) are 1 bp and 2 bps, where *D* is depth imbalance, and *c* indicates that all other conduits in both pairs are the same. To randomize the depth imbalance variable, the observation (*D1, c*) is flipped to have the value of (*D4, c*). Thus, the $\tilde{PI}^k_{s,a}$ value for (*D1, c*) is assigned as 2 bps instead of 1 bp. We repeat this process 500 times for each conduit k and report the mean and standard deviation of the MDA for each variable. In column 3, we report the rank of the variable based on its relative importance to price discovery. Column 4 reports the p-values from the t-tests of differences between the MDAs of two consecutive conduits. For instance, the first reported p-value of 0.00 relates to the difference between the MDA of orders and the MDA of the state of the limit order book.

	Mean	Std. Dev.	Rank	p-Value
	[1]	[2]	[3]	[4]
State of the Book	10.38	3.07	1	-
Order	9.25	2.74	2	(0.00)
Price History	8.79	3.27	3	(0.06)
Order Size	6.59	2.56	4	(0.00)
Order Arrival Frequency	4.57	1.43	5	(0.00)
Bid-Ask Spread	4.34	1.46	6	(0.06)
Time of Day	2.86	0.31	7	(0.00)

Table 7
An Application: HFTs v. Retail Investors

The table compares the informativeness of trades submitted by high-frequency traders (HFTs) and retail investors. Using brokerage identifiers, we divide sample brokerages into three categories: (i) those, whose clientele is predominantly HFTs, (ii) those, whose clientele is predominantly retail traders, and (iii) others. In Panel A, we compare the average price impacts of trades from the first two brokerage categories. In Panel B, using the same two brokerage categories, we estimate the following regression model for each stock i and each state-action pair s,a:

$$\hat{PI}^{HFT}_{i,s,a} = \alpha + \beta \times \hat{PI}^{Retail}_{i,s,a} + \varepsilon_{i,s,a},$$

where $\hat{PI}^{HFT}_{i,s,a}$ is the estimated price impact of trades originating from HFT brokerages, and $\hat{PI}^{Retail}_{i,s,a}$ is the estimated price impact of trades originating from retail brokerages. The β coefficient measures the relative informedness of HFT trades compared to Retail trades. If $\beta = 1$, then HFTs and retail traders have the same price impact functions, implying that they are equally informed. We fist estimate a baseline model that compares simple price impact functions for the two account types. The simple price impact functions are comparable to the averages in Panel A. The baseline estimate is followed by a set of models that control for the previously identified price discovery conduits in the information set of both HFTs and retail investors. We add the conduits one by one, in the order of importance identified in the MDA analysis. The t-statistics are for the hypothesis that the β coefficient is equal to 1.

Panel A: Average price impacts		
HFTs	3.03	
Retail traders	1.88	
Panel B: Regression results		
	Coefficient	t-Statistic
Baseline	1.58	15.45
+State of the Book	1.01	2.54
+Price History	0.97	-7.54
+Order Arrival Frequency	0.90	-37.72
+Bid-Ask Spread	0.88	-49.05
+Time of Day (the full model)	0.87	-62.08

Figure 1. VAR impulse response and RL convergence functions

The figure provides examples of price impact estimates obtained using the traditional VAR and the RL. Panel A reports the VAR impulse response function (IRF, dashed line) and the RL convergence function (solid line) for a typical buy market order. To plot the RL convergence function, we use the price impact estimates obtained after each iteration of the learning rule. Specifically, when solving equation (1), the price impact estimate *Q(s,a)* updates to a new value after each iteration until convergence. By construction, after the initial iteration the estimate contains only the immediate price impact, while upon convergence it contains both the immediate and all the subsequent impacts. Thus, as we observe the progression of the estimates between the first iteration and convergence, we see how the price impact estimate evolves from the immediate (and possibly temporary) to the long-term permanent impact. Panel B demonstrates the flexibility of RL convergence functions using an example of a temporary price impact that is in the direction opposite the permanent price impact. In this panel, we use RL to plot the price impact of a buy order that executes when the order book has a negative imbalance. Even though the order has a temporarily positive price impact, this impact soon reverses followed by a permanent price impact that is negative.

