Estimating permanent price impact

Richard Philip[1]

[1]Discipline of Finance, The University of Sydney

August 4, 2017

Abstract

Traditional methods used to estimate the permanent price impact of a trade could be misspecified when a nonlinear relationship between permanent price impact and trade size exists, potentially changing research inferences. We augment existing vector autoregression (VAR) estimation methods to model the nonlinear relationship between permanent price impact and trade size. However, when we include additional variables to capture today's trading environment, it is difficult to estimate permanent price impact via traditional methods. We propose an alternative technique, a flexible reinforcement learning (RL) algorithm, which makes fewer assumptions about the data generating process and results in less estimation issues.

[0]This paper has benefited from the comments of Sean Anthonitz, Amy Kwan, Michael Goldstein, Talis Putnins, Ioanid Rosu, Steve Satchell, Andriy Shkilko and Bart Yueshen

1 Introduction

In asymmetric information models, traders interact with a specialist to set market prices. One critical component of asymmetric information models is that a trade conveys information. To empirically estimate the information content of a trade, researchers typically use a trade's permanent price impact, estimated via the impulse response function of a vector autoregression (VAR) framework. The VAR framework represents a bivariate linear time series model of trades and quote revisions. While researchers use a change in the midquote price to represent a quote revision, for the trade variable, researchers typically choose between 1) *trade sign*, 2) *signed trade size* or 3) a polynomial trade variable. Interestingly, the literature provides little guidance on the choice of trade variable. However, we demonstrate that the choice of trade variable alters conclusions for empirical research relying on price discovery models.

The existing theoretical and empirical literature provides justification for all of the three trade variables. Theoretical models suggest that permanent price impact is linear with trade size (see Kyle (1985) and Huberman and Stanzl (2004)), which supports using *signed trade size* as the trade variable. In contrast, the empirical literature shows that trade size contributes little incremental explanatory power above the trade direction when estimating permanent price impact (see Hasbrouck (2007), Jones et al. (1994)), which supports using *trade sign* as the trade variable. Lastly, Hasbrouck (1991a) argues that the assumed linearity between the trade variable and quote updates may be a tenuous approximation and could lead to a misspecified model. To overcome this misspecification, Hasbrouck (1991a) advocates the use of polynomial terms as the trade variable to capture any nonlinear relationships between quote updates and trade size. Despite these findings, polynomial terms are rarely used in practice.

Further, the existing VAR framework may not be applicable to today's modern trading environment (see O'Hara (2015)). Hasbrouck (1991a) asserts that the VAR framework is misspecified if the dealer possesses an informational advantage in their knowledge of the limit order book. In the modern market environment, co-located market making high frequency traders (HFT) have this information advantage due to their speed advantage.[1] However, as we demonstrate, it is difficult to include additional explanatory variables which capture information contained in the limit order book in the traditional VAR model due to an assumed linearity.

In this study, we propose a new flexible VAR system to capture nonlinear relationships between permanent price impact and trade size, which is easier to implement than a polynomial VAR. We highlight that using *trade sign* or *signed trade size* as the trade variable leads to incorrect conclusions, due to an assumed linearity between quote updates and the trade variable. Furthermore, we demonstrate the difficulty of including additional control variables within a VAR framework. The VAR model is misspecified if additional control variables are omitted or modelled linearly. We provide a solution based on a reinforcement learning (RL) algorithm, which suffers fewer estimation issues. A complete, commented R file to implement our RL algorithm is available in the internet appendix.

RL techniques offer several advantages over traditional VAR methods. In comparison to VAR models, RL is a model free method to estimate a Markov Decision Process and only assumes that

[1] Brogaard et al. (2014) and Goldstein et al. (2016) document high frequency traders (HFT) exploit their informational advantage by conditioning trades on the shape of the order book.

the data is Markov. Thus, RL models are highly flexible and can be used for complex processes, such as stock returns that are difficult to realisticly model through a set of equations. The RL technique is widely used in the computer science literature and its flexibility is highly suitabile for modelling the limit order book (see Easley et al. (2013)). When there are only a low number of model specifications (e.g., trade size or trade sign), we show that the RL model yields similar empirical results to the VAR model. However, when the model specifications become more complex, by including both trade size and order book variables, the RL model outperforms the VAR model. Additionally, we can use the RL model to decompose the permanent price impact into the private information component and a component reflecting the trader's ability to exploit public information contained in the limit order book, which the VAR framework does not allow.

We document several findings, which have important implications for empirical market microstructure research. First, when we apply our flexible VAR system to empirical data, we document a strong nonlinear relationship between quote updates and trade size. Specifically, as a function of trade size, the permanent price impact is positive, increasing and concave for all 20 sample stocks.[2] This nonlinearity is particularly important for research comparing price discovery across trading venues (Barclay et al. (2003)), between lit and dark markets (Comerton-Forde and Putnins (2015)) or between different trader types (Brogaard et al. (2016)) as erroneous research inferences can occur if these trading venues or market participants have differing average trade sizes. Additionally, using the flexible VAR system, we document a strong autocorrelation in trade sizes. Specifically, we find that a small trade is more likely to follow a small trade while a large trade is more likely to precede a large trade.

Second, using an example of two equally uninformed market participants, who have different average trade sizes, we demonstrate the potential issues of making inferences from a misspecified VAR model. Since both market participants are uninformed with no private information, they should have equivalent permanent price impacts. However, depending on the choice of trade variable, the traditional VAR model provides conflicting results. When we use *trade sign* as the trade variable in the VAR specification, we find that the market participant with a larger average trade size has a larger permanent price impact. In contrast, we obtain the opposite conclusion when we estimate a VAR model using *signed trade size* as the trade variable (i.e., the market participant with a smaller average trade size has a larger price impact). This discrepancy in conclusions is attributed to the nonlinear relationship between permanent price impact and trade size. On the other hand, when we estimate the permanent price impact of a trade for both market participants using our flexible nonlinear VAR system, we correctly conclude that both participants cause equal permanent price impact.

Third, we re-estimate a multi-factor VAR model which includes a variable to control for information contained in the shape of the limit order book.[3] We compare the permanent price impact of two traders who have no private information about future prices. The first trader conditions their trades on public information contained in the shape of the limit order book, while the second trader does not. Since both traders have no private information we expect the traders to have the same permanent price impact. However, we show that the VAR model again provides mixed results depending on the chosen specification. Thus, the traditional VAR fails to decompose the permanent price impact into a component attributed to private information, and a component reflecting the

[2]These findings are consistent with Hasbrouck (1991a) and Engle and Patton (2004).

[3]We use the depth imbalance which is shown to predict short term price movements, see Engle and Patton (2004), Cao et al. (2009) and Goldstein et al. (2016).

trader's ability to exploit public information contained in the limit order book.

On the other hand, using the RL model the correct inferences are made. Further, we demonstrate the information contained in the shape of the limit order book is a significant component of permanent price impact. Notably, the cross sectional variation of permanent price impact due to the shape of the order book is 3.5 times larger than the cross sectional variation of permanent price impact due to trade size. These findings highlight the importance of controlling for the shape of the limit order book information when estimating the private information contained in a trade.

Fourth, we contribute to the growing body of literature suggesting that limit order arrivals also contribute to price discovery (see O'Hara (2015) and Brogaard et al. (2016)). Using the RL model, we report similar findings to (Brogaard et al., 2016) for estimates of permanent price impacts for limit order submissions and cancellations. Using the RL model, we also extend their work and document that market orders contribute 57.3% of price discovery, while limit order submissions and cancellations contribute 30.1% and 12.7% to price discovery, respectively.

Fifth, our results also support the predictions from Mendelson and Tunca (2004), who provide one theoretical explanation for why a nonlinear relationship between permanent price impact and trade size exists. They suggest agents are endogenous in the volumes they trade; Agents trade large volumes when the market is liquid, resulting in small price impacts per unit traded, and small volumes when the market is illiquid, resulting in large price impacts per unit traded. We provide empirical evidence supporting this theory and show a trade executing against the thin side of the order book has more permanent price impact than a trade executing against the thick side of the order book. Futhermore, we demonstrate that traders minimize their price impact by conditioning their trades on the available liquidity. Traders are 13.7 times more likely to submit a large order when the market is liquid than when it is illiquid and 9.9 times more likely to submit a small order than a large order when the market is illiquid. These findings extend Collin-Dufrense and Fos (2015) who suggest that informed traders select times of higher liquidity when they trade. Our results show large traders also select times of higher liquidity when they trade.

Finally, the concave monotonic relationship we document between permanent price impact and order size provides one possible explanation for why it is generally accepted that trade size contributes little incremental explanatory power above the trade sign variable when estimating the VAR model. A VAR model estimated using *signed trade size* generates a linear relationship between order size and price impact. Conversely, a VAR model estimated using *trade sign* generates a step function, which better approximates the concavity in the price impact function we document.

This paper is organized as follows. Section 2 presents the VAR framework and explains the generalization we propose. Section 3 describes the data used in the analysis. Section 4 compares different VAR specifications under multiple scenarios and discusses the results. Section 5 develops the estimation technique that employs an RL algorithm while Section 6 discusses the results of the RL estimation technique. Section 7 explains why permanent price impact has a nonlinear relationship with size. Finally, Section 9 summarizes and concludes.

2 The VAR Model

2.1 Nonlinear VAR model

To study the information content of a trade, Hasbrouck (1991a) proposes the following VAR system,

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

Common implementations of this model define r_t as the change in the natural logarithm of the midpoint that follows a trade at time, t and x_t is the trade indicator variable. The immediate price impact of a trade is given by β_0 while the permanent price impact of a trade is obtained from the vector moving average (VMA) representation of the VAR model:

$$
\begin{bmatrix} r_t \\ x_t \end{bmatrix} = \begin{bmatrix} a(L) & b(L) \\ c(L) & d(L) \end{bmatrix} \begin{bmatrix} \epsilon_{1,t} \\ \epsilon_{2,t} \end{bmatrix} \tag{2}
$$

where L is a lag operator. The permanent price impact of a trade, $r^x_{t+\infty}$, is measured by the cumulative impulse response function.

$$
r^x_{t+\infty} = \sum_{i=1}^{\infty} b_i \epsilon_{2,t} \tag{3}
$$

The model assumes a linear relationship exists between the trade indicator variable, which is commonly the *trade sign* or *signed trade size*, and quote revisions. For clarity, we refer to estimation of Equation 1 using *trade sign* (*signed trade size*) for the trade variable as the *trade sign VAR* (*trade size VAR*).

Hasbrouck (1991a) highlights that the model is misspecified when a nonlinear relationship between quotes and *signed trade size* and/or *trade sign* exists. As such, Hasbrouck (1991a) suggests using a quadratic VAR specification to capture any nonlinear relationships between trade size and permanent price impact. To implement a quadratic VAR specification we replace the x_t terms in Equation 1 by a set $\{x^0_t, x_t, x^2_t\}$, where x^0_t equals 1 for buyer initiated trades and -1 for seller initiated trades. x_t is the signed trade size and $x^2_t = x^0_t|x_t|^2$. However, such a transformation raises three possible concerns. First, if the relationship between trade size and permanent price impact is not quadratic then this model remains misspecified and the concerns found in Equation 1 persist. Second, estimating the impulse response function is not straight forward. When a trade occurs, the error variables for x^0_t, x_t and x^2_t are all simultaneously shocked. The implementation of these simultaneous shocks is constrained by the model specification and knowledge of how future residuals unfold.[4] Third, there is a large number of parameters to estimate. Given there are 7 variables, a model which has 5 lags requires 245 coefficient estimates. Because of these potential estimation issues, we propose a similar approach to Engle and Patton (2004) to capture any nonlinear relationships. Specifically, we discretize trade size into n quantiles to obtain the following VAR, which for clarity we refer to as the *nonlinear VAR*.

$$\begin{aligned} r_t &= \sum_{i=1}^{\infty} \alpha_i r_{t-i} + \sum_{j=1}^{n}\sum_{i=0}^{\infty} b_{ji} x^j_{t-i} + \epsilon_{1,t} \\ x^k_t &= \sum_{i=1}^{\infty} \delta^k_i r_{t-i} + \sum_{j=1}^{n}\sum_{i=1}^{\infty} c^k_{ji} x^j_{t-i} + \epsilon^k_t \quad \text{for } k = 1, ..n \end{aligned} \tag{4}$$

where x^j_t is +1 (-1) if a trade is buyer (seller) initiated and its trade size falls in quantile j, and 0 otherwise. In contrast to the previously outlined VARs, which assume a relationship between trade size and permanent price impact, Equation 4 only assumes that trades of equal size, have equal price impact.

