The Conduits of Price Discovery: A Machine Learning Approach*

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

May 31, 2020

Abstract. Theory models suggest that market conditions should have substantial effects on order submission strategies and price discovery. Empirical analyses of such conditional effects are methodologically challenging and therefore uncommon. We bypass these challenges using a machine learning technique that allows for multiple conditioning variables in the presence of non-linearities. The analysis confirms theory predictions in that price discovery is affected, and often dominated, by such conditions as the state of the limit order book and prior order history. Furthermore, the technique allows us to rank the importance of conduits through which information flows into prices. The current state of the limit order book stands out as the primary conduit.

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

JEL: G14; G15

*We thank Jonathan Brogaard, Michael Brolley, David Cimon, Andreas Park, and Konstantin Sokolov for valuable comments. Shkilko acknowledges support of the Social Sciences and Humanities Council of Canada and the Canada Research Chairs program for support. Some of this project was completed while Shkilko held a fractional appointment at the University of Sydney. The most current version of this 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 many 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 generally followed a similar pattern, with early studies focusing on trade-driven price discovery and more recent work emphasizing a significant role of limit orders.

In addition to modeling the direct impact 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 (2019) show that the current state of the limit order book may play an important role in price discovery, affecting the informativeness of both market and limit orders. In turn, Riccó, Rindi, and Seppi (2020) predict that price discovery may 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 entirely 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 (2019) argue that machine learning techniques are well-suited to shed light on market structure issues that traditional methodologies struggle to address. Philip (2019) shows that RL is particularly useful in microstructure applications characterized by non-linearities and tends to outperform conventional methods. To put it plainly, RL is a flexible technique able to uncover relations among multiple theory-based variables (e.g., price impacts, the state of the book, and prior order history) by allowing a computer to learn what these relations are instead of being told what they might be. Section 2 describes RL estimation in detail.

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 and on prior order history. The state of the book explains a sizable portion of the permanent price impact of newly submitted market and limit orders, nearly four times as much as order submission itself. In turn, prior order history is about two-thirds as important as an order submission. Further, consistent with Riccó, Rindi, and Seppi (2020), the information content of arriving orders is often 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 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 the most important conduit for information flow into prices. It determines the lion's share of price discovery, 60.6%. Individual order submissions and cancellations follow as a relatively distant second, at 13.5%, with individual market orders dominating the effect. Prior order history contributes 10.9%, and order size and time effects add, respectively, 8.6% and 6.4%. To our knowledge, these results represent the first attempt to determine the pecking order of price discovery conduits.

In addition to ranking the means by which information flows into prices, our analyses contribute to a line of research into trading behavior of informed agents. A number of theoretical models seek to understand this behavior, often delivering dual 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. Akey, Grégoire, and Martineau (2019), and Shkilko (2020) show that some informed agents move prices quickly, while Collin-Dufresne and Fos (2015), Kacperczyk and Pagnotta (2019), and Garriott and Riordan (2020) 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 then show that an informed agent willing to exploit stochastic noise to disguise her trading may stay undetected by the market. Our data contain strong evidence of such behavior; the most informed agents tend to time periods of returns opposite the direction of their own trading.

Theoretical models allow price discovery 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. It is quite possible that this pattern has changed in recent years due to the proliferation of indexing institutions that tend to rebalance portfolios at the end of the day diluting informed trading (Bogousslavsky and Muravyev (2020)). Our analyses support this possibility; 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 (2020), who suggest that as the probability of execution declines through the day, informed trading should shift away from limit orders.

Our contribution to the finance literature is three-fold. First, we empirically examine recent theory predictions that the state of the limit order book, prior order history, and time effects are important characteristics, on which price discovery should be conditioned. Second, we describe a machine learning technique for price discovery analyses that is considerably more flexible than the conventional methodologies, especially when it comes to multiple conditioning variables and dealing with sizeable non-linearities in the data. Finally, we provide the first ranking of price discovery conduits, that is, the channels through which information flows into prices.

Related literature. A number of theory models describe the price discovery process. 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 (2019)), Riccó, Rindi, and Seppi (2020)). 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 a significant role of limit orders (Bloomfield, O'Hara, and Saar (2005), 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 and ranking the various channels through which information flows into prices.

Our analyses also contribute to research into trading behavior of informed agents. Kyle (1985) suggests that a privately informed agent will trade slowly attempting to avoid detection and thus moving the price gradually. In turn, Collin-Dufresne and Fos (2016) argue that the agent may be willing to time stochastic market conditions to avoid detection even better. We find evidence consistent with both behaviors. The most informed transactions in our data are contrarian; they time periods of returns opposite their own direction of trading. Yet there is also evidence of gradual momentum-like patterns characteristic of traders described in Kyle (1985).

Finally, this study contributes to the literature on intraday patterns. Early research (e.g., Wood, McInish, and Ord (1985), Chung, Van Ness, and Van Ness (1999)) shows that volatility and spreads follow a U-shaped pattern, consistent with the possibility that information flows into prices in the morning and in the afternoon. Notably, Upson and Van Ness (2017) report that the spread pattern has changed in recent years; spreads are wide in the morning and decline through the day. It is possible, and perhaps likely, that the pattern change is due to changes in the timing of price discovery. We shed light on this possibility by showing that price discovery through market orders is indeed most pronounced in the morning and then gradually declines. Notably, the information content of limit orders decays much more rapidly through the day compared to market orders. This result is consistent with theory predictions of Riccó, Rindi, and Seppi (2020), who note that as the day draws to a close, the probability of limit order executions declines, and informed traders switch to more aggressive order submission strategies.

2. Reinforcement Learning model

In this section, we describe the components and estimation of the RL model first proposed by Philip (2019).

2.1 Model parameters

Estimating price impacts via the RL framework 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 simple application of the model, there is only one state: the market is open. As such, this specification is similar to a typical VAR model of price impact. In contrast, a more complex RL setup could allow for different states of the limit order book. For example, one can specify five states to describe the state of the book at the time an action takes place: a large negative depth imbalance, a small negative depth imbalance, no depth imbalance, a small positive depth imbalance, and a large positive depth imbalance.

In turn, actions a are the choices available to the market participant. In a simple RL model, the agent chooses between two possible actions: to buy or to sell. A more complicated model may allow for several actions, including limit order submissions, cancellations, market orders, and also a choice of order size.

Once states and actions are defined, they are combined to form state-action pairs (s,a). As an illustrative example, a model with 3 states (negative depth imbalance, no depth imbalance, positive depth imbalance) and 2 actions (to buy or to sell) results in $3 \times 2 = 6$ state-action pairs. Each observation in the dataset is then classified into one of the 6 state-action pair categories. The number of observations may vary across state-action pairs. For example, there is likely to be more observations in the (negative depth imbalance, sell) pair than in the (negative depth imbalance, buy) pair. This imbalance does not impair RL estimation.