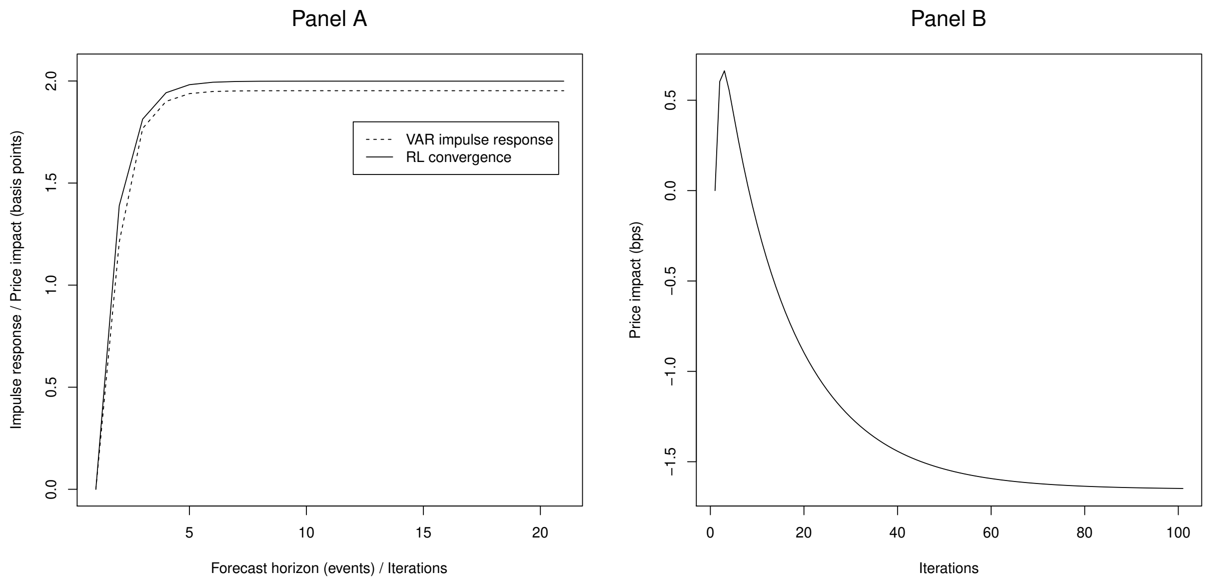

Figure 2. Price impact functions for different order types and sizes

The figure plots price impact functions for market orders (Panel A), limit order submissions (Panel B), and limit order cancellations (Panel C). We estimate price impact functions using the RL model with one state (market open) and 6 order types: sell market orders, buy market orders, sell limit order submissions, buy limit order submissions, sell limit order cancellations, and buy limit order cancellations. Each order type is further categorized into 5 size quintiles, resulting in $6 \times 5 = 30$ overall actions. The estimated model has $1 \times 30 = 30$ state-action pairs. We report the average price impact across 20 sample stocks for each state-action pair.