2.2 Nonlinear VAR model with multiple market participants

The concern for misspecification due to nonlinear relationships is most apparent when comparing the difference in price impact between two market participants. To simulataneously estimate the price impact of two market particpants, most researchers rely on the VAR model of Barclay et al. (2003). Barclay et al. (2003) develop a three equation VAR system, which includes a trade indicator variable for each of the market participants. A 10 lag version of this model is written as:

$$\begin{aligned} r_t &= \sum_{i=1}^{10} \alpha_i r_{t-i} + \sum_{i=1}^{10} \beta_i x^e_{t-i} + \sum_{i=0}^{10} \zeta_i x^m_{t-i} + \epsilon_{1,t} \\ x^e_t &= \sum_{i=1}^{10} \delta_i r_{t-i} + \sum_{i=1}^{10} \phi_i x^e_{t-i} + \sum_{i=1}^{10} \nu_i x^m_{t-i} + \epsilon_{2,t} \\ x^m_t &= \sum_{i=1}^{10} \vartheta_i r_{t-i} + \sum_{i=1}^{10} \psi_i x^e_{t-i} + \sum_{i=0}^{10} \eta_i x^m_{t-i} + \epsilon_{3,t} \end{aligned} \tag{5}$$

Where r_t is the log return during interval t, x^m_t is the signed sum of trades (+1 buys, -1 sells) during the interval t for market participant m and x^e_t is the signed sum of trades during interval t for market participant e. Because *trade sign* is the chosen trade variable, this model specification assumes either, (1) trades with different size will cause the same permanent price impact, or (2) both market participants have the same trade size. If assumption (1) is violated, to allow a like for

[4] A full discussion of this concern is presented in the appendix.

like comparison between the two market participants, we must have market particpants who have the same trade size (assumption (2)).

Alternatively, if the trade variable is the *signed trade size* then, for similar reasons, the model assumes (1) price impact is linear with trade size, or (2) both market participants have the same trade size. Hasbrouck (1991a) and Engle and Patton (2004) demonstrate that the first assumption does not hold for either trade variable specification. The second assumption is unlikely to hold for many empirical studies. For example Brogaard et al. (2016) report and average trade size of $4,299 and $8,526 for HFT and non-HFT respectively. Similarly, Comerton-Forde and Putnins (2015) report an average size of dark trades between $10,000 to $150,000 in their sample, while lit trades have an average size of $5,000 to $13,000. Lastly, (Barclay et al., 2003) report ECN trades are smaller, on average, than market maker trades. Specifically, 51% of small trades occurs on ECNs, but this percentage declines to only 2% for large trades.

In this study we refer to Equation 5 as the *two participant trade sign VAR* when using *trade sign* as the trade variable. In contrast, when we use *signed trade size* as the trade variable we refer to the model as the *two participant trade size VAR*.

To illustrate the potential pitfalls of violating either of the two assumptions made by a two participant VAR, we model two participants, which differ in average trade size but have the same permanent price impact that is not linear with trade size. The top left panel of Figure 1, Panel A shows the relationship between permanent price impact and trade size for both the large trader and the small trader (subpanel (*i*)). We assume that both the large and the small trader have the same nonlinear permanent price impact as a function of trade size (black line). Thus, a correctly specified model should conclude that the large and the small trader have the same permanent price impact.

The solid black line, for the large trader in Panel A (subpanel *(ii)*), shows the permanent price impact of the large trader when we estimate a linear VAR model using *trade sign* as the trade variable. Similarly, the estimated permanent price impact for the small trader is depicted with the solid black line in subpanel *(iii)*. Comparing between the lines for the large and small trader, we arrive at the incorrect inference that the large trader has a larger price impact, or more private information, than the small trader.

In Figure 1, Panel B, we replace the *trade sign* variable with *signed trade size*. Again, both the large and the small trader have the same nonlinear permanent price impact as a function of trade size and thus, both traders should have the same permanent price impact after controlling for trade size. However, when we estimate a linear VAR model using *signed trade size* as the trade variable, we observe the gradient of the estimated price impact function of the small trader (subpanel *(vi)*) is larger than the gradient of the large trader (subpanel *(v)*). Accordingly, we draw the incorrect conclusion that the small trader has a larger price impact than the large trader, for equivalent sized trades; the opposite inference to what we make when we use *trade sign* as the trade variable.

Importantly, linear specifications of the VAR model using either *trade sign* or *signed trade size* as the trade variable lead to incorrect conclusions. Further, we show that the result flips depending on the choice of the trade variable. If we use *trade sign* as the trade variable in the VAR, we find that the large trader has a larger price impact than the small trader. In contrast, when we use *signed trade size* in the model, we find that the smaller trader has the larger price impact.

Insert Figure 1 here

To overcome these estimation issues we propose a flexible VAR system that relaxes the two assumptions of Barclay et al. (2003). Specifically, we extend the nonlinear VAR model proposed in Equation 4 to a multiple market participant case which we refer to as the *two participant nonlinear VAR*

$$\begin{aligned} r_t &= \sum_{i=1}^{\infty} \alpha_i r_{t-i} + \sum_{j=1}^{n} \sum_{i=0}^{\infty} b_{ji}^{e} x_{t-i}^{e,j} + \sum_{j=1}^{n} \sum_{i=0}^{\infty} b_{ji}^{m} x_{t-i}^{m,j} + \epsilon_{1,t} \\ x_t^{e,j} &= \sum_{i=1}^{\infty} \delta_i^j r_{t-i} + \sum_{j=1}^{n} \sum_{i=1}^{\infty} c_{ji}^{e} x_{t-i}^{e,j} + \sum_{j=1}^{n} \sum_{i=1}^{\infty} c_{ji}^{m} x_{t-i}^{m,j} + \epsilon_{2,t}^{j} \\ x_t^{m,j} &= \sum_{i=1}^{\infty} \delta_i^j r_{t-i} + \sum_{j=1}^{n} \sum_{i=1}^{\infty} c_{ji}^{e} x_{t-i}^{e,j} + \sum_{j=1}^{n} \sum_{i=1}^{\infty} c_{ji}^{m} x_{t-i}^{m,j} + \epsilon_{3,t}^{j} \end{aligned} \tag{6}$$

where $x_t^{y,j}$ is +1 (-1) if a trade has volume which lies is quantile j, and is a buy (sell) initiated by market participant y, or 0 otherwise, where $y \in (e, m)$. The impulse response function arising from a shock to $x_t^{y,j}$ will provide an estimate for the permanent price impact of a trade of specific size j initiated by market participant y. By comparing the impulse response function from shocking $x_t^{e,j}$ with the impulse response function from shocking $x_t^{m,j}$, we are able to make a fair comparison of the permanent price impacts between participants e and m for trades of equal size, j. Using Equation 6 we can then disentangle if one market participant has a larger price impact due to larger trade size or more private information.

3 The Data

We use full order book data for the Australian Securities Exchange (ASX) extracted from the SIRCA database for the period January 1, 2015 to February 28, 2015.[5] The SIRCA database contains every trade, order submission, cancellation and amendment allowing us to fully reconstruct the order book. In our analysis, we consider the 20 largest stocks by market capitalization. Table 1 presents the summary statistics for the 20 stocks.

Insert Table 1 here

For the purpose of our study, data from the ASX offers several advantages over other markets. First, the Australian market remains largely consolidated during our sample period, allowing for the correct chronological ordering of trades. Because the ASX executes almost 90% of lit volume, unobserved trades on competing exchanges are less likely to contaminate our price impact measures. Second, we can accurately infer trade direction from the reconstructed limit order book and do not

[5]This sample is long enough for precise estimations, see Dufour and Engle (2000))

need to rely on trade classification algorithms, such as Lee and Ready (1991).[6] Third, most datasets report large orders that execute against several resting limit orders as multiple trades. [7] Johnson et al. (2015) demonstrates that using multiple trades rather than the single large order can bias empirical studies. Using the granularity of our data set, we accurately group multiple trade reports into a single transaction even when a single large order trades against orders with different prices. Finally, we can accurately determine the best bid and ask prices immediately before a trade. [8]

We prepare the data for our analysis as follows. First, we reconstruct the limit order book. When a series of consecutive trades are reported with the same time stamp, trading in the same direction and occurring from the same broker we treat them as one trade. Second, for every trade, we determine the prevailing bid and ask immediately prior to the trade. Third, we only include trades during continuous trading hours and omit all trades occurring at the opening and closing auction. Omitting the opening price avoids contamination from large price moves caused by the arrival of overnight news. This preparation process leaves a sample of more than 140,000 observations for the most actively trade stock (CBA) and over 29,000 observations for the least actively traded stock in the sample (AMP).

4 Estimation and Results

In this section, we study the relationship between trade size and permanent price impact. Using BHP as a representative stock, we estimate the *nonlinear VAR* using 5 lags and 8 trade size quantiles.[9] Table 2, Columns 1 to 8 present the first lag coefficient estimates for the trade quantile equations and Column 9 shows the first lag coefficients for the quote revision equation. For example, Column 1 reports the first lag coefficients corresponding to the smallest trade quantile equation defined as

$$x_t^1 = \sum_{i=1}^{5} \delta_i^k r_{t-i} + \sum_{i=1}^{5} c_{1i}^1 x_{t-i}^1 + \sum_{i=1}^{5} c_{2i}^1 x_{t-i}^2 + + \sum_{i=1}^{5} c_{3i}^1 x_{t-i}^3 + ... + + \sum_{i=1}^{5} c_{8i}^1 x_{t-i}^8 + \epsilon_t^1 \tag{7}$$

Specifically, the first lag coefficient for the smallest trade quantile variable is 0.17, the first lag coefficent for the second smallest trade quantile variable is 0.06, while the first lag coefficient for the largest trade quantile variable is 0.01. Lastly, the first lag coefficient for the quote revision variable is 65.10. Consistent with earlier findings, our lag coefficient estimates for all trade size quantile equations are positive and significant, indicating that signed trades exhibit a strong positive autocorrelation (see Hasbrouck (1991a), Engle and Patton (2004) and Dufour and Engle (2000))

[6]see Theissen (2001)

[7]Johnson et al. (2015) demonstrate that using data sources like NYSE daily TAQ and the consolidated tape a single large order cannot be recovered from a sequence of reported trades.

[8]In contrast, Lee and Ready (1991) suggest using a 5 second delay between the reported quote and trade price to reflect the fact that the mechanism reporting quotes to the tape is not time synchronized with the trade-reporting mechanism. Over time, the required delay suggested to apply has changed to reflect the markets progression in technology but the problem still persists.

[9]The results are robust to choice of parameters. 8 trade size quantiles provided a balance between parameters to estimate and flexibility to capture nonlinearity.

Insert Table 2 here

We also observe evidence of autocorrelation in trade size. Table 2, Column 1, reports the first lag coefficients for the small sized trade equation ($k = 1$) and we find that the coefficients *decrease* monotonically between $c_1^k = 0.17$ and $c_8^k = 0.01$. In contrast, Column 8 reports the first lag coefficients for the large sized trade equation, where we observe the coefficients *increase* monotonically between $c_1^k = 0.01$ and $c_8^k = 0.08$. This result demonstrates that a small trade is more likely to be followed by another small trade, while a large trade is more likely to precede a large trade.

Next, we estimate the permanent price impact of different sized trades using the long term impulse response function of the *nonlinear VAR* we propose. Table 3 reports the permanent price impact for different sized trades for each individual stock. Our results confirm a nonlinear relationship exists between permanent price impact and trade size for all sample stocks. Specifically, we show that permanent price impact increases with trade size monotonically at a decreasing rate. The results demonstrates that in the modern market enivornment, permanent price impact, as a function of trade size, is positive, increasing and concave, consitent with Hasbrouck (1991a) and Engle and Patton (2004).

Insert Table 3 here

We now compare our *nonlinear VAR* results with those obtained using the *trade sign VAR* or *trade size VAR.* Figure 2 plots the relationship between permanent price impact and trade size for the representative stock, BHP, using the three different VAR specifications. Estimates of the permanent price impact as a function of trade size using the *nonlinear VAR* are concave. In contrast, when we use the *trade sign VAR* estimates of the permanent price impact produce a step function. This step function reflects the assumption that trades of differing size have equal permanent price impact. Alternatively when we use the *trade size VAR* we obtain a linear relationship between permanent price impact and trade size.