2.2 Estimating the RL model

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

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

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

In equation (1), estimating the model requires three inputs: immediate price impact R, transition probability T, and the discount factor γ. The immediate price impact refers to changes in the midpoint price as a result of an action in a given state. We require an immediate impact for each state-action pair and thus, our illustrative example requires 6 immediate impacts. To compute the immediate impact for each state action pair, we first compute the immediate impact for every observation in the data as the change in log midpoint price from one observation to the next.[1]

[1] Observations may be based on all trades as in our illustrative example but may also include all order submissions and cancellations.

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 of the state-action pair (negative depth imbalance, sell), we take all sale transactions at the time of a negative depth imbalance and compute the average change in the log midpoint.

In turn, the transition probability is the probability of observing the future state-action pair $(s', \bar{a})$ after observing the current state-action pair (s, a). Because there are $(s \times a)^2$ possible transitions from (s, a) to $(s', \bar{a})$, our example requires calculating $6 \times 6 = 36$ transition probabilities. For example, to calculate the probability of transitioning from the current state-action pair (negative depth imbalance, sell) to the future state-action pair (positive depth imbalance, buy), we determine the number of times we observe a sale transaction when a negative depth imbalance exists in the order book, followed by a buy transaction when a positive depth imbalance exists, and express this number as a fraction of all observations in the (negative depth imbalance, sell) state-action pair.

Finally, γ is a discount factor representing the time value of money and lies between 0 and 1.[2] 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, $R(s, a)$, and the permanent price impact of all subsequent actions triggered by the initial action captured by the $\sum_{s' \in S} \sum_{a' \in A} T((s, a), (s', \overline{a})) Q(s', \overline{a})$ term. The permanent price impact of all subsequent actions depends on the probability of observing subsequent actions $\bar{a}$ while in s' and the permanent price impact of these subsequent actions.

Because both the right-hand side (RHS) and the left-hand side (LHS) of equation (1) contain a $Q(s, a)$ term, we can obtain estimates of $Q(s, a)$ using an iterative learning rule known as Q-learning. Specifically, for the first iteration, we initialize all $Q(s, a)$ estimates on the RHS of

[2] Precisely, $0 <= \gamma < 1$, as γ must be less than 1 to ensure convergence.

equation (1) to zero and solve for $Q(s,a)$ on the LHS of equation (1).[3] 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 of equation (1) converges, which implies that $Q(s,a)$ values do not change with subsequent iterations.[4]

3. Data

We use full order book data for the Australian Securities Exchange (ASX) – Australia's main equity market. The data come from the SIRCA database for a January 4, 2016 to February 29, 2016 sample period. The sample includes 20 largest stocks listed on the ASX. The SIRCA database contains every trade, order submission, cancellation and amendment allowing us to fully reconstruct the order book.

The ASX data are particularly suitable for our purposes for two reasons. First, the Australian market remains largely consolidated during the sample period, allowing for a correct chronological ordering of trades. Such ordering is notably more difficult when a researcher uses data from the U.S., where equity markets are geographically dispersed, and transmission latencies may lead to incorrect message sequencing in consolidated data (e.g., Easley, López de Prado, and O'Hara (2012); Chakrabarty, Jain, Shkilko, and Sokolov (2020)). Second, the ASX executes almost 90% of Australian exchange volume. As such, activities on the competing exchanges, which are not in the ASX dataset, are less of a concern compared to, say, using Nasdaq ITCH data that cover a smaller share of the U.S. market.

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.

[3]Thus, in the first iteration, the price impact of state-action pair (s,a) is equal to the immediate reward of state-action pair (s,a) as equation (1) collapses to $Q(s,a) = R(s,a)$.

[4]For an illustrative example of the iterative learning rule, see Philip (2019).

Upson, McInish, and Johnson (2018) show that empirical results may be biased without such consolidation. 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 approximately 72,000 trades over the sample period to the Commonwealth Bank of Australia (CBA) with over 386,000 trades. The sample stocks also cover a wide range of stock prices, from AUD 3.11 to 103.89, and a range of average trade sizes, from 90 to 6,728 shares.

[Table 1]

4. RL and VAR: a comparison

When measuring price discovery, the finance literature frequently relies on vector autoregression (VAR) techniques developed by Hasbrouck (1991a), Hasbrouck (1991b), Hasbrouck (1995). As such, it is important to show that the price impact estimates from the RL model are comparable to the VAR estimates. To do so, we estimate a 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 (market open) and two actions (buy or sell). These two actions are analogous to the $+/-1$ trade sign indicator in the VAR model.

We report the price impact estimates in Table 2, with estimates based on the RL model in Panel A and those based on the VAR model in Panel B.[5] The average price impact estimated via the traditional VAR model ranges from 0.94 (CBA) to 5.20 (STO) basis points. 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) basis points.

Panel B presents the corresponding price impact estimates from the RL model. In contrast to the VAR model, 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 a price impact of -2.49 basis points for sells and 2.53 points for buys. Observing across all buy and sell estimates, the magnitude of price impact estimates 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 price impact estimate for STO.

To further compare between the 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 estimates: 37 out of 40 buy and sell price impact estimates based on the RL model fall within the 95% confidence interval of the VAR estimates. Importantly, the correlation between the VAR estimates and the average of the RL buy and sell estimates is over 0.99.

[Insert Table 2]

[5] For the VAR permanent price impact estimates, we report returns of the cumulative impulse response function after 30 events.

5. Non-linearities and expansion of 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 (2019) shows that ignoring the nonlinear effects of trade sizes 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 the previous section, we focus 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 so, we expand the RL analysis to include trades and limit orders. For consistency of exposition, we refer to trades as *market orders*. When it comes to limit orders, we examine all top of the book order submissions and cancellations.

We estimate the RL model using one state (market open) 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.[6] In addition, for each order category we use 5 size quintiles, resulting in $6 \times 5 = 30$ unique actions. The final model has $1 \times 30 = 30$ state-action pairs. To ensure a fair comparison, we use the same size cutoffs for each action.[7]

Figure 1 plots the price impact estimates for market orders (Panel A), limit order submissions (Panel B) and limit order cancellations (Panel C). Each point on the graph represents the average price impact across the sample stocks for each action. The x-axis plots the largest value for each size quintile, with negative values representing sell actions and positive values representing buy actions.

[Insert Figure 1]

[6] A buy (sell) limit order cancellation is a cancellation from the bid (ask) side of the order book.
[7] The cutoffs based on size quintiles are 39, 236, 460, 873 and 4,192 shares.

Panel A contains an *S*-shaped pattern indicating that price impact increases with trade size, albeit at a decreasing rate.[8] This finding is consistent with the economic intuition that large trades should cause greater price impacts than small trades. Perhaps equally importantly, increasing the trade size from 999 shares to 1,000 shares is associated with a smaller increase in price impact than increasing the trade size from 1 share to 2 shares.