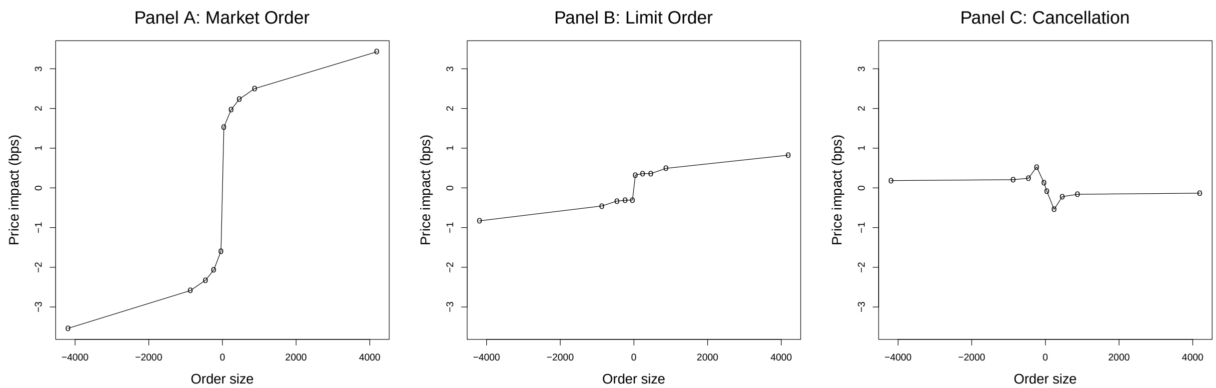

Figure 3. Price impacts conditional on order book imbalance

The figure plots the price impact functions for market orders (Panel A), limit order submissions (Panel B), and limit order cancellations (Panel C) based on the book imbalance at the time of the action. We estimate the price impact function using the RL model with 10 imbalance states at the time of the action, which are formed into 10 deciles, and 6 actions: sell market orders, buy market orders, sell limit order submissions, buy limit order submissions, sell limit order cancellations, and buy limit order cancellations. Each action is further categorized into 5 size quintiles, resulting in $6 \times 5 = 30$ overall actions. The estimated model has $10 \times 30 = 300$ state-action pairs. We compute the average price impact across the sample stocks for each state-action pair and report the average price impact for the most negative (*D1*, dotted line with cross markers) and the most positive (*D10*, dashed line with circle markers) imbalance deciles. The solid lines represent the base case from Figure 2, and the horizontal lines represent the unconditional imbalance effects.

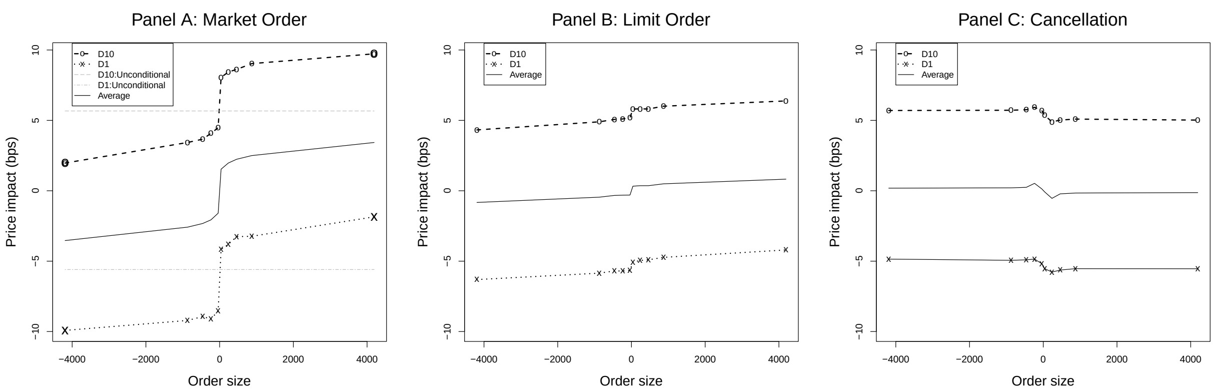

Figure 4. Price impacts conditional on price history

The figure plots the price impact functions for market orders (Panel A), limit order submissions (Panel B), and limit order cancellations (Panel C) based on price history. We estimate the price impact function using the RL model with 10 states based on the midpoint returns in the 10 actions prior to the current action, which are formed into deciles from most negative returns to most positive returns, and 6 order types: sell market orders, buy market orders, sell limit order submissions, buy limit order submissions, sell limit order cancellations, and buy limit order cancellations. Each order type is categorized into 5 size quintiles, resulting in $6 \times 5 = 30$ overall actions. The estimated model has $10 \times 30 = 300$ state-action pairs. We report the average price impact across the sample stocks for each state-action pair. To save space, we report the average price impacts for the most negative (*Ret1*) and most positive (*Ret10*) return deciles.