Figure 2 demonstrates that when we use *trade sign* as the trade variable in the VAR model defined by Equation 1 then a step function occurs. This step function partially captures the nonlinear relationship between price impact and trade size and may explain why, despite economic considerations, much of the literature uses *trade sign* as the trade variable when estimating a VAR model.

Insert Figure 2 here

4.1 Comparing Multiple market participants or trade venues

In the previous section, we establish a nonlinear relationship between permanent price impact and trade size. In this section, we investigate the economic implications if these nonlinear relationships are not captured. In Section 4.1.A, using an example of two equally uninformed investors who differ only in their average trade size, we estimate their respective price impacts using first, the traditional *two participant trade sign (trade size) VAR* and second, the *two participant nonlinear VAR* we propose. Additionally, the VAR framework is misspecified if a trader uses knowledge of the limit order book in their trading decisions (Hasbrouck, 1991a). In Section 4.1.B, we introduce an additional variable, namely the order book depth imbalance, in a multi-factor VAR framework to investigate the potential pitfalls of model misspecification.

A. Two uninformed market participants with differing trade size

We estimate the permanent price impact of trades from two uninformed market participants, who are similar in all aspects except for trade size. One participant is the large size trader, *ls*, while the other is the small size trader, *ss*. To form the two market participants, we allocate all trades with a trade size above the median *ls*, and all trades below the median trade size to *ss*. Barclay et al. (2003) conclude that small trades are more informed than large trades and thus, one potential issue is that *ss* could be more informed than *ls*. To reduce the possibility that one trader is more informed than the other trader, we randomly select 25% of all trades from *ls* and re-categorise these trades to *ss*. Similarly, we randomly allocate 25% of *ss* trades to *ls*. This random allocation of trades also removes the discrete jump from *ss* to *ls* at the median trade size. Using this construction method, we argue both market participant categories are equally uninformed. As such, we expect trades of equal signed size should have the same permanent price impact for both market participants.

Figure 3, plots the estimated impulse response functions for the two market participants estimated via the *two participant trade sign VAR* (Panel A) or the *two participant trade size VAR* (Panel B) for the representative stock, BHP. Using the *two participant trade sign VAR* we find that the permanent price impact for market participant ls is larger than the impact caused by a trade initiated by market participant ss (Panel A). In contrast, using the *two participant trade size VAR* we show that the permanent price impact for market participant ls is now smaller than the impact caused by a trade initiated by market participant ss (Panel B).[10]

Insert Figure 3 here

Figure 3 demonstrates that our conclusions change depending on whether we use a *trade size* VAR or a *trade sign* VAR. Using a *trade size* VAR, we find that the trader with the smaller order size has the larger price impact. Conversely, using a *trade sign* VAR, we show that the larger order size trader has the larger price impact. These conflicting results arise because of the nonlinear relationship between the trade variable and the trades permanent price impact.

These results have significant implications for research inferences when comparing trades originating from different market participants or trades across different market venues. For example, when comparing retail traders or HFT, who have a small average trade size, against institutional trader, who have a large average trade size, incorrect inferences could occur if the VAR model is misspecified. A similar concern occurs if we compare the permanent price impact of trades across different exchanges, or between lit and dark trading venues, due to a documented difference in

[10]The data used to estimate the impulse response function presented in Panel B is the same data used when estimating the impulse response functions presented in Panel A.

average trade size among these venues (see Comerton-Forde and Putnins (2015) and Barclay et al. (2003))

To overcome the assumption of linearity, we estimate our proposed *two participant nonlinear VAR* using 5 lags and 8 trade size quantiles for BHP for the large and small trader. We estimate the permanent price impact for trades of different size initiated by different market participants by observing the long term impulse response function for quote revisions corresponding to a shock to the respective trade size quantiles.

Figure 4 depicts the permanent price impact caused by trades of different size for the two market participants. We observe that both market participants have similar price impact functions. This result is consistent with the expectation that both market participants have the same permanent price impact for comparable size trades.[11] Accordingly, when comparing the permanent price impact between two market participants we advocate the use of the *two participant nonlinear VAR* to ensure correct inferences are made.

Insert Figure 4 here

B. Multi-factor VAR system

The existing literature suggests that permanent price impact is not just a function of trade size but multiple dimensions. For example, Dufour and Engle (2000) demonstrate that the time duration between trades affects the permanent price impact of trades, while Engle and Patton (2004) show the volume available at the best bid and ask impacts quote revisions. This the permanent price impact of a trade is a function of both the trader's private information and their knowledge of the order book.

We estimate different VAR specifications for two uninformed market participants that differ by the market conditions under which they trade. Specifically, one participant (ai) executes their orders *against* the direction of the imbalance in the order book. In contrast, the other market participant (wi) executes their orders *with* the direction of the imbalance in the order book. Brogaard et al. (2014) and Goldstein et al. (2016) demonstrate that HFT execute trades in the direction of the order book imbalance. Furthermore, Engle and Patton (2004) show that the depth imbalance at the time of a trade can affect the magnitude of the permanent price impact. As such, when comparing the private information content of a trade between market participants the VAR model should also control for the public information contained within the limit order book.

We define imbalance of the order book using the same metric as Goldstein et al. (2016)

$$DI = \frac{\sum_{i=1}^{5} Vol_{Bid,i} - \sum_{i=1}^{5} Vol_{Ask,i}}{\sum_{i=1}^{5} Vol_{Bid,i} + \sum_{i=1}^{5} Vol_{Ask,i}} \tag{8}$$

where $Vol_{Bid,i}$ is the total volume on the bid at price level i and $Vol_{Ask,i}$ is the total volume on the ask at price level i. We allocate all trades that are buyer initiated when $DI > 0$ and

[11] We provide more rigorous testing of these findings in Appendix B.

seller initiated when $DI < 0$ to the market participant that trades with the imbalance, wi. All remaining trades in the sample trade against the imbalance and we assign to market participant ai. For similar reasons presented in Section 4.1.A, we randomly select 25% of all trades and those in the ai (wi) category, we re-categorised to be in wi (ai). We argue that neither category is more informed about the fundamental value of the stock via private information. Accordingly, we expect that both market participants should have the same permanent price impact attributed to private information.

To estimate the permanent price impact of the two market participants we use two VAR specifications common in the literature. The first is the *two participant trade sign VAR* and the second includes the order book imbalance prior to the trade as an additional explanatory variable. [12]

Figure 5, Panel A depicts the estimated impulse response functions for our representative stock for both market participants using the *two participant trade sign VAR*. The market participant that trades *with* the order book imbalance has a larger permanent price impact than the market participant that trades *against* the order book imbalance. This difference is because the impulse response function captures both private information and the trader's ability to time trades around publicly available information contained in the limit order book.

To distinuish between the private and public information content of a trade, we include the order book imbalance as an additional explanatory variable in the VAR. Engle and Patton (2004) propose a similar VAR, which we define as follows:

$$
\begin{aligned}
r_t &= \sum_{i=1}^{10} \alpha_i r_{t-i} + \sum_{i=1}^{10} \beta_i x^e_{t-i} + \sum_{i=0}^{10} \zeta_i x^m_{t-i} + \omega_t DI + \epsilon_{1,t} \\
x^e_t &= \sum_{i=1}^{10} \delta_i r_{t-i} + \sum_{i=1}^{10} \phi_i x^e_{t-i} + \sum_{i=1}^{10} \nu_i x^m_{t-i} + \psi_t DI + \epsilon_{2,t} \\
x^m_t &= \sum_{i=1}^{10} \vartheta_i r_{t-i} + \sum_{i=1}^{10} \psi_i x^e_{t-i} + \sum_{i=0}^{10} \eta_i x^m_{t-i} + \Gamma_t DI + \epsilon_{3,t}
\end{aligned} \tag{9}
$$

Figure 5, Panel B plots the estimated cumulative impulse response functions for our representative stock for both market participants using the VAR model defined by Equation 9. In contrast to our earlier results using the *two participant trade sign VAR*, we find that the market participant who trades against the imbalance has the bigger permanent price impact.[13]

Insert Figure 5 here

These results demonstrate how conflicting conclusions occur depending on the VAR specification and the corresponding assumptions. When using the *two participant trade sign VAR*, we

[12] The results for the *two participant trade size VAR* are similar to those reported for the *two participant trade sign VAR*.

[13] We provide more rigorous testing of these findings in Appendix C

assume the shape of the order book has no effect on quote revisions. In contrast, Equation 9 assumes that the order book imbalance affects quote revisions linearly. However, it is possible that permanent price impact has a non nonlinear relationship with depth imbalance.

Previously, we demonstrate that quote revisions are not linear in trade size and that we can specify a VAR model that captures the nonlinear relationship by discretizing trades into size buckets. Extending this concept a VAR model that captures nonlinear terms in trade size and order book imbalance could contain discretized trade size and order book imbalance dummies along with their interactions. However, the number of coefficient estimates for such a model has quadratic growth. Accordingly, the number of parameters to estimate becomes infeasible for large amounts of discretization.

5 The RL Model

In the previous section, we demonstrate how nonlinear relationships between trade size, market conditions and permanent price impact can lead to incorrect conclusions about the information content of a trade. Furthermore, we show estimation difficulties arise when we include a nonlinear factor in a VAR system. In this section, we present a model for calculating permanent price impact using reinforcement learning (RL) which does not suffer the same concerns as a VAR model. RL is a type of machine learning founded in the computer science literature, and is well suited to computing the permanent price impact of a trade. RL assumes the data is Markov, but is otherwise unrestrictive, which allows a large degree of flexibility as to what explanatory variables should be included and can capture any nonlinear relationships that exist. Easley et al. (2013) introduce RL to the field of market microstructure and demonstrate its efficacy for modelling the order book and advocate its use when solving the problem of buying a specified volume of shares in a specified amount of time.

Typically, for an RL model, an agent has knowledge of its current environment via an indication of the current state it lies within, s. The agent then chooses an action, a, which changes the state of the environment. This state transition rewards the agent through a scalar instant reward, r. The agent should choose actions that increase the long run sum of the instant rewards received for each action. We present an algorithm that enables the agent to learn the optimal action to take at each state, which is known as Q learning. We adapt the Q learning algorithm to model permanent price impact of a trade. Formally an RL model consists of:

- A discrete set of environment states, S.
- A set of possible actions, A.
- A reward function $R : S \times A \to \Re$
- A state transition function $T : S \times A \to \Pi(S)$

The environment states include any variables which reflect the current environment. The set of actions reflect the available options an agent can make. The reward function is the expected instantaneous reward as a function of the current state and action made. Lastly, $\Pi(S)$ maps transitions from states to probabilities. $T(s, a, s')$ represents the probability of making a transition from state s to state s' using action a. For RL, the optimal value of a state is typically computed as

$$V^*(s) = \max_\pi \mathbb{E}\big(\sum_{t=0}^{\infty} \gamma^t r_t\big) \tag{10}$$

where r_t is the immediate reward at time t and $0 < \gamma < 1$ is a discount factor. $V^*(s)$ is the expected infinite discounted sum of reward the agent receives if they start in that state and executes the optimal policy defined by π moving forward. This optimal value is unique and can be defined as the solution to the simultaneous equations

$$V^*(s) = \max_a \Big(R(s,a) + \gamma \sum_{s' \in S} T(s,a,s')V^*(s')\Big), \forall s \in S \tag{11}$$

Given the optimal value function, the optimal policy can be specified as

$$\pi^*(s) = \operatorname*{argmax}_a \Big(R(s,a) + \gamma \sum_{s' \in S} T(s,a,s')V^*(s')\Big) \tag{12}$$

For every state there are possible actions that can be made, thereby resulting in $S \times A$ possible experience tuples, $\langle s, a \rangle$. For every experience tuple there is an associated value $Q^*(s, a)$, which is the expected infinite discounted sum of reward that the agent will gain if it takes action a in state s, then follows the optimal policy path. Using (11) we note that $Q^*(s, a)$ can be expressed recursively as

$$Q^*(s,a) = R(s,a) + \gamma \sum_{s' \in S} T(s,a,s') \max_{a'}(Q^*(s',a')) \tag{13}$$

Equation (13) is the basis for Q learning and the associated Q learning rule is

$$Q_{t+1}(s_t,a_t) = Q_t(s_t,a) + \alpha\Big(R(s_t,a_t) + \gamma \max_a(Q_t(s_{t+1},a)) - Q(s_t,a)\Big) \tag{14}$$

where α is the learning rate. The Q learning rule is a value iteration update and it is shown that under appropriate conditions the Q values will converge to Q^* with probability 1 (see Watkins and Dayan (1992) and Tsitsiklis (1994)).