Panel B reports a similar *S*-shaped pattern for limit order submissions, although the magnitude of the *S* function is smaller compared to trades. This 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 reverse *S*-shape, consistent with the notion that a sell (buy) order cancellation signals a future price increase (decrease). We test these relations formally in Table 3.

[Insert Table 3]

Table 3 reports the average price impact estimate for each (order category, size) action. To simplify exposition, we reverse the signs of buys and sells for limit order cancellations such that a sell (buy) cancellation is interpreted similarly to a buy (sell) submission. We also test for differences in price impact estimates between market orders and limit order submissions (Column 4) and limit order submissions and cancellations (Column 5). The data confirm that, for a given size, market orders have a larger price impact than limit order submissions or cancellations, which is consistent with the results in Figure 1. More specifically, price impacts of individual market orders are 4.2 to 7 times larger than those of individual limit order submissions. In the meantime, price impacts of individual limit order submissions are 0.6 to 6.4 times larger than those of cancellations. These differences are statistically significant for 90% of (order category, size) actions.

[8]Estimating price impact using a VAR model with +1 for buys and -1 for sells implies that price impact follows a step function. That is, the VAR assumes that buys and sells of different sizes have equal price impacts, although with opposite signs.

The data show 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? We note that limit order submissions are often triggered by new information. In the meantime, cancellations often depend on queuing. That is, if a trader learns that her limit order is too far down in the queue (such knowledge is typically gained by most professional traders after submission), she may choose to convert it to a marketable order. The resulting cancellation is less informed than the original submission because it is driven by queue positioning rather than exogenous information.

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 from other order types. To examine this possibility, we investigate the contribution to total price discovery attributed to market and limit orders, after adjusting for the frequency of each order type. Table 4 reports that market orders, limit order submissions and cancellations respectively contribute 53.8%, 35.7%, and 10.5% to total price discovery.

At first glance, it may appear that these results contradict Brogaard, Hendershott and Riordan (2019, BHR), who analyze Canadian data and report that limit orders lead price discovery. A closer look however suggests that our results are complementary. First, Australian trading is less fragmented than trading in Canada, and therefore cross-market market making and arbitrage activities, which typically involve a larger number of limit order submissions and cancellations (e.g., van Kervel (2015); Shkilko and Sokolov (2020)), are less prevalent on the ASX. Specifically, Table 4 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 Toronto Stock Exchange (TSX). Consistently, there is less price discovery from limit orders on the ASX, relative to the TSX. Second, BHR show that it is mainly the HFT price discovery that occurs via limit orders, while non-HFT price discovery occurs through market orders. Our sample includes both trader categories, and as such delivers results for an average trader.

[Insert Table 4]

6. Conditioning price discovery

The results in the previous section show that there is a concave relation between trade size and price impact. It is natural to now ask what other factors contribute to price discovery. To answer this question, we follow the theoretical predictions of Goettler, Parlour, and Rajan (2009), Roşu (2019), and Riccó, Rindi, and Seppi (2020) and extend the RL model to include three additional sources of information: current state of the limit order book (Section 6.1), prior order history (Section 6.2), and time of day (Section 6.3). For these analyses, we estimate the RL model with the same 30 actions described in Section 5 while changing the states to reflect the information source. For instance, we distinguish the state with a large negative limit order book imbalance from that with a large positive imbalance. Finally, in Section 6.4, we report the proportional contribution of each information source to price discovery, effectively answering the question: Which information conduits are most important in determining future price movements?

6.1 Current state of the limit order book

For this investigation, we estimate the RL model using 10 states (depth imbalance deciles at the time of each action) and 30 actions, as described in Section 5, resulting in a model with $10 \times 30 = 300$ state-action pairs. We calculate depth imbalance (DI) at the time of each market order submission, limit order submission, and limit order cancellation as follows:

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

where $Vol_{Bid,i}$ is the total bid volume, and $Vol_{Ask,i}$ is the total ask volume at price level i, for a total of ten price levels (five offer and five bid levels). We then divide *DI*s into ten deciles and assign each order submission and cancellation to one of the deciles, from the most negative imbalance to the most positive one. Figure 2 plots the price impact estimates for the most negative (*D1*, cross markers) and most positive (*D10*, circle markers) deciles for market orders (Panel A), limit order submissions (Panel B), and limit order cancellations (Panel C).

[Insert Figure 2]

We make two observations from Figure 2. First, similar to Figure 1, and for both positive and negative imbalances, we observe an *S*-shaped price impact function for market orders and limit order submissions and a reverse *S*-shaped function for cancellations. Second, and more importantly, the difference in price impact between the start (*Sell Q5*) and the end (*Buy Q5*) of the *S*-shape function, which represents the contribution of order size, is much smaller than the difference between the *D1* and *D10* functions, which represents the contribution of depth imbalance. For example, when a large positive depth imbalance exists (*D10*), a large market buy order (*Buy Q5*) has a price impact of 9.73 bps, while a large market sell order (*Sell Q5*) has a price impact of 1.97 bps. The difference of 7.76 bps represents the contribution of order size to price impact while holding the depth imbalance constant. In contrast, when we hold the action constant and vary the state, the contribution of depth imbalance is much larger than that of order size. For example, the price impact for large market buy orders in state *D10* is 9.73 bps, while the price impact of a similarly sized market buy order 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 2 reveals a curious regularity. Namely, the information content of arriving orders is often opposite their direction and aggressiveness, and this relation depends crucially on the state of the book. For instance, large buy orders, both market and limit, have negative price impacts when the order book imbalance is negative and vice versa. More specifically, a large market order to buy has a -1.85 bps price impact when the book imbalance is significantly negative, while a large market order to sell has a 1.97 bps price impact when the book imbalance is significantly positive. This result is consistent with the theoretical predictions of Riccó, Rindi, and Seppi (2020) and, to our knowledge, is the first empirical confirmation of these predictions.

To highlight these differences further, Table 5 reports the difference in price impact estimates between *D10* and *D1* (columns 3, 6, and 9) and between *Buy Q5* and *Sell Q5* (the bottom row). The results confirm that depth imbalance plays a considerably more important role in price discovery than order size. For market orders, the average difference between *D10* and *D1* is 12.34 bps (column 3), while the average difference between *Buy Q5* and *Sell Q5* is 7.76 bps for *D10* and 8.07 bps for *D1* (the bottom row). That is, for market orders depth imbalance contributes 56% more to price discovery than order size (the difference between 12.34 bps and $(7.76+8.07)/2 = 7.92$ bps).

[Insert Table 5]

For limit order submissions and cancellations, the difference is even more prominent. For submissions, depth imbalance contributes five times more to price discovery than order size. Specifically, the average contribution of order size is $(2.06+2.11)/2 = 2.09$ bps, and the average contribution of depth imbalance is 10.74 bps. By way of similar calculations, the contribution of depth imbalance for cancellations is 15 times greater than that of order size.