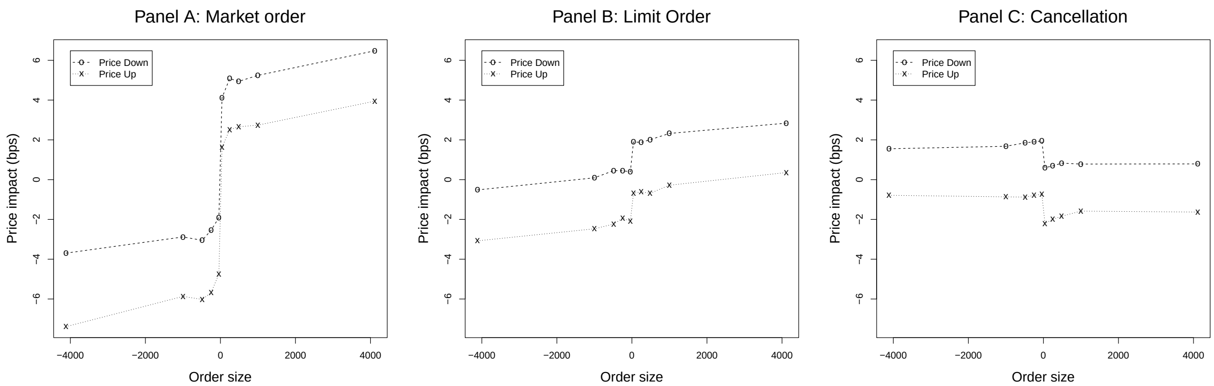

Figure 5. Price impacts conditional on the bid-ask spread

The figure plots the price impact functions for market orders (Panel A), limit order submissions (Panel B), and limit order cancellations (Panel C) based on the size of the spread at the time of the action. We estimate the price impact function using the RL model with 2 states defined by the size of the bid-ask spread. *Narrow* (*Wide*) contains observations when the spread equals (is greater than) AUD 0.01. The RL model contains 6 order types: sell market orders, buy market orders, sell limit order submissions, buy limit order submissions, sell limit order cancellations, and buy limit order cancellations, with each order type further categorized into 5 size quintiles, resulting in $6 \times 5 = 30$ actions. The estimated model has $2 \times 30 = 60$ state-action pairs. We compute the average price impact across the sample stocks for each state-action pair and report the average price impacts for the *narrow* and *wide* states.