We now demonstrate how this RL technique can be used to measure permanent price impact. Initially, consider the agent to be a trader and the environment states reflect the market environment. For the simplest case we assume there is only one state, which is analogous to the market being open. Brogaard et al. (2014) and Goldstein et al. (2016) show traders make trades dependant on market conditions, which are reflected in the shape of the order book. The RL model can be extended to account for these differing market conditions by allowing for several states.

In the simplest form of the model, the actions available to the agent or trader are to buy and sell, which allows us to estimate the permanent impact of a buy order and sell order. Because RL is highly flexible, we can extend the model to many possible actions. For example, if we want to compute the permanent price impact of a large order versus a small order, we can include four possible actions (large buy order, small buy order, large sell order and small sell order) into the model. To model permanent price impact using the RL model we make the instant reward $R(s, a)$ the change in log midpoint caused by taking action a that transitions the data from s at time t to s' at time $t + 1$.

The permanent price impact of a trade implies the trader makes only one trade, but the RL model outlined above assumes the trader makes more trades in the future. To restrict the trader to only one action, we modify (13) to

$$Q^*(s, a) = R(s, a) + \gamma \sum_{s' \in S} T(s, a, s'\overline{a}) Q^*(s', \overline{a}) \tag{15}$$

The right hand term of equation (13) is simplified to $\gamma \sum_{s' \in S} T(s, a, s'\overline{a}) Q^*(s', \overline{a})$. This simplification means the trader only has one trade, and cannot make any more trades in the future. The trader's action will transition the market into a new state s', once in that state, other participants reactions $\overline{a}$ occur with probability $T(s, a, s'\overline{a})$. With this simplification, the Q-learning rule is now

$$Q_{t+1}*(s_t, a_t) = R(s_t, a_t) + \gamma \sum_{s' \in S} T(s, a, s'\overline{a}) Q_t(s', \overline{a}) \tag{16}$$

The recursive nature of this model is its inherent strength, we make no assumptions about the data generating process and the data is left free to explain itself.

To simultaneously estimate the difference in permanent price impact caused by multiple market participants, we modify equation (16) to

$$Q_{t+1}(s_t, a_t, c_t) = R(s_t, a_t, c_t) + \gamma \sum_{s' \in S} T(s_t c_t, a_t, s'\overline{a}\overline{c}) Q_t(s', \overline{a}, \overline{c}) \tag{17}$$

Here the experience tuple now includes which market participant, c, made the trade/action at time t. If there are two market participants c_1, c_2, and both have no private information about the future price of the stock, then they should cause the same permanent price impact for trades with the same size and market conditions or $Q(s_i, a_j, c_1) = Q(s_i, a_j, c_2)$ for all i and j. Conversely, if the traders in c_1 are more informed than traders in c_2, we expect buys (sells) by traders in c_1 to have a larger positive (negative) permanent price impact than equivalent buys (sells) by traders in c_2, or more formally, $Q(s_i, a_j, c_1) > Q(s_i, a_j, c_2)$ for buys and $Q(s_i, a_j, c_1) < Q(s_i, a_j, c_2)$ for sells. Thus, any discrepancy between respective Q values for differing market participants can be attributed to private information. To formally test whether one market participant contributes more to price discovery due to private information, we estimate the following regression:

$$Q(s_i, a_j, c_1) = \beta Q(s_i, a_j, c_2) + \epsilon \tag{18}$$

If $\beta > 1$ then market participant c_1 has more private information and has a larger permanent price impact for trades with the same size and market conditions than c_2. The converse is true if $\beta < 1$. Lastly, if $\beta = 1$ then neither market participant has more private information.

Consider a market where one participant trades very large volumes while the other participant trades small volumes. Economic considerations suggest the large trader will cause more permanent price impact for their trades. However, is this larger impact due only to the trade size differential or knowledge of private information? The RL model enables us to disentangle these two effects. Equation (18) determines if two market participants differ in private information while observing differences in the transition probabilities, $T(s_t c_t, a_t, s'\overline{ac})$, attributes difference in price impact to size or market conditions.

In a similar spirit to Hasbrouck (1991b), we propose a simple and intuitive metric, which measures the proportion of total price movement attributable to a certain action or trade. The price contribution is computed as

$$PC_a = \frac{\sum_{s' \in S} Q(s, a) n(s, a)}{\kappa} \tag{19}$$

Where $n(s, a)$ is the number of times action a is made while in state s and κ is the total amount of price movement observed in the sample computed as $\sum_{s' \in S, a \in A} |R(s, a)|$.

We can extend Equation (19) to measure the proportional price contribution of each market participant. For example, in the case of two market participants the permanent contribution to price made by participant c_i is

$$PC_{c_i} = \frac{\sum_{s' \in S, a \in A} Q(s, a, c_i) n(s, a, c_i)}{\kappa} \tag{20}$$

For a model in which trades are the only actions available, all of the actions or trades in the sample must account for all movement in price in the sample. This implies that $\kappa = \sum_{s' \in S, a \in A} |R(s, a)| = \sum_{s' \in S, a \in A} Q(s, a) n(s, a)$ [14]

[14] Empirically slight deviations may occur due to discretization of the data

Equations (19) and (20) provide examples of how the RL model attributes what proportion of price discovery comes from various sized trades, or different market participants. It is clear that determining what proportion of price discovery is attributed to a specific action or market participant is the ratio of price movements that action or market participant causes to the total price movement in the sample. This concept can be extended to additional actions available to the market such as limit order submission or cancellations.

6 Estimation Results

6.1 Base RL Model

In this section we compare a trade's permanent price impact with estimated via a VAR model and the proposed RL model. We specify a baseline RL model with two possible actions, $A \in (+1, -1)$, where $a_t = +1$ for buyer initiated trades and $a_t = -1$ for seller initiated trades. The model assumes there is only one market state (ie. open) to parallelsHasbrouck's (1991a) VAR model using *trade sign* as the trade indicator variable. Table 4, Columns 1 and 2 report the RL models estimated permanent price impact, for a sell and buy, respectively. For each sample stock, we also report the permanent price impact estimated using the *trade sign VAR*. Table 4 also provides the corresponding 95% confidence interval for the estimated permanent price impact obtained using the *trade sign VAR*. For all sample stocks, except one, the RL's estimated permanent price impact lies within the 95% confidence interval of the *trade sign VAR* permanent price impacts. These results demonstrate the baseline RL model and Hasbrouck VAR model yield similar results.

Insert Table 4 here

6.2 Accounting for trade size

In this section we demonstrate how the RL model can capture the nonlinear dynamics between trade size and permanent price impact and compare the RL estimates with the *nonlinear VAR* estimates. To capture any nonlinearities, we expand the number of possible actions for the RL model from buy and sell to a larger set. Accordingly, we estimate the RL model using fifteen possible actions, $A \in (1, ..., 15)$, where $a_t = i$ if the directional trade size falls in decile i. Thus, we assign a large sell order as $a_t = 1$, while we assign a large buy order as $a_t = 15$. Similar to the model estimated in Section 6.1, we select only one market state (open).

Figure 6 plots the RL models estimated permanent price impact for each trade decile for our representative stock, BHP. For comparison, Figure 6 also depicts the permanent price impact estimated via the *nonlinear VAR*. Observing Figure 6, we see that both models yield a similar relationship between trade size and price impact. One discrepancy between the VAR model and the RL model is the smoothness of the relationship between trade size and price impact; the relationship estimated by the RL model is smoother relative to the relationship estimated using the VAR model. Economic considerations suggest that price impact should increase monotonically with trade size. The RL model captures this monotonic relationship for all trade deciles. In contrast, the VAR model estimates appear more noisy and does not recover a smooth monotonic relationship.

Insert Figure 6 here

6.3 Two uninformed participants; one that strategically trades using public information

In Section 4.1.B we demonstrate the econometric concerns of using a VAR model when comparing two market participants who condition their trades on public information contained in the limit order book. In this section we test if the RL method suffers similar concerns. Similar to the RL model estimated in Section 6.2, we model ten possible actions or trade sizes, $A \in (1, .., 10)$ where $a_t = i$ if the directional trade size falls in decile i. However, we now specify 5 environment or market states $s \in (1, .., 5)$, where $s_t = i$ if the depth imbalance just before trade t lies in quantile i. In addition we allow for two market participant categories such that $C \in (wDI, aDI)$. These extensions results in 100 (10*5*2) experience tuples $\langle s, a, c \rangle$.

We assign trades for the two market participant categories using the same method found in Section 4.1.B. Similarly, we argue both market participants should have the same permanent price impact for a trade of similar size under comparable market conditions.

For our representative stock, BHP, Figure 7, Panel A plots the RL model's estimated permanent price impact for every experience tuple for the market participant which trades *against* the imbalance. For comparison, Figure 7, Panel B, contains the equivalent permanent price impact for every experience tuple for the participant who trades *with* the imbalance. Visually, it appears the two participants have equal permanent price impacts for trades of the same size and market conditions. We confirm this notion by obtaining a coefficient not statistically different from 1 when we regress all impact tuples for the participant who trades *with* the imbalance against the equivalent tuple for the participant who trades *against* the imbalance. We repeat this process for all sample stocks and obtain the same result. A coefficient not statistically different from one suggest that both participants cause equivalent price impacts under trades of similar conditions. These results are in contrast to the VAR models tested in Section 4.1.B where the market participants were concluded to have different permanent prices impacts and the participant which caused the larger permanent price impact was dependent on the VAR specification.

Insert Figure 7 here

From Figure 7 we note there is approximately 3.5 times more cross sectional variation in permanent price impact across the depth imbalance metric than the cross sectional variation across trade size. These results suggest that the shape of the order book has significantly more influence on permanent price impact than the size of a trade. This finding reinforces the notion that price discovery does not just occur via market orders but also via limit orders and is consistent with Brogaard et al. (2016). Furthermore, given HFT's trade conditional on the shape of the order book (see Goldstein et al. (2016) and Brogaard et al. (2014)) empirical research conducted on HFT price discovery of should also control for the current state of the order book.

7 Why is price impact nonlinear with size?

Mendelson and Tunca (2004) provide one theoretical motivation for why a nonlinear relationship between permanent price impact and trade size exists, by suggesting agents are endogenous in the volumes they trade. Specifically, agents time their trades by trading large volumes when the market is liquid, resulting in small price impacts per unit traded, and small volumes when the market is illiquid, resulting in large price impacts per unit traded. In this section we use the RL model to test if traders behave consistently with Mendelson and Tunca (2004).

We estimate the RL model with 16 possible actions, $A \in (1, .., 16)$, where $a_t = i$ if the directional trade size falls in decile i. This means if $a_t = 1$ then the actions is an extremely large sell order, conversely if $a_t = 16$ we have an extremely large buy. We also include 5 environment or market states, $s \in (1, .., 5)$ which reflect the amount of liquidity available at the time of the market order submission, with $s_t = 1(5)$ representing the lowest (highest) liquidity available quintile.[15]

Figure 8 presents the permanent price impact caused for the differing states of market liquidity available and demonstrates that small orders in an illiquid market cause a larger permanent price impact than large orders in a liquid market. Specifically, when the market is illiquid, a small market buy order for less than 10 shares has a permanent price impact of approximately 2 basis points. In contrast, when the market is highly liquid, a large market buy order of approximately 4000 shares has a permanent price impact of approximately 1 basis point. This finding supports Mendelson and Tunca's (2004) hypothesis that a trade executing against the thin side of the order book has more permanent price impact than a trade executing against the thick side of the order book.

Insert Figure 8 here

We test whether trades change their behavior based on the amount of liquidity available in the market as predicted by Mendelson and Tunca (2004). Consistent with their predictions Table 5 shows that traders are more likely to trade large size during periods of high liquidity and trade small size during periods of low liquidity. Our results show a trader is 13.7 times more likely to submit a large market order when the market is extremely liquid compared to when it is extremely illiquid. Similarly, when the market is extremely illiquid, traders are 9.8 times more likely to submit a small order than a large order.

Insert Table 5 here

Collectively, these results demonstrate the permanent price impact of a trade is conditional on the current liquidity available and that agents trade endogenously depending on the prevailing market conditions. Specifically, traders choose to time their order submission to minimise permanent price impact by submitting large orders when liquidity is high and small orders when liquidity is low. These results support Mendelson and Tunca's (2004) hypothesis that traders alter their behaviour due to the time varying nature of market liquidity and provide one possible explanation for why permanent price impact is non linear with trade size.