The effects of the state of the limit order book discussed thus far are conditional on an order submission or cancellation. Is it possible to gauge the unconditional effect? 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 effect of the book may be approximated. For this, 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 5. Averaging these gives the unconditional price impact of 11.3 bps as illustrated by horizontal grey lines in Figure 2. To clarify, this result suggests that when the limit order book is in state *D1* or *D10*, the future price should be expected to fall (rise) by 5.6 bps.

Overall, the results in this section highlight the importance of the current state of the order book in influencing prices. The literature often assumes that market orders come from informed investors, whose information may be short-lived or highly valuable and therefore requires 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 (2020), 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.

6.2 Prior order 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 (2020) 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 a similar RL model with 10 states and 30 actions as before. This time however, the states are determined based on midpoint returns 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 positive returns, respectively. As in the previous analyses, to ensure consistency across all actions, we use the same cutoff values for the return deciles.

We plot the price impact estimates for the most negative (*Ret1*) and most positive (*Ret10*) return deciles in Figure 3. For market orders in Panel A, price reversal patterns appear to play a non-trivial role. Specifically, large buys that follow large negative returns (circle markers) have greater price impacts than large buys that follow large positive returns (cross markers). Similarly, large sells following large positive returns have larger price impacts than large sells following large negative returns.

These results are consistent with two growing strands of literature that examine the behavior of informed investors. The first strand, represented by the theory models of Holden and Subrahmanyam (1992), Kaniel and Liu (2006), Chau and Vayanos (2008), Baruch, Panayides, and Venkataraman (2017), and Yang and Zhu (2019) as well as the empirical results of van Kervel and Menkveld (2019), Akey, Grégoire, and Martineau (2019), and Shkilko (2020) suggests that prices adjust to the presence of informed traders rather efficiently. As such, price movements caused by informed trading are likely to be followed by subsequent price movements in the same direction. The second strand, outlined in the model of Collin-Dufresne and Fos (2016), and the empirical results of Collin-Dufresne and Fos (2015), Kacperczyk and Pagnotta (2019), and Garriott and Riordan (2020) suggest that informed investors tend to time the market to reduce trading costs. As such, the informed are likely to trade after prices move in the opposite direction of their trading interest, causing reversals. Both of the above-mentioned literature strands find support in our results.

[Insert Figure 3]

Table 6 reports the price impact estimates for each action in the *Ret1* and *Ret10* return deciles. To analyze the relative contributions of prior order history and order size to price discovery, Table

6 also shows the difference in price impact estimates between large sells and large buys (bottom row) and *Ret10* and *Ret1* (Columns 3, 6 and 9). For market orders and limit order submissions, we find that order size contributes more to price discovery than historical returns while we report the opposite result for cancellations. For market orders, we find that order size is five times more important for price discovery than historical returns. In contrast, historical returns are 62% more important than order size for price discovery due to cancellations.

[Insert Table 6]

6.3 Time effects

So far, we have shown that the current state of the limit order book as well as price history contribute to price discovery. In this section, we ask if the time of order submission matters as proposed by Riccó, Rindi, and Seppi (2020). To examine this issue, we augment the RL model to include the time of day effects by splitting the trading day into three states: 10:00 (market open) - 10:30 (*morning*), 11:30 - 14:30 (*midday*) and 15:30 - 16:00 (market close) (*afternoon*). As before, we estimate an RL model with 30 actions, resulting in a final model consisting of $3 \times 30 = 90$ state-action pairs.

Early studies of intraday market activity document a U-shaped pattern in return volatility and spreads (greater values at the beginning and at the end of trading day) (e.g., Wood, McInish, and Ord (1985). More recent research suggests that the pattern has changed into an *S*-shaped, with higher values early and lower values near market close (e.g., Upson and Van Ness (2017); Barardehi and Bernhardt (2019)). The change is not surprising given the proliferation of index funds that trade heavily at the end of the day to rebalance their portfolios (Bogousslavsky and Muravyev (2020)). Consistent with this recent literature, our results in Figure 4 show that trading activity in the morning contributes more to price discovery than the afternoon and midday sessions. We test these relations formally in Table 7.

[Insert Figure 4]

[Insert Table 7]

Table 7 reveals that the average difference in magnitude of the price impact between the morning and afternoon sessions are 1.87, 0.65 and 0.25 basis points for market orders, limit order submissions and cancellations, respectively. However, the relative importance of time of day effects is generally lower than that of order size. While time of day effects contribute approximately the same amount as order size to price impact for cancellations, the time of day effects for market orders is substantially less. For market orders, time of day effects are 75% less important than order size.[9]

6.4 Contribution to price discovery

The results above show that the state of the order book, prior order history, and time of day all have effects on price formation. In this section, we probe deeper to understand the *relative* importance of these variables. To do so, we predict a price impact metric using all variables identified so far and then examine the contribution of each variable to the predicted metric, $\hat{PI}$.

Specifically, to obtain $\hat{PI}s$, we estimate an RL model with 30 actions and 300 states. 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 as described in Section 5. The 300 states comprise 10 depth imbalance deciles from Section 6.1, 10 return deciles from Section 6.2, and 3 time of day states from Section 6.3. We estimate this model to obtain $\hat{PI}s$ for $300 \times 30 = 9,000$ state-action pairs.

Next, using $\hat{PI}s$ obtained for each state-action pair, we move to determine the relative impor-

[9]The average contribution of order size is $(9.66+5.20)/2 = 7.43$ basis points and the average contribution of depth imbalance is 1.87 basis points.

tance of the price discovery conduits by estimating the following regression:

$$\begin{aligned}\hat{PI}_{s,a} &= \alpha_0 + \beta_1 SgnOrder_{s,a} + \beta_2 SgnOrderSize_{s,a} + \beta_3 DepthImb_{s,a} \\ &\quad + \beta_4 Ret_{s,a} + \beta_5 SgnTime_{s,a},\end{aligned} \tag{4}$$

where $\hat{PI}_{s,a}$ is the predicted price impact for state-action pair (s,a), and each of the regressors are ordinal variables corresponding to the previously identified price discovery determinants. *SgnOrder* is an indicator variable equal to 1 for buys and -1 for sells for market orders and limit order submissions. For cancellations, *SgnOrder* is an indicator variable equal to 1 for sell cancellations and -1 for buy cancellations. *SgnOrderSize* is the order size quintile, multiplied by *SgnOrder*. *DepthImb* is the depth imbalance decile at the time of the action. *Ret* is the historical stock return decile, based on midpoint returns in the 10 actions prior to the current action. *SgnTime* is equal to 1 for the morning session, 2 for the midday session, and 3 for the afternoon session multiplied by *SgnOrder*.

Because $\hat{PI}_{s,a}$ is a directional variable, the regressors must also be directional for the regression to be meaningful. Such variables as the order size and time of day are non-directional (always positive), and we therefore must multiply them by the order direction. More specifically, to investigate the contribution of time of day to price impact, we need to distinguish between a buy trade in the morning session, which typically has a positive price impact, from a sell trade in the morning session, which typically has a negative price impact. On the other hand, *DepthImb* and *Ret* represent deciles that span from negative to positive values, and therefore are directional variables that do not need to be adjusted.