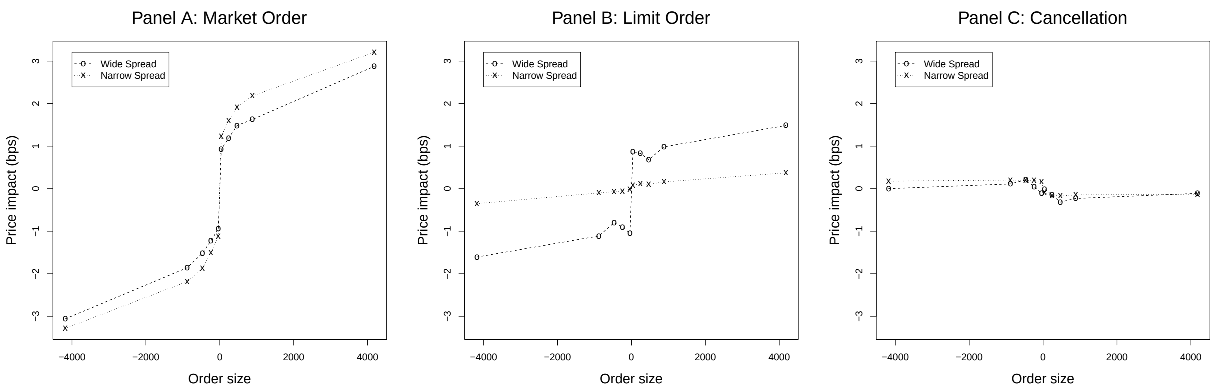

Figure 6. Price impacts conditional on order arrival frequency

The figure plots the price impact functions for market and limit orders conditional on the spread at the time of order submission and on order arrival frequency. We estimate the price impact function using the RL model with 2 states defined by the size of the spread and 5 states defined by the order arrival frequency, giving a total of $2 \times 5 = 10$ states. *Narrow* (*Wide*) contains observations when the spread equals (is greater than) AUD 0.01. The order arrival frequency states are based on the time between the current action and the previous action of the same type, which are formed into 5 quintiles. The RL model contains 6 order types: sell market orders, buy market orders, sell limit order submissions, buy limit order submissions, sell limit order cancellations, and buy limit order cancellations, with each order type further categorized into 5 size quintiles, resulting in $6 \times 5 = 30$ overall actions. The estimated model has $10 \times 30 = 300$ state-action pairs. We compute the average price impact across the sample stocks for each state-action pair and report the average price impacts for the *Narrow* and *Wide* subsamples. and for the highest (*FastArrival*) and lowest (*SlowArrival*) order arrival frequency quintiles.

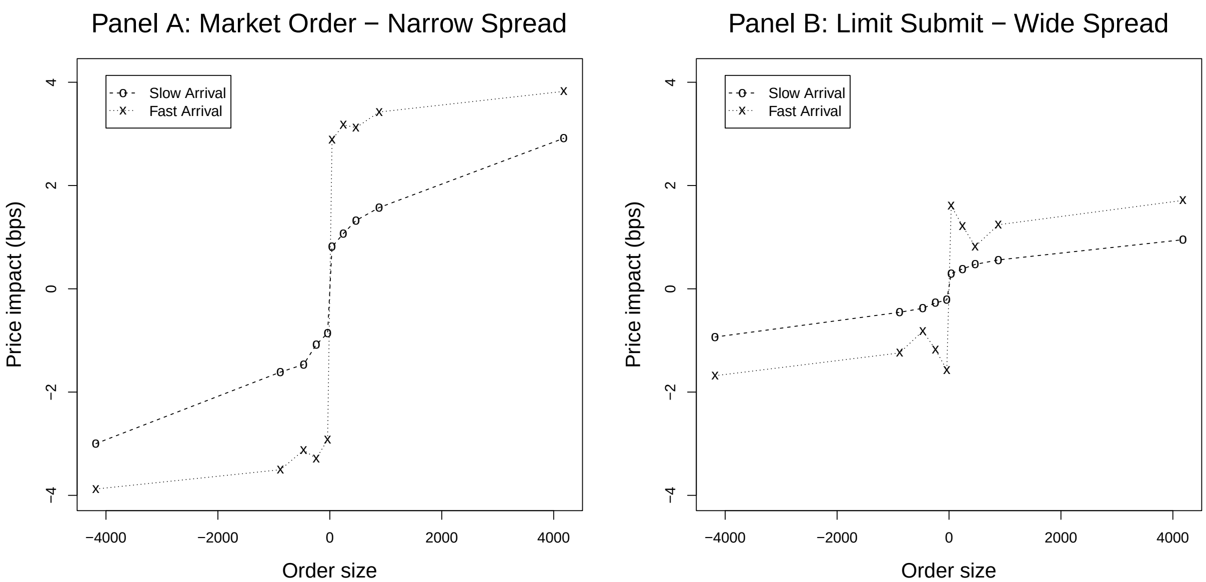

Figure 7. Price impact functions conditional on time of day

The figure plots the price impact functions for market orders (Panel A), limit order submissions (Panel B), and limit order cancellations (Panel C) based on the time of day. We estimate the price impact function using the RL model with 3 states based on the time of the action (*AM*, *Midday*, *PM*) and 6 order types: sell market orders, buy market orders, sell limit order submissions, buy limit order submissions, sell limit order cancellations, and buy limit order cancellations. Each order type is categorized into 5 size quintiles, resulting in $6 \times 5 = 30$ overall actions. The estimated model has $3 \times 30 = 90$ state-action pairs. We report the average price impact over the 20 stocks for each state-action pair.