[15] we compute the amount of liquidity available as the total volume available at the best bid (ask) for market sell (buy) orders

8 Price contribution of various market actions

As another application of the RL model, we extend the findings of Brogaard et al. (2016), who show that limit order submissions and cancelations, in addition to market orders, cause a permanent price impact. We extend the RL model to include six possible actions: market order submission (buy or sell), limit order submission (submitted to best bid or best ask) and limit order cancelation (from best bid or best ask).

Table 6 reports the permanent price impact attributable to each of the possible actions. Similar to Brogaard et al. (2016) we find both the submission and cancellation of limit orders convey information resulting in permanent price impact. Consistent with economic considerations, buy limit order submissions and sell limit order cancellations result in a positive permanent price impact. Conversely, when a limit sell order is submitted or limit buy order cancelled then a negative permanent price impact occurs. Notably, the submission of a market order has the largest permanent price impact of around 0.69 basis points. The submission of a limit order causes permanent price impact of roughly 0.13 basis points which is approximately 19% the magnitude of impact caused by a market order. Cancellations are shown to have the smallest permanent price impact of 0.085 basis point which is only 12% the size of a market orders permanent price impact. These findings are comparatively similar to those reported by Brogaard et al. (2016).

Insert Table 6 here

Extending Brogaard et al. (2016), we use Equation 19 to estimate the proportion of total price discovery attributable to each action. Table 6 shows that market orders contribute 57.2% to price discovery, while limit order submissions and cancelations contribute 30.1% and 12.7%, respectively. While the average permanent price impact of a limit order (0.13 bps for submissions and 0.085 bps for cancelations) is significantly smaller than that of a market order (0.69 bps), limit order submissions and cancelations contribute close to 45% to overall price discovery. The proportionately higher market-wide contribution is due to the large number of limit order submissions and cancellations, relative to market orders submissions.

9 Conclusion

We demonstrate the importance of trade size when investigating the dynamics of trades and quote revisions. Specifically, we show that the VAR model of Hasbrouck (1991a) is misspecified when we use *trade sign* or *signed trade size* as the trade variable. This misspecification can lead to incorrect conclusions. Accordingly, we address this misspecification by generalizing the Hasbrouck (1991a) VAR model to capture any nonlinear relationships that may exist between quote updates and trade size. We demonstrate that this generalized VAR model performs well under basic specifications, arriving at correct inferences.

However, all VAR models suffer from estimation issues when their specification becomes complex. As such, we propose using a RL algorithm to estimate the permanent price impact of a trade. For low specifications, the VAR model and RL model yield similar permanent price impact estimates. In contrast, when the model specifications become more complex, the RL model outperforms the VAR model. Using the RL algorithm, we demonstrate the shape of the order book has strong effects on the dynamics of trades and quote revision. Notably, a trade executing against the thin side of the order book has a higher permanent price impact than a trade which executes against the thick side of the order book. However, the higher permanent price impact is due to a trader's knowledge of public limit order book information, rather than their access to private information. We demonstrate that traders condition their trades on the liquidity available: agents trade large volumes during periods of high liquidity and small volumes when the market is illiquid. This finding is consistent with Mendelson and Tunca (2004) and provides one explanation for the nonlinear dynamic between quote revisions and trade size.

References

Barclay, M. J., Hendershott, T., and McCormick, D. T. (2003). Competition among Trading Venues: Information and Trading on Electronic Communications Networks. *Journal of Finance*, 58(6):2637–2666.

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

Brogaard, J., Hendershott, T., and Riordan, R. (2016). Price discovery without trading: Evidence from limit orders.

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

Collin-Dufrense, P. and Fos, V. (2015). Do prices reveal the presence of informed trading? *The Journal of Finance*, 70(4):1555–1582.

Comerton-Forde, C. and Putnins, T. J. (2015). Dark trading and price discovery. *Journal of Financial Economics*, 118(1):70 – 92.

Dufour, A. and Engle, R. F. (2000). Time and the price impact of a trade. *The Journal of Finance*, 55(6):2467–2498.

Easley, D., de Prado, M. L., and O'Hara, M. (2013). *High Frequency Trading - New Realities for Traders Markets and Regulators*. Oxford University Press, USA.

Engle, R. F. and Patton, A. J. (2004). Impact of trades in an error-correction model of quote prices. *Journal of Financial Markets*, 7(6):1–25.

Goldstein, M., Kwan, A., and Philip, R. (2016). High Frequency Trading.

Hasbrouck, J. (1991a). Measuring the information content of stock trades. *The Journal of Finance*, 46(1):179–207.

Hasbrouck, J. (1991b). The summary informativeness of stock trades: An econometric analysis. *The Review of Financial Studies*, 4(3):571–595.

Hasbrouck, J. (2007). *Empirical Market Microstructure : The Institutions, Economics, and Econometrics of Securities Trading: The Institutions, Economics, and Econometrics of Securities Trading*. Oxford University Press, USA.

Huberman, G. and Stanzl, W. (2004). Price manipulation and quasi-arbitrage. *Econometrica*, 72(4):1247–1275.

Johnson, H., McInish, T., and Upson, J. (2015). Orders versus trades on the consolidated tape.

Jones, C. M., Kaul, G., and Lipson, M. L. (1994). Transactions, volume, and volatility. *Review of Financial Studies*, 7(4):631–651.

Kyle, A. S. (1985). Continuous auctions and insider trading. *Econometrica*, 53(6):1315–1335.

Mendelson, H. and Tunca, T. I. (2004). Strategic trading, liquidity, and information acquisition. *Review of Financial Studies*, 17(2):295–337.

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

Theissen, E. (2001). A test of the accuracy of the lee/ready trade classification algorithm. *Journal of International Financial Markets, Institutions and Money*, 11(2):147 – 165.

Tsitsiklis, J. N. (1994). Asynchronous stochastic approximation and q-learning. *Machine Learning*, 16(3):185–202.

Watkins, C. J. C. H. and Dayan, P. (1992). Q-learning. *Machine Learning*, 8(3):279–292.

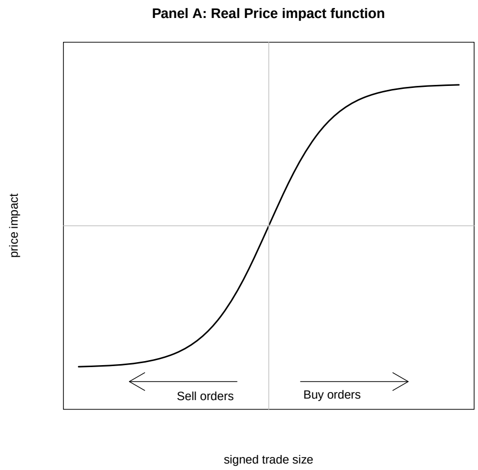

(i) Permanent price impact function of ***both*** *traders*

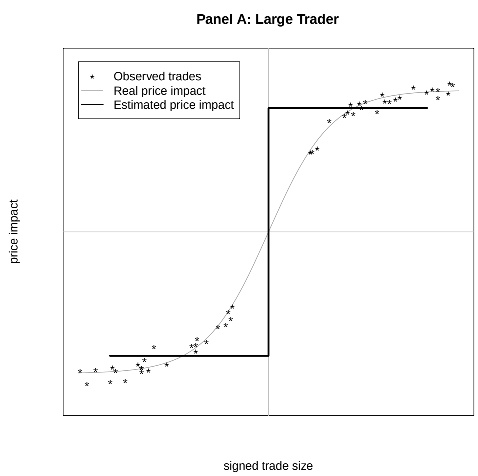

(ii) Estimated permanent price impact of a ***large*** *trader using a* ***trade sign*** *VAR*

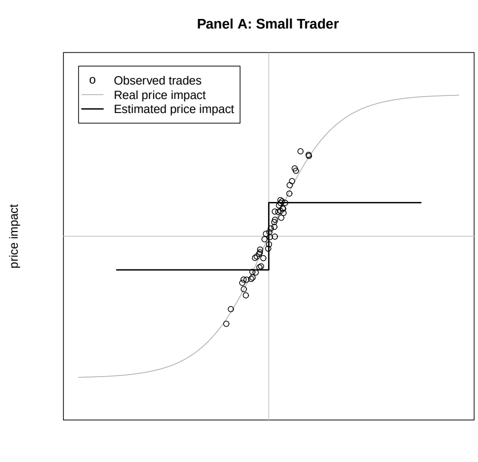

(iii) Estimated permanent price impact of a ***small*** *trader using a* ***trade sign*** *VAR*

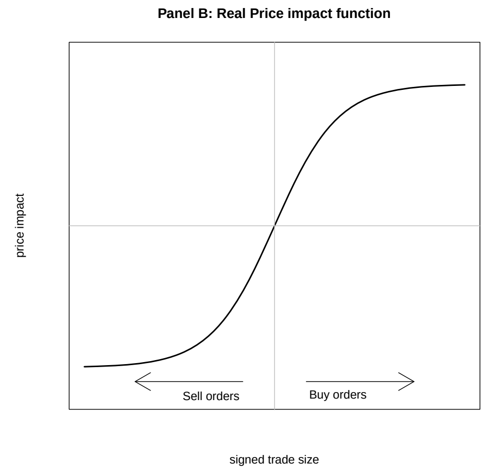

(iv) Permanent price impact function of ***both*** *traders*

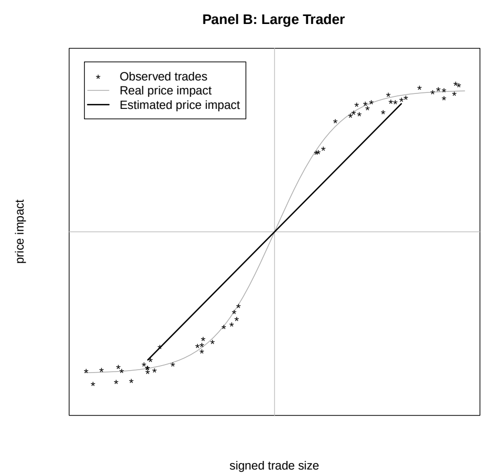

(v) Estimated price impact of a ***large*** *trader using a* ***signed trade size*** *VAR*

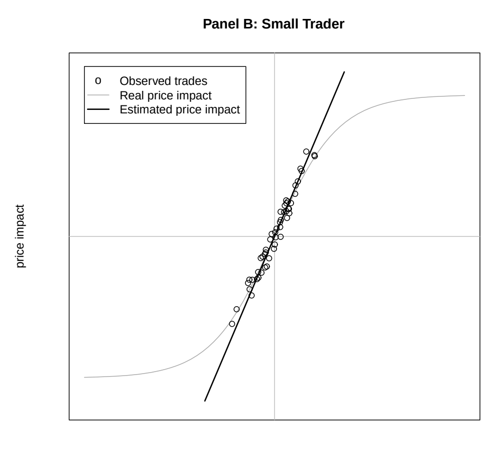

(vi) Estimated price impact of a ***small*** *trader using a* ***signed trade size*** *VAR*

Figure 1: Plots a nonlinear price impact function of two market participants that differ by their order sizes. Panel A depicts the average impact of each participant if trade direction is considered. Panel B shows the estimated impact of each participant if impact is assumed to be linear with size.