The results for all market and limit order submissions and cancellations are in Table 8. We report the regression coefficients in Column 1, standardized regression coefficients in Column 2, and relative contributions of each variable to overall price discovery in Column 3. The standardized coefficients help us account for the variation in the scales of explanatory variables. Specifically, they reduce the scaling effect of some variables being expressed in deciles and others – in terciles. Finally, to compute relative contributions, we divide the magnitude of each variable's standardized coefficient by the sum of the magnitude of all standardized coefficients.

[Insert Table 8]

For all market and limit orders in Panel A, the regression approach confirms that all previously identified variables are important for price discovery; every regressor is statistically significant. Focusing on their relative contributions, the data show that the mere occurrence of an order determines about 15% of its price impact. Importantly, the current state of the order book contributes 57.4%, while the remaining variables – order size, prior order history, and time of day – determine between 7.2% and 10.2% of the price impact each.

Taken together, the results suggest that the current state of the limit order book plays a major role in price discovery through both market and limit orders, but particularly the limit orders, for which the state of the book determines the lion's share of price impact. Notably, prior order history and time of day are also quite important in their effect on prices. These results call for more theoretical research to identify new conditional determinants of price discovery.

6.5 Robustness

In the previous section, we use regression analysis to determine the relative contribution of each variable to price discovery. For additional robustness, and to provide a machine learning alternative to the regression method, we follow Easley, Lopez de Prado, O'Hara, and Zhang (2019) and compute the *mean decreased accuracy* (MDA) metric for each identified conduit. The MDA is commonly used in the machine learning literature, but is relatively new to finance. As we explain in more detail shortly, in our setting it measures the decrease in accuracy of the price impact estimate if one of the information variables (i.e., *Order*, *Order size*, etc.) is estimated with error. Importantly, the MDA does not assume a specific functional form and, compared with the regression analysis, can account for deeper relationships such as variable interactions. As such, it also does not require signing the regressor variables to account for order direction.

Estimating the MDA requires two parameters: the true predicted price impact and the randomized predicted price impact. The true predicted price impact, $\hat{PI}_{s,a}$, is the estimated price impact for each state-action pair (s,a) based on the ML model with 9,000 state-action pairs described in Section 6.4. In turn, to obtain the randomized predicted price impact, we randomize one of the five information variables k while holding all other variables constant. The randomized predicted price impact, $\tilde{PI}^k_{s,a}$, is the estimate associated with the randomly chosen state-action pair created by changing variable k.

For example, to test the relative importance of *DepthImb*, let us assume that the true price impacts in state-action pairs (*DepthImb D1, c*) and (*DepthImb D4, c*) are 1 bp and 2 bp, respectively, where *c* indicates that all other variables in both pairs are the same. To randomize the depth imbalance variable, the observation (*DepthImb D1, c*) is flipped to have the value of (*DepthImb D4, c*). Thus, the $\tilde{PI}^k_{s,a}$ value for (*DepthImb D1, c*) is assigned as 2 bp instead of 1 bp. Using these assignments, the MDA for variable k is estimated as follows:

$$MDA^k = \sum_{(s,a)=1}^{(S,A)} \left(\frac{|(\hat{PI}_{s,a} - \tilde{PI}^k_{s,a})|}{\hat{PI}_{s,a}} \right) / (S,A) \tag{5}$$

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 9.

[Insert Table 9]

Since the MDA measures deterioration in prediction accuracy, a variable with a high MDA is more important than a variable with a low MDA. Consistent with the earlier regression results, the MDA metric produces the same ranking for the relative importance of the price discovery conduits. Specifically, we find that depth imbalance contributes most to price discovery (MDA

= 5.06), followed by order occurrence (2.20), return (1.10), order size (0.58), and time of day (0.28).

7. 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, or time of day. In the meantime, conditional empirical analyses have not yet entered the mainstream due to methodological limitations. Specifically, the VAR models 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 address these limitations, we describe and apply a machine learning technique known as Reinforcement Learning (RL), which is particularly useful for applications characterized by non-linearities and multiple conditioning states. 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 current state of the limit order book and prior order history. Furthermore, 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.

Our analyses shed light on the predictions of multiple models of trader behavior. We show that the most informed agents tend to time periods of returns opposite the direction of their own trading. Finally, our intraday analyses suggest that the most informed market orders tend to be submitted during the morning hours, and their informativeness gradually declines through the day. Notably, limit order informativeness declines more sharply, consistent with the notion that as the probability of execution declines during the day, informed trading may shift to market orders.

Our contribution to the literature is threefold. First, we shed light on predictions of several theory models that have not been tested due to methodological 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. Specifically, we show that the current state of the limit order book is the dominant channel through which information flows into prices, followed by individual order submissions, prior order history, and the time effects.

References

Akey, P., V. Grégoire, and C. Martineau, 2019, "Price Revelation from Insider Trading: Evidence from Hacked Earnings News," *Working Paper, University of Toronto*.

Barardehi, Y., and D. Bernhardt, 2019, "Re-Appraising Intraday Trading Patterns: What You Didn't Know You Didn't Know," *Working Paper*, Chapman University.

Baruch, S., M. Panayides, and K. Venkataraman, 2017, "Informed trading and price discovery before corporate events," *Journal of Financial Economics*, 125(3), 561–588.

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.

Bogousslavsky, V., and D. Muravyev, 2020, "Should We Use Closing Prices? Institutional Price Pressure at the Close," *Working Paper*, Boston College.

Brogaard, J., T. Hendershott, and R. Riordan, 2019, "Price Discovery without Trading: Evidence from Limit Orders," *The 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.

Chakrabarty, B., P. K. Jain, A. Shkilko, and K. Sokolov, 2020, "Unfiltered Market Access and Liquidity: Evidence from the SEC Rule 15c3-5," *Management Science*, forthcoming.

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

Chung, K. H., B. F. Van Ness, and R. A. Van Ness, 1999, "Limit orders and the bid–ask spread," *Journal of Financial Economics*, 53(2), 255–287.

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, 2019, “Microstructure in the Machine Age,” *Working Paper*, Cornell University.

Easley, D., M. M. López de Prado, and M. O’Hara, 2012, “Flow Toxicity and Liquidity in a High-frequency World,” *The 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.

Garriott, C., and R. Riordan, 2020, “Trading on Long-Term Information,” *Working Paper*, The Bank of Canada.

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.

———, 1995, "One Security, Many Markets: Determining the Contributions to Price Discovery," *The Journal of Finance*, 50(4), 1175–1199.

Holden, C. W., and A. Subrahmanyam, 1992, "Long-Lived Private Information and Imperfect Competition," *The Journal of Finance*, 47(1), 247–270.

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.

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.

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

Philip, R., 2019, "Estimating permanent price impact via machine learning," *Journal of Econometrics*, forthcoming.

Riccó, R., B. Rindi, and D. J. Seppi, 2020, "Information, Liquidity, and Dynamic Limit Order Markets," *Working Paper*, Bocconi University.

Roşu, I., 2019, "Liquidity and Information in Limit Order Markets," *Journal of Financial and Quantitative Analysis*, pp. 1–48.

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., T. H. McInish, and H. Johnson, 2018, "Orders versus Trades on the Consolidated Tape," *Working Paper*, University of Texas at El Paso.

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.

van Kervel, V., and A. J. Menkveld, 2019, "High-Frequency Trading around Large Institutional Orders," *The Journal of Finance*, 74(3), 1091–1137.

Wood, R. A., T. H. McInish, and J. K. Ord, 1985, "An Investigation of Transactions Data for NYSE Stocks," *The Journal of Finance*, 40(3), 723–739.

Yang, L., and H. Zhu, 2019, "Back-Running: Seeking and Hiding Fundamental Information in Order Flows," *The Review of Financial Studies*, Forthcoming.

Table 1
Summary statistics

This table reports summary statistics for our sample stocks. Our sample period covers January 4, 2016 to February 29, 2016 for the top 20 stocks listed on the ASX based on market capitalizations on January 4, 2016. We report the number of trades (*NTrades*), average trade price in AUD (*Price*), 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 price impact estimates

This table reports the price impacts estimated via the VAR and RL models. Our sample period covers January 4, 2016 to February 29, 2016 for the top 20 stocks listed on the ASX based on market capitalizations on January 4, 2016. 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 open) and two actions (to buy or to sell). For each stock, we report the price impact estimate for each buy and sell action. The bold fonts indicate that the magnitude of the RL price impact estimates fall 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 impact estimates for different actions and order sizes

This table reports the average price impacts for market orders (*MO*), limit order submissions (*LO*) and cancellations (*Cancel*). Our sample period covers January 4, 2016 to February 29, 2016 for the top 20 stocks listed on the ASX based on market capitalizations on January 4, 2016. We estimate the price impact function using the RL model with one state (market open) and 30 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 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 over the 20 stocks for each state-action pair. We test for differences between *MO* and *LO* (Column 5) and *LO* and *Cancel* (Column 6) using a T-test. * and ** represents significance at the 5% and 1% levels, respectively.

Action	Market orders (MO)	Limit orders (LO)	Cancellations (Cancel)	MO - LO	LO - Cancel
Sell Q5	-3.54	-0.83	-0.13	-2.71**	-0.70**
Q4	-2.59	-0.46	-0.16	-2.13**	-0.29**
Q3	-2.33	-0.33	-0.22	-1.99**	-0.11*
Q2	-2.07	-0.31	-0.54	-1.76**	0.23
Q1	-1.59	-0.30	-0.08	-1.29**	-0.22**
Q1	1.53	0.32	0.13	1.21**	0.19**
Q2	1.97	0.36	0.53	1.61**	-0.17
Q3	2.24	0.36	0.24	1.88**	0.12*
Q4	2.50	0.49	0.21	2.01**	0.29**
Buy Q5	3.43	0.82	0.18	2.60**	0.64**

Table 4
Contribution of each action to price discovery

This table reports the contribution of each action to overall price discovery, after adjusting for the frequency of each action. Our sample period covers January 4, 2016 to February 29, 2016 for the top 20 stocks listed on the ASX based on market capitalizations on January 4, 2016. We estimate the price impact function using the RL model with one state (market open) and 30 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 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 multiply the estimated price impact of each action by the frequency of the action occurring in the dataset. The final row reports the total frequency and contribution of market orders, limit order submissions and cancellations to overall price discovery.

Action	Market orders		Limit orders		Cancellations	
	Freq. (%)	Cont. (%)	Freq. (%)	Cont. (%)	Freq. (%)	Cont. (%)
Sell Q5	2.29	10.38	4.92	5.80	2.65	0.53
Q4	1.20	4.24	5.14	3.44	3.03	0.78
Q3	0.65	2.13	5.58	3.06	4.23	1.42
Q2	1.62	4.85	5.55	2.64	3.29	1.73
Q1	2.59	6.20	5.37	2.48	2.02	0.34
Q1	2.52	5.54	5.61	2.79	1.98	0.45
Q2	1.55	4.28	5.25	2.80	3.16	1.9
Q3	0.66	2.05	5.47	3.19	4.16	1.67
Q4	1.24	4.13	5.19	3.67	3.14	0.96
Buy Q5	2.32	10.04	4.92	5.78	2.67	0.72
Total (%)	16.65	53.83	53.01	35.65	30.34	10.51

Table 5
Price impact estimates conditional on depth imbalance

This table reports the average price impacts for market orders, limit order submissions and cancellations based on the depth imbalance at the time of the action. Our sample period covers January 4, 2016 to February 29, 2016 for the top 20 stocks listed on the ASX based on market capitalizations on January 4, 2016. We estimate the price impact function using the RL model with 10 states based on depth imbalance at the time of the action, which are formed into 10 deciles, and 30 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 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 over the 20 stocks for each state-action pair and report the average price impact for the most negative (*D1*) and most positive (*D10*) depth imbalance deciles. We report the difference between *D10* and *D1* in Columns 4, 7 and 10, and the difference between the largest buys (*Buy Q5*) and sells (*Sell Q5*) in the last row.

	Market orders			Limit orders			Cancellations		
Action	D10 [1]	D1 [2]	D10 - D1 [3]	D10 [4]	D1 [5]	D10 - D1 [6]	D10 [7]	D1 [8]	D10 - D1 [9]
Sell Q5	1.97	-9.92	11.89	4.32	-6.31	10.63	5.70	-4.86	10.56
Q4	3.43	-9.21	12.64	4.90	-5.85	10.75	5.72	-4.94	10.66
Q3	3.67	-8.94	12.60	5.07	-5.70	10.76	5.76	-4.91	10.66
Q2	4.08	-9.11	13.18	5.10	-5.70	10.80	5.93	-4.85	10.78
Q1	4.48	-8.55	13.03	5.19	-5.66	10.85	5.69	-5.18	10.88
Q1	8.06	-4.17	12.22	5.81	-5.08	10.88	5.36	-5.53	10.89
Q2	8.44	-3.81	12.24	5.82	-4.95	10.77	4.89	-5.78	10.67
Q3	8.62	-3.26	11.88	5.81	-4.90	10.71	5.02	-5.63	10.65
Q4	9.03	-3.23	12.27	6.01	-4.73	10.74	5.09	-5.54	10.63
Buy Q5	9.73	-1.85	11.58	6.38	-4.20	10.58	5.02	-5.54	10.56
Average			12.34			10.74			10.69
Buy Q5 - Sell Q5	7.76	8.07		2.06	2.11		-0.68	-0.68	

Table 6
Price impact estimates conditional on historical stock returns

This table reports the average price impacts for market orders, limit order submissions and cancellations based on historical stock returns. Our sample period covers January 4, 2016 to February 29, 2016 for the top 20 stocks listed on the ASX based on market capitalizations on January 4, 2016. We estimate the price impact function using the RL model with 10 states based on 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 30 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 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 compute the average price impact over the 20 stocks for each state-action pair and report the average price impact for the most negative (*Ret1*) and most positive (*Ret10*) return deciles. We report the difference between *Ret10* and *Ret1* in Columns 4, 7 and 10, and the difference between the largest buys (*Buy Q5*) and sells (*Sell Q5*) in the last row.