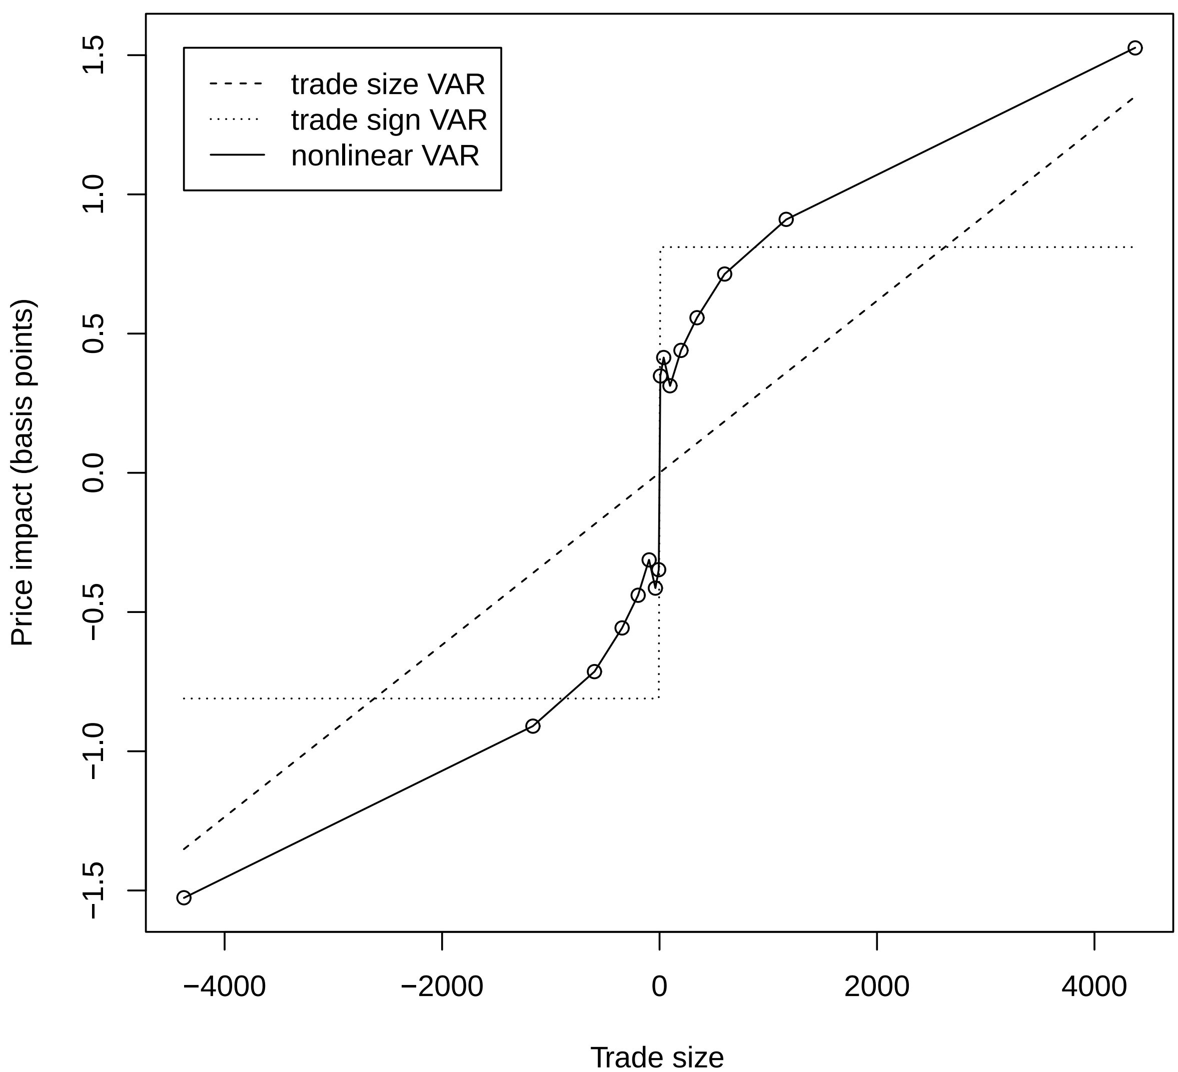

Figure 2: Depicts the permanent price impact of different sized trades for BHP. Three different specifications of a VAR model are estimated and the long term impulse response function is the permanent price impact. Model 1 is estimated using Equation 1 where the trade indicator variable is signed trade size. Model 2 is estimated using Equation 1 where the trade indicator variable is the trade sign. Model 3 is estimated using Equation 4

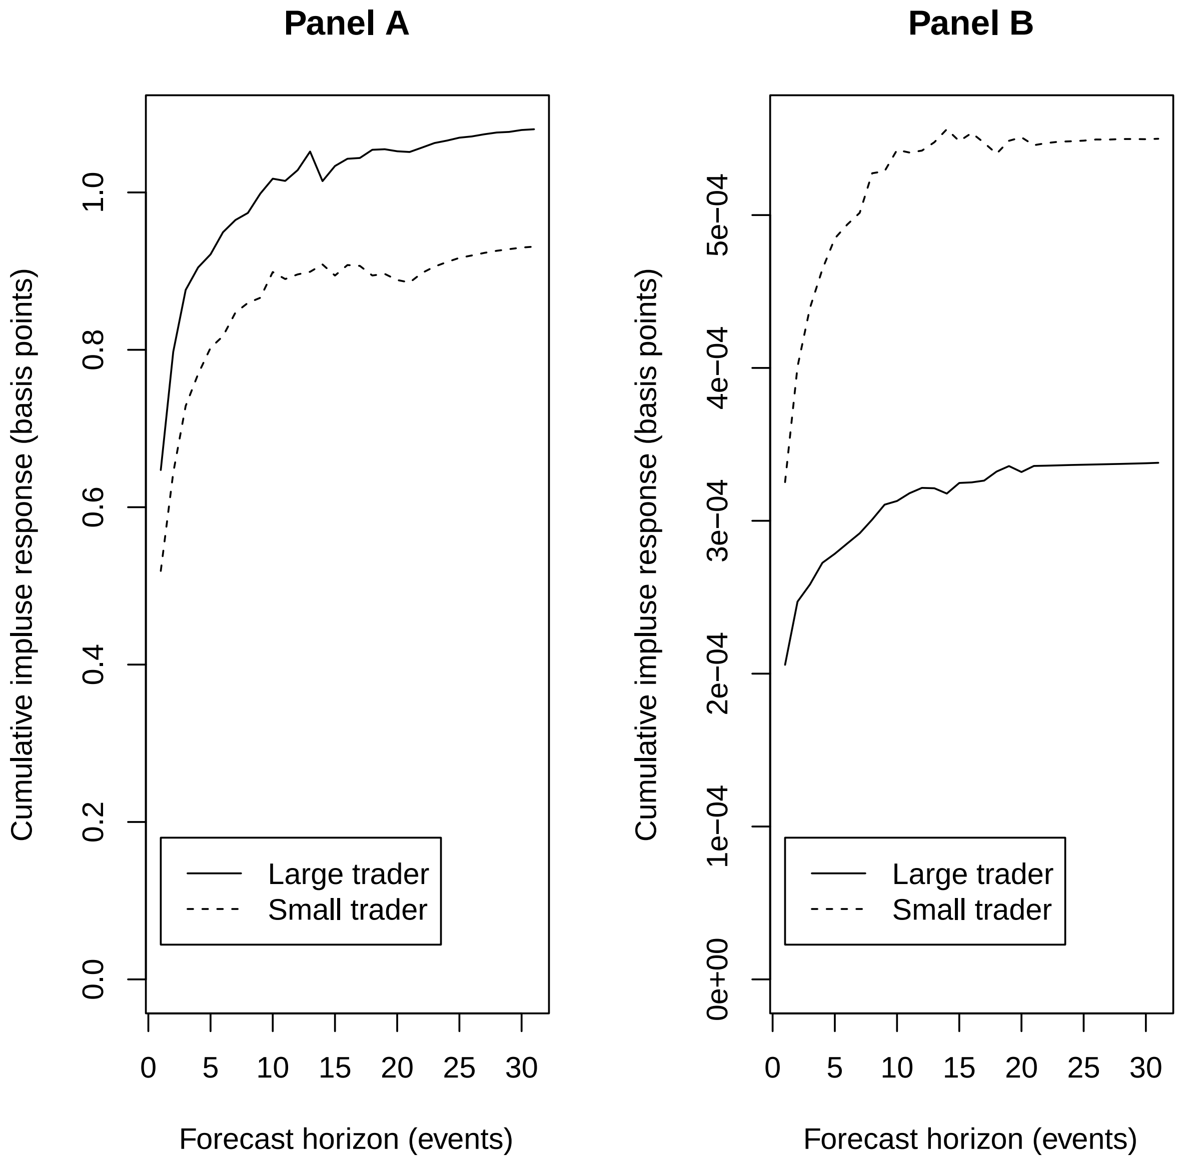

Figure 3: Plots the cumulative impulse response function for two uninformed market participants. One participant trades a larger volume per order than the other participant. Panel A plots the cumulative impulse response function for the Hasbrouck (1991a) VAR model when the trade variable is trade sign, while Panel B depicts the cumulative impulse response functions when the trade variable is the signed trade size.

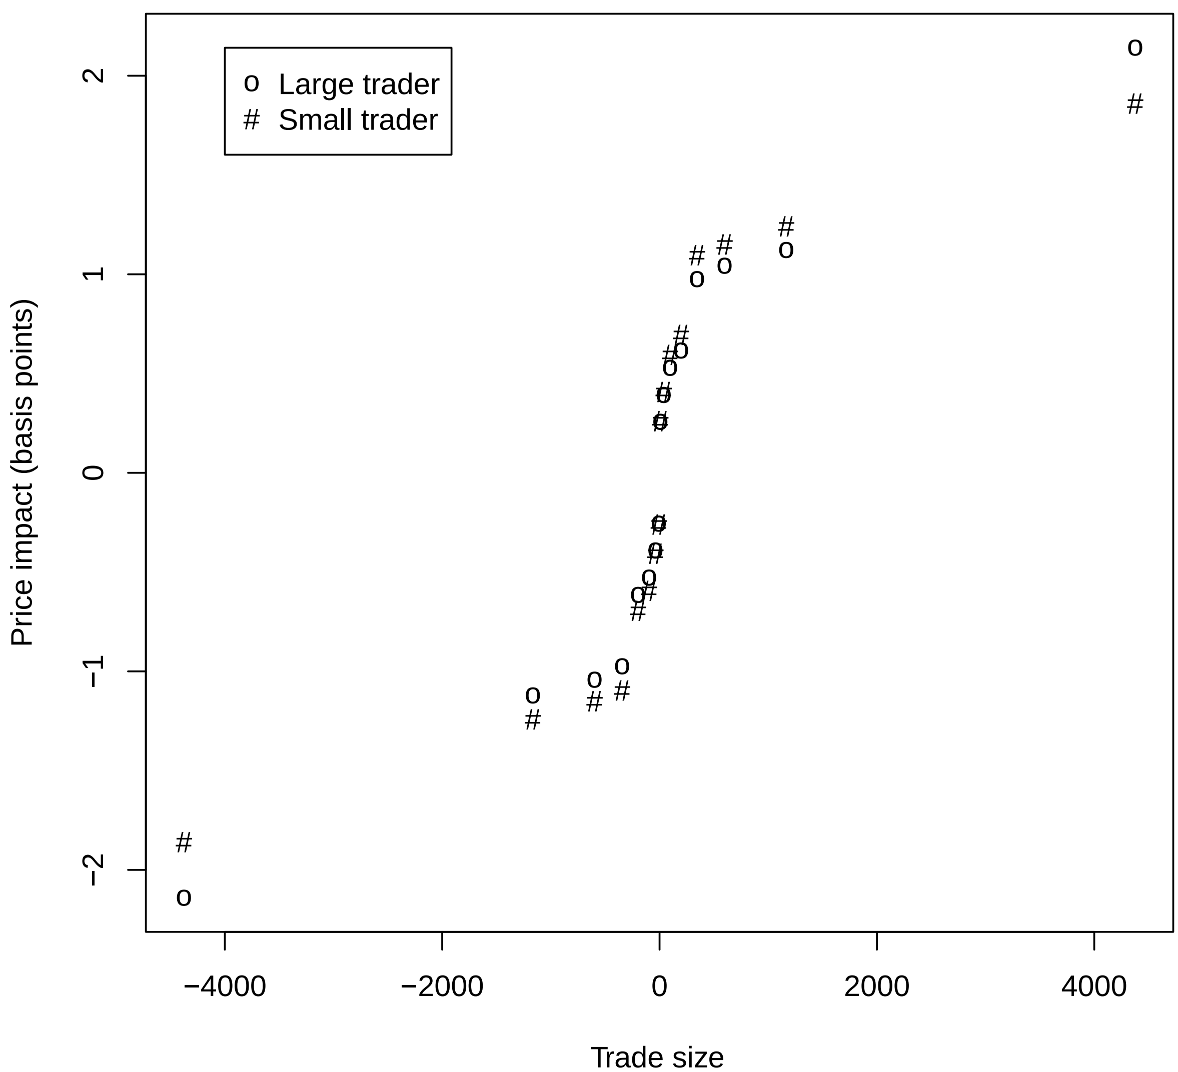

Figure 4: Plots the permanent price impact caused for trades of various size estimated using the VAR model defined by Equation 6. The permament price impact is estimated for two uninformed market participants. One participant has a small average trade size, whereas the other market participant has a large average trade size.

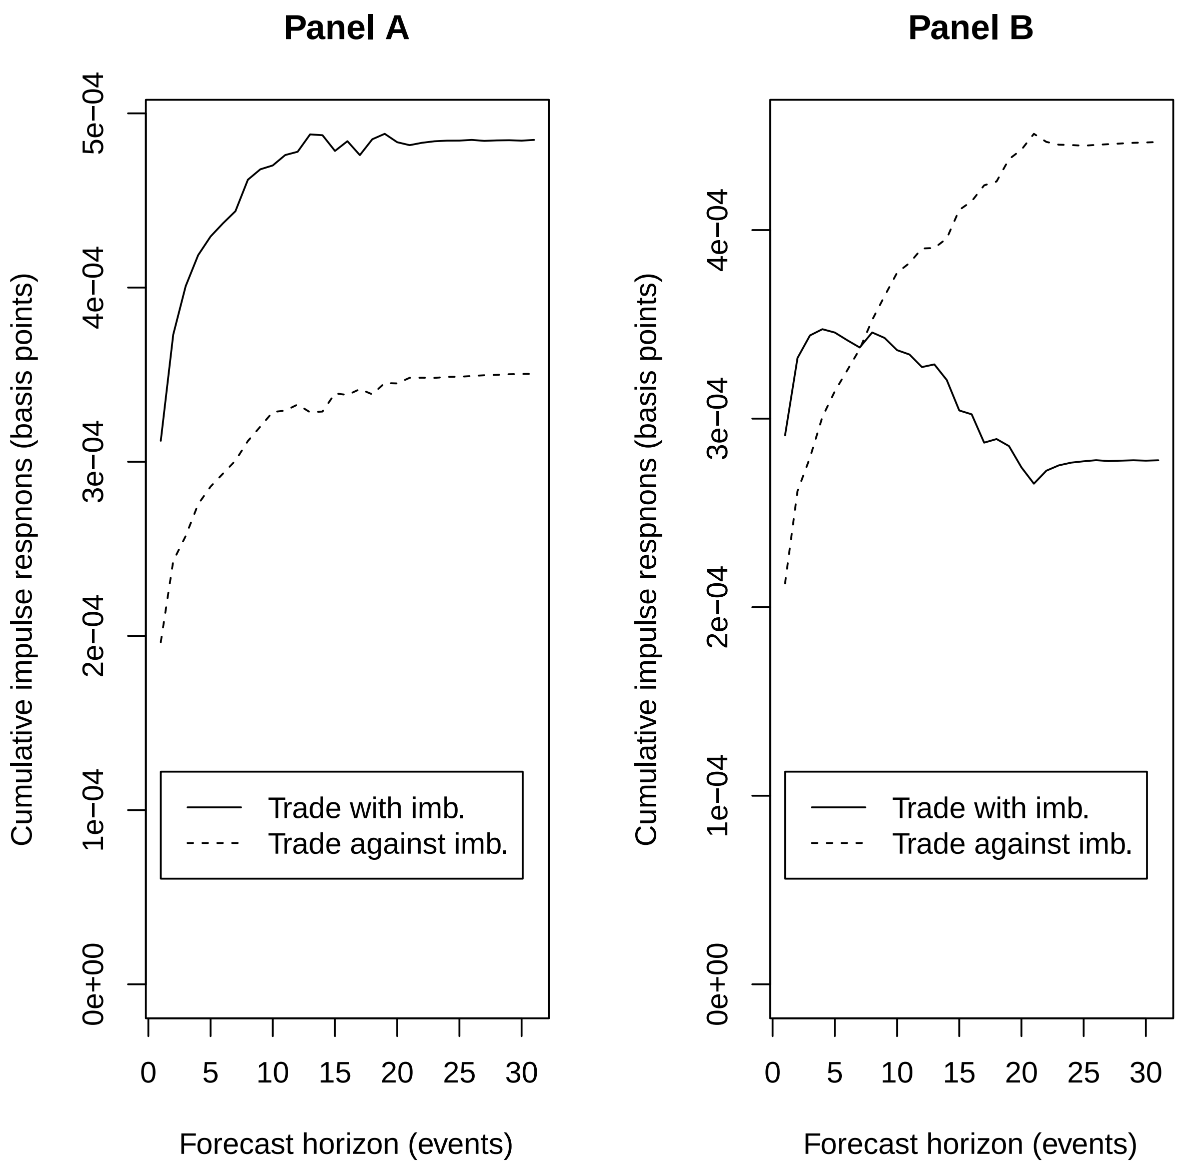

Figure 5: Plots the cumulative impulse response function for two market participants. One participant trades in the direction of the order book imbalance while the other does not. Panel A plots the cumulative impulse response function for the Hasbrouck (1991a) VAR model. Panel B depicts the impulse response functions for the VAR model which includes the order book imbalance as an exogenous variable.

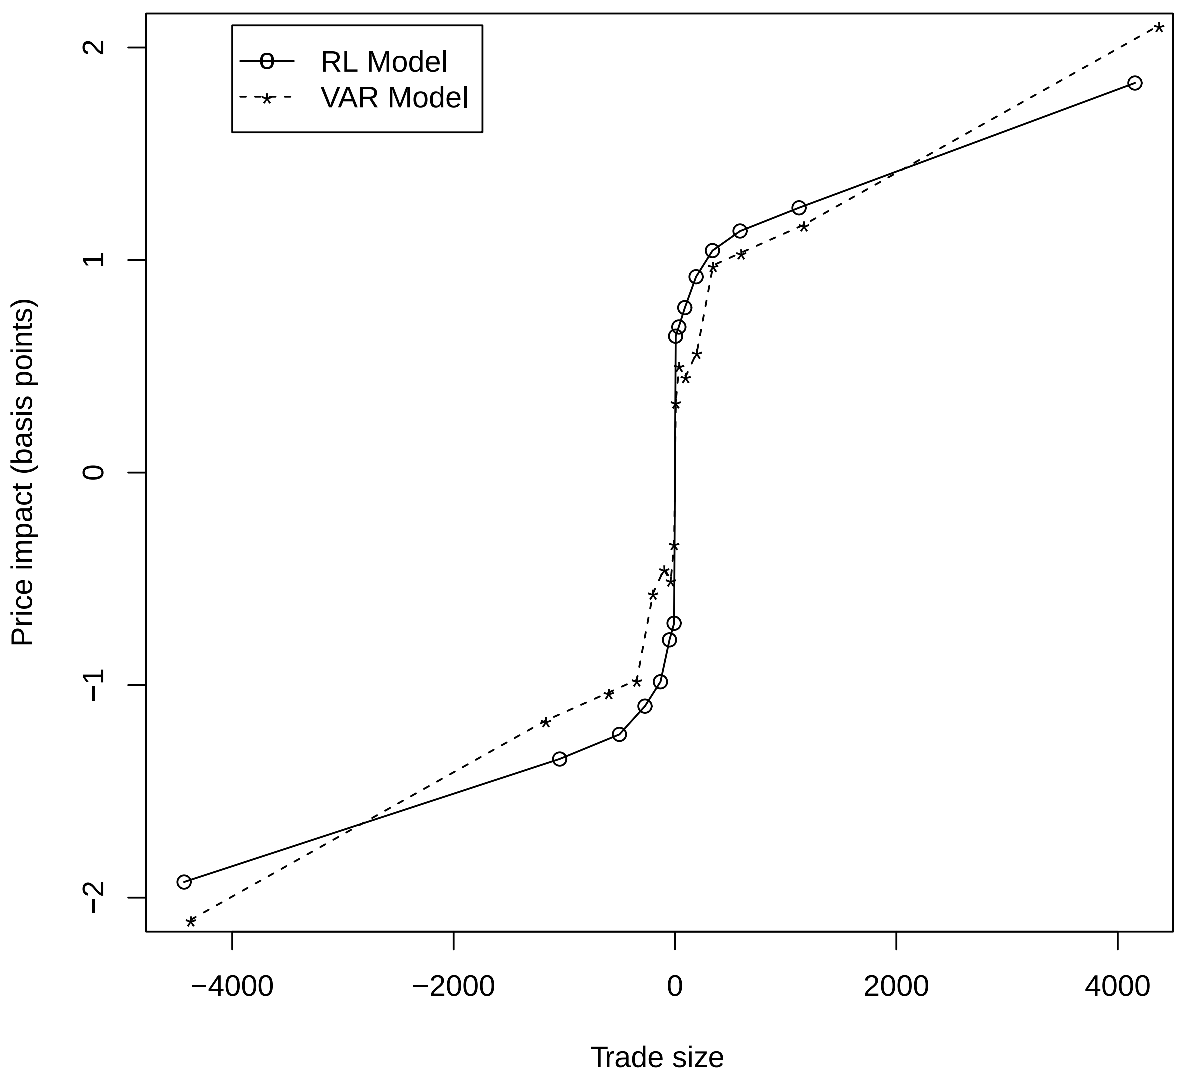

Figure 6: Depicts the RL models estimated permanent price impact for all trade volume decile. The figure also plots the permanent price impact for various trade volume deciles estimated using the VAR model defined by Equation 4.

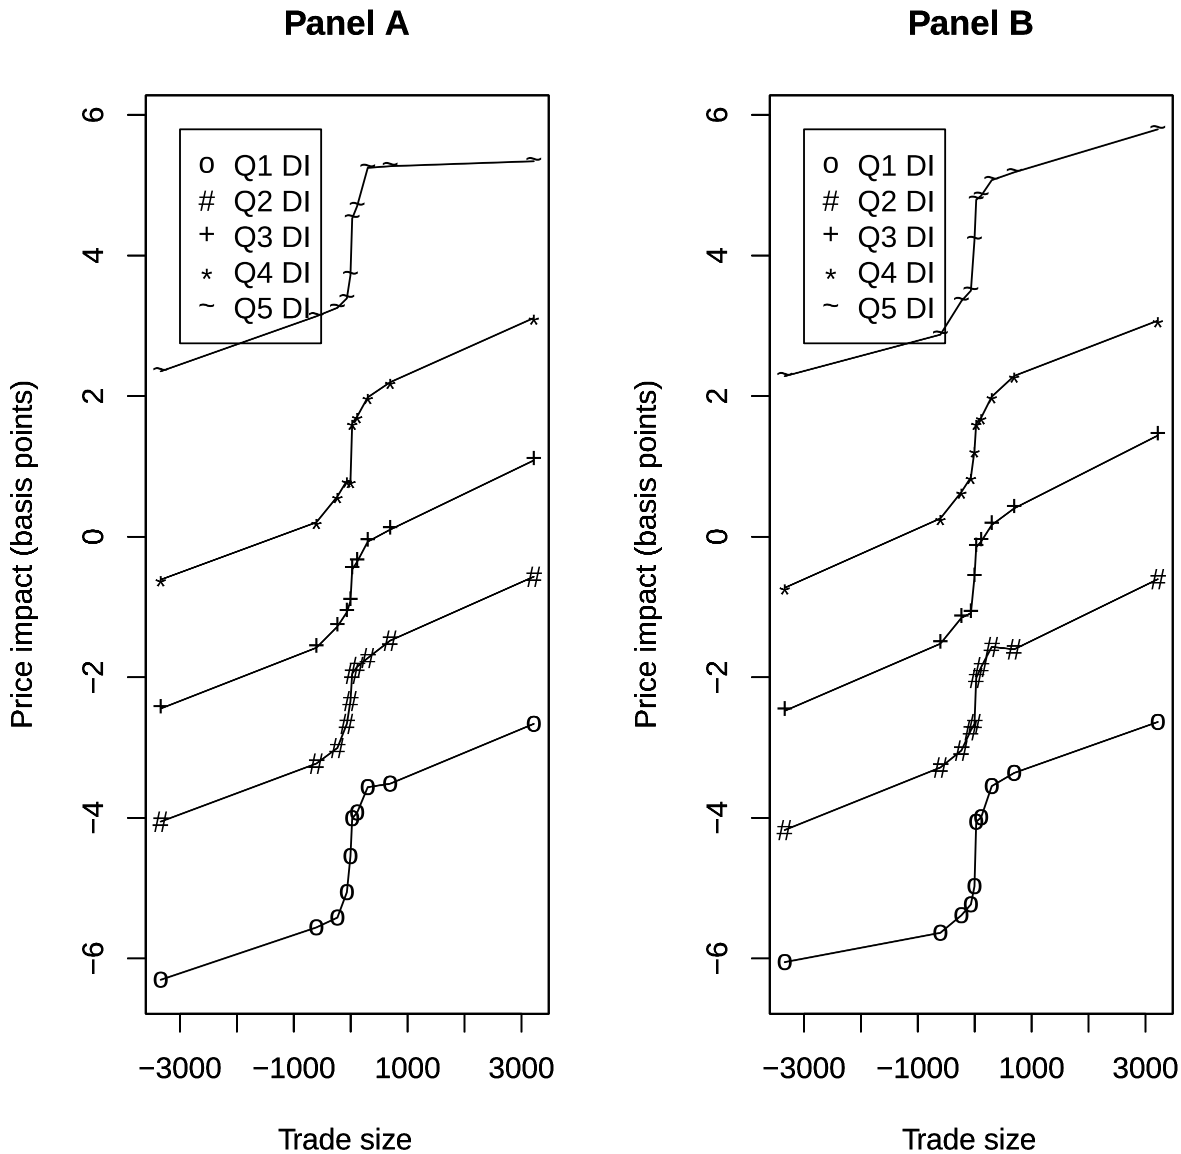

Figure 7: Plots the permanent price impact, estimated using the RL model, for various trade sizes and order book imbalances. Q1 DI is for trades that occur when the depth imbalances fall into the smallest quintile, or the most negative depth imbalances. Q5 DI is for trades that occur when the depth imbalances are in the top quintile. Panel A plots the permanent price impact for different trade sizes under different market states for the market participant that trades against the imbalance. Panel B plots the permanent price impact for different trade sizes under different market states for the market participant that trades with the imbalance.