	Market Orders			Limit orders			Cancellations		
Action size	Ret1	Ret10	Ret1 - Ret10	Ret1	Ret10	Ret1 - Ret10	Ret1	Ret10	Ret1 - Ret10
Sell Q5	-3.69	-7.39	3.70	-0.52	-3.08	2.56	1.55	-0.79	2.34
Q4	-2.88	-5.87	2.99	0.09	-2.47	2.56	1.68	-0.86	2.54
Q3	-3.05	-6.02	2.97	0.46	-2.23	2.69	1.85	-0.88	2.73
Q2	-2.54	-5.67	3.13	0.45	-1.94	2.38	1.90	-0.79	2.69
Q1	-1.91	-4.74	2.83	0.40	-2.10	2.49	1.95	-0.73	2.68
Q1	4.12	1.62	2.49	1.90	-0.68	2.57	0.60	-2.20	2.80
Q2	5.08	2.51	2.57	1.88	-0.62	2.49	0.69	-1.99	2.68
Q3	4.94	2.66	2.28	2.01	-0.69	2.70	0.83	-1.84	2.67
Q4	5.24	2.74	2.50	2.32	-0.28	2.60	0.78	-1.58	2.37
Buy Q5	6.48	3.95	2.53	2.84	0.36	2.48	0.79	-1.63	2.42
Average			2.79			2.55			2.59
Buy Q5 - Sell Q5	10.17	11.34		3.36	3.43		-0.76	-0.84	

Table 7
Price impact estimates conditional on the time of the day

This table reports the average price impacts for market orders, limit order submissions and cancellations based on the time of the action. Our sample period covers January 4, 2016 to February 29, 2016 for the top 20 stocks listed on the ASX based on market capitalizations on January 4, 2016. We estimate the price impact function using the RL model with 3 states based on the time of the action (*AM*, *Midday*, *PM*) and 30 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 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 compute the average price impact over the 20 stocks for each state-action pair and report the average price impact for the *AM* and *PM* sessions. We report the difference between *AM* and *PM* in Columns 4, 7 and 10, and the difference between the largest buys (*Buy Q5*) and sells (*Sell Q5*) in the last row.

	Market Orders			Limit orders			Cancellations		
Action	AM	PM	AM - PM	AM	PM	AM - PM	AM	PM	AM - PM
Sell Q5	-4.82	-2.66	-2.16	-1.57	-0.52	-1.04	0.42	0.12	0.29
Q4	-3.60	-1.73	-1.87	-0.99	-0.28	-0.71	0.40	0.14	0.26
Q3	-3.35	-1.46	-1.89	-0.73	-0.19	-0.54	0.52	0.32	0.20
Q2	-3.03	-1.20	-1.83	-0.66	-0.19	-0.47	0.49	0.35	0.14
Q1	-2.40	-0.95	-1.45	-0.69	-0.22	-0.48	0.56	-0.02	0.59
Q1	2.33	0.85	1.48	0.66	0.21	0.45	-0.44	0.07	-0.51
Q2	3.01	1.00	2.01	0.69	0.23	0.46	-0.41	-0.34	-0.08
Q3	3.12	1.28	1.84	0.74	0.21	0.53	-0.43	-0.33	-0.10
Q4	3.48	1.59	1.89	1.05	0.29	0.77	-0.27	-0.10	-0.17
Buy Q5	4.84	2.54	2.30	1.58	0.54	1.04	-0.28	-0.07	-0.21
Average			1.87			0.65			0.25
Buy Q5 - Sell Q5	9.66	5.20		3.15	1.06		-0.69	-0.20	

Table 8
Contribution to price discovery

This table reports estimation results for an OLS regression of price impact against all information variables specified in equation (4). Our sample period covers January 4, 2016 to February 29, 2016 for the top 20 stocks listed on the ASX based on market capitalizations on January 4, 2016. The dependent variable is the predicted price impact estimated via our full RL model based on 9,000 state-action pairs. The independent variables are ordinal variables relating to our earlier analyses. For market orders and limit order submissions, *SgnOrder* is an indicator variable equal to 1 for buys and -1 for sells. For cancellations, *SgnOrder* is an indicator variable equal to 1 for sell cancellations and -1 for buy cancellations. *SgnOrderSize* is the order size quintile, multiplied by *SgnOrder*. *DepthImb* is the depth imbalance decile at the time of the action. *Ret* is the historical stock return decile, based on midpoint returns in the 10 actions prior to the current action. *SgnTime* is equal to 1 for the morning session, 2 for the midday session and 3 for the afternoon session and is multiplied by *SgnOrder*. Columns 1 and 2 present the regression coefficients and standardized regression coefficients, respectively. Column 3 presents the relative contribution of each variable, which is the magnitude of the standardized coefficient divided by the sum of the magnitudes of all standardized coefficients. T-statistics are in parentheses.

	(1)	(2)	(3)
SgnOrder	0.877	0.877	15.3%
	(15.41)	(15.41)	
SgnOrderSize	0.172	0.57	9.9%
	(14.47)	(14.47)	
DepthImb	1.148	3.297	57.4%
	(196.34)	(196.34)	
Ret	-0.415	-0.587	10.2%
	(-34.94)	(-34.94)	
SgnTime	-0.191	-0.413	7.2%
	(-9.29)	(-9.29)	
R-squared	0.907	0.907	

Table 9
MDA contribution to price discovery

This table reports the Mean Decreased Accuracy (MDA) estimated as follows for each variable 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 9,000 state-action pairs, and $\tilde{PI}^k_{s,a}$, is the estimate created by randomly changing variable k. For example, to test the relative importance of *DepthImb*, let us assume that the true price impacts in state-action pairs (*DepthImb D1, c*) and (*DepthImb D4, c*) are 1 bp and 2 bp, respectively, where c indicates that all other variables in both pairs are the same. To randomize the depth imbalance variable, the observation (*DepthImb D1, c*) is flipped to have the value of (*DepthImb D4, c*). Thus, the $\tilde{PI}^k_{s,a}$ value for (*DepthImb D1, c*) is assigned as 2 bp instead of 1 bp. We repeat this process 100 times for each variable k and report the mean and standard deviation of the MDA for each variable.

	Mean	Std. Dev.
Order	2.20	0.21
OrderSize	0.58	0.06
DepthImb	5.06	0.39
Ret	1.10	0.08
Time	0.28	0.02

Figure 1. Price impact functions for different actions and order sizes

These figures plot the price impact functions for market orders (Panel A), limit order submissions (Panel B) and limit order cancellations (Panel C). Our sample period covers January 4, 2016 to February 29, 2016 for the top 20 stocks listed on the ASX based on market capitalizations on January 4, 2016. We estimate the price impact function using the RL model with one state (market open) and 30 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 categorized into 5 size quintiles, resulting in 6x5=30 overall actions. The estimated model has 1x30=30 state-action pairs. We report the average price impact over the 20 stocks for each state-action pair.

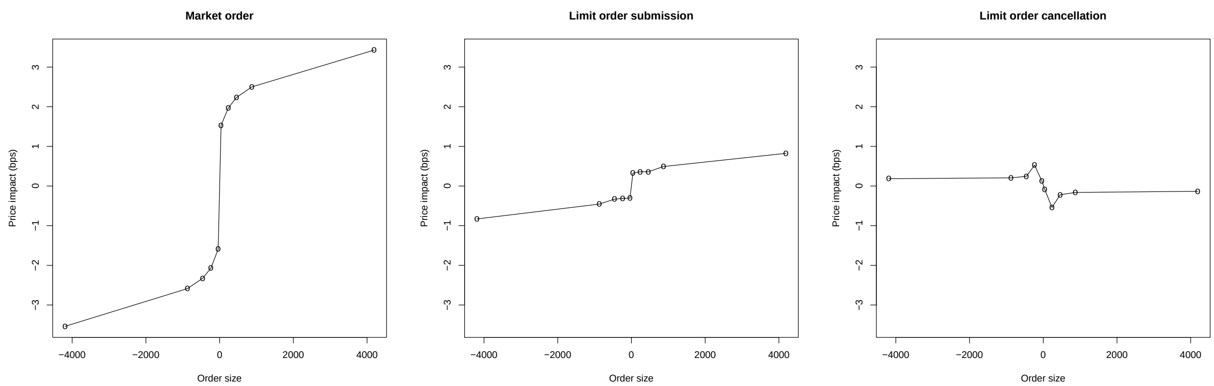

Figure 2. Price impact functions conditional on depth imbalance

These figures plot the price impact functions for market orders (Panel A), limit order submissions (Panel B) and limit order cancellations (Panel C) based on the depth imbalance at the time of the action Our sample period covers January 4, 2016 to February 29, 2016 for the top 20 stocks listed on the ASX based on market capitalizations on January 4, 2016. We estimate the price impact function using the RL model with 10 states based on depth imbalance at the time of the action, which are formed into 10 deciles, and 30 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 categorized into 5 size quintiles, resulting in 6x5=30 overall actions. The estimated model has $10 \times 30 = 300$ state-action pairs. We compute the average price impact over the 20 stocks for each state-action pair. To save space, we report the average price impacts for the most negative (*D1*) and most positive depth (*D10*) imbalance deciles.

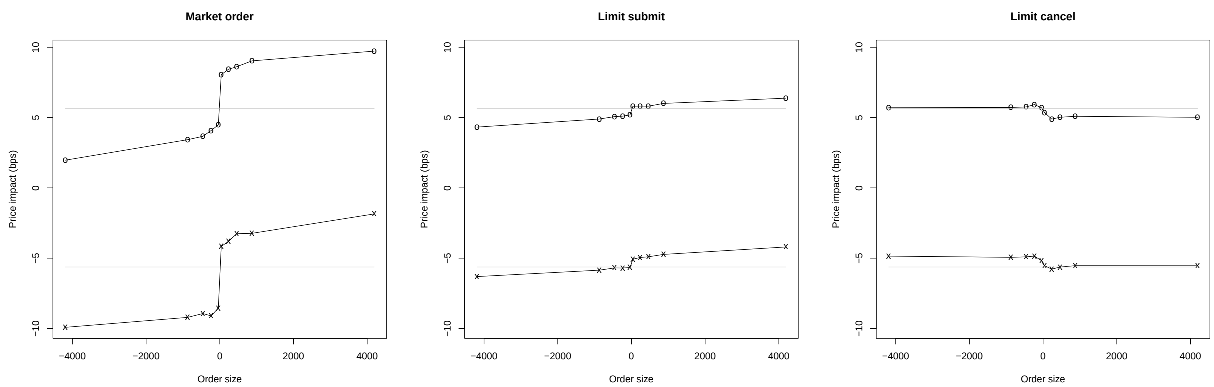

Figure 3. Price impact functions conditional on historical stock returns

These figures plot the price impact functions for market orders (Panel A), limit order submissions (Panel B) and limit order cancellations (Panel C) based on historical stock returns. Our sample period covers January 4, 2016 to February 29, 2016 for the top 20 stocks listed on the ASX based on market capitalizations on January 4, 2016. 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 30 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 categorized into 5 size quintiles, resulting in 6x5=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. To save space, we report the average price impacts for the most negative (*Ret1*) and most positive (*Ret10*) return deciles.

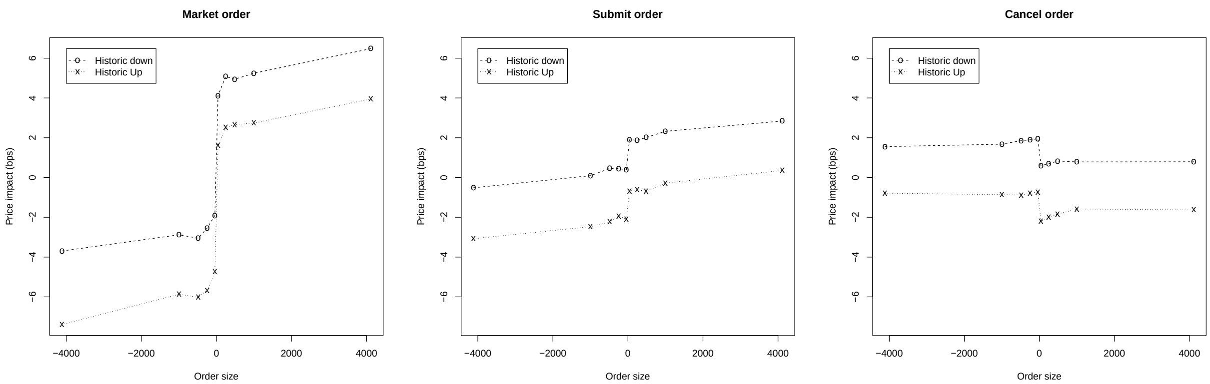

Figure 4. Price impact functions conditional on the time of the day

These figures plot 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 the action. Our sample period covers January 4, 2016 to February 29, 2016 for the top 20 stocks listed on the ASX based on market capitalizations on January 4, 2016. We estimate the price impact function using the RL model with 3 states based on the time of the action (*AM*, *Midday*, *PM*) and 30 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 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. deciles.