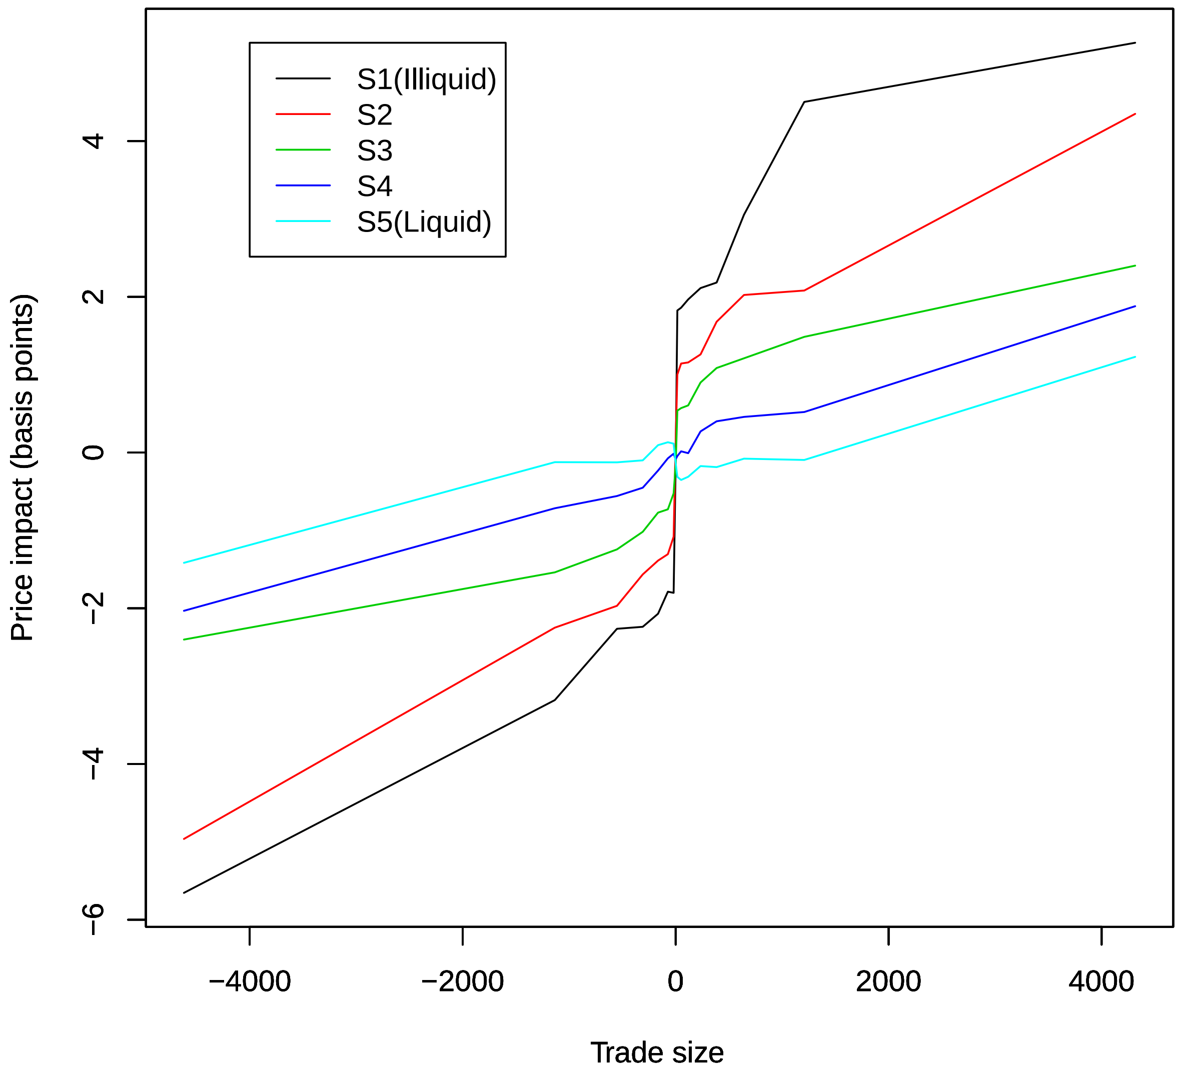

Figure 8: Plots the permanent price impact for various trade sizes under different states of market liquidity. S1 is for trades that occur when the liquidity available is in the smallest quintile. S5 is for trades that occur when the liquidity available is in the top quintile.

Stock	N	Average Price	Average Spread	Average Size	Median Size	Max Size
BHP	119200	28.31	0.009	865.34	265	150000
CBA	140106	85.45	0.012	229.93	90	24433
WBC	90564	33.38	0.010	651.35	200	73099
ANZ	85584	32.02	0.010	680.91	240	50000
NAB	99321	34.00	0.010	677.51	220	59445
TLS	41001	6.21	0.009	5786.35	583	614862
WOW	96460	30.69	0.009	514.05	169	40990
RIO	121586	56.51	0.011	223.45	94	21360
CSL	108930	85.70	0.004	111.23	49	28333
WPL	115734	34.81	0.007	343.69	111	50000
NCM	73039	12.82	0.009	844.82	201	704467
QBE	51599	10.75	0.009	1111.34	205	76853
AMP	29447	5.52	0.009	2533.54	278	263718
SUN	46214	14.15	0.009	743.62	190	124613
ORG	42472	10.88	0.010	988.33	224	43869
BXB	39948	10.49	0.010	950.93	174	70198
MQG	102431	58.67	0.013	165.33	66	18000
STO	60930	7.57	0.009	1594.68	290	198500
AMC	63146	12.95	0.009	1057.49	260	72259
IAG	32832	6.32	0.009	1989.37	296	199127

Table 1: Summary Statistics for the trades and quotes data of the 20 stocks with the highest market capitilisation. The sample contains 40 trading days from January 1, 2015 to Feburary 28, 2015. Average price is the average executed price in dollars, Average Spread is the average distance betwen the bid and ask reported in dollars. Size is trade size reported in shares, with the average, median, maximum and minimum reported.

	$x_t^{k=1}$	$x_t^{k=2}$	$x_t^{k=3}$	$x_t^{k=4}$	$x_t^{k=5}$	$x_t^{k=6}$	$x_t^{k=7}$	$x_t^{k=8}$		r_t
$c_{1,1}^k$	0.17	0.05	0.04	0.03	0.02	0.03	0.02	0.01	b_1	-0.00
	60.19	**18.16**	**12.86**	**9.21**	**8.22**	**8.80**	**7.53**	**4.43**		**-37.69**
$c_{2,1}^k$	0.06	0.10	0.05	0.04	0.04	0.03	0.02	0.01	b_2	-0.00
	21.72	**33.69**	**17.99**	**14.53**	**12.13**	**11.26**	**8.25**	**4.64**		**-28.14**
$c_{3,1}^k$	0.03	0.05	0.10	0.06	0.04	0.04	0.03	0.02	b_3	-0.00
	11.71	**17.75**	**34.70**	**19.11**	**13.55**	**12.49**	**9.47**	**7.57**		**-24.87**
$c_{4,1}^k$	0.03	0.04	0.05	0.09	0.05	0.05	0.04	0.02	b_4	-0.00
	9.09	**14.61**	**17.48**	**29.96**	**18.79**	**15.65**	**14.23**	**7.44**		**-19.96**
$c_{5,1}^k$	0.02	0.04	0.05	0.06	0.15	0.08	0.07	0.06	b_5	-0.00
	7.49	**12.43**	**15.69**	**20.39**	**53.40**	**28.73**	**23.06**	**19.64**		**-11.62**
$c_{6,1}^k$	0.01	0.03	0.04	0.05	0.08	0.10	0.08	0.06	b_6	-0.00
	4.30	**11.08**	**13.64**	**18.84**	**28.03**	**35.41**	**25.88**	**19.68**		**-11.60**
$c_{7,1}^k$	0.01	0.02	0.03	0.05	0.06	0.06	0.07	0.05	b_7	-0.00
	4.53	**6.19**	**9.48**	**15.67**	**22.03**	**21.19**	**25.21**	**16.35**		**-8.66**
$c_{8,1}^k$	0.01	0.01	0.01	0.02	0.05	0.05	0.05	0.08	b_8	0.00
	2.98	**2.99**	**2.36**	**8.37**	**16.38**	**15.93**	**17.13**	**26.90**		**16.04**
δ_1^k	-65.10	-52.67	-61.09	-56.72	-62.12	-44.22	-30.84	-10.68	α	-0.18
	-13.36	**-11.06**	**-12.61**	**-11.75**	**-12.99**	**-9.18**	**-6.39**	**-2.21**		**-58.21**

Table 2: First lag coefficient estimates and t-statistics (in bold) for the VAR model defined by Equation 4. x_t^k is +1 (-1) if a trade is buyer (seller) initiated and its trade size falls in quantile k and 0 otherwise; r_t is the change in the natural logarithm of the midpoint that follows a trade at time, t. Column 1 represents the first lag coefficients for the smallest trade quantile equation($k = 1$). Column 8 reports the first lag coefficients for the largest trade quantile equation ($k = 8$). Column 9 reports the first lag coefficients for the quote revision equation. The sample covers January 1, 2015 to Feburary 28, 2015 and is for our representative stock, BHP

		q1	q2	q3	q4	q5	q6	q7	q8
AMC	Price impact	0.63	0.73	0.80	0.66	0.16	0.89	0.93	2.57
	Trade size	7	36	91	191	356	619	1343	5709
AMP	Price impact	2.63	2.76	2.75	2.20	2.04	2.26	2.70	4.23
	Trade size	8	40	114	252	605	1255	3123	17191
ANZ	Price impact	0.37	0.45	0.40	0.53	0.62	0.71	0.92	1.42
	Trade size	7	33	85	179	293	490	968	3810
BHP	Price impact	0.35	0.41	0.31	0.44	0.56	0.71	0.91	1.53
	Trade size	7	38	97	198	345	598	1166	4375
BXB	Price impact	1.16	1.16	1.25	1.14	1.03	1.34	1.71	2.91
	Trade size	7	32	78	151	310	594	1408	6312
CBA	Price impact	0.23	0.28	0.26	0.37	0.35	0.46	0.60	0.85
	Trade size	4	18	39	72	105	163	313	1245
CSL	Price impact	0.41	0.42	0.49	0.58	0.70	0.65	0.87	1.23
	Trade size	3	11	25	43	69	97	159	687
IAG	Price impact	1.79	2.15	2.21	2.17	1.77	1.88	2.34	3.60
	Trade size	9	47	106	224	509	1022	2637	16822
MQG	Price impact	0.35	0.44	0.48	0.53	0.59	0.77	0.90	1.25
	Trade size	3	12	28	52	89	121	207	789
NAB	Price impact	0.21	0.35	0.23	0.36	0.43	0.53	0.83	1.30
	Trade size	7	36	80	161	279	472	909	3567
NCM	Price impact	0.81	0.99	1.00	0.73	1.02	1.27	1.77	3.21
	Trade size	6	26	71	156	280	498	1036	4311
ORG	Price impact	1.03	1.30	1.36	1.22	1.20	1.56	2.19	3.52
	Trade size	8	34	86	171	343	663	1449	5360
QBE	Price impact	0.86	1.03	1.17	1.02	1.18	1.45	1.84	3.20
	Trade size	8	33	89	182	381	712	1598	6890
RIO	Price impact	0.35	0.44	0.48	0.60	0.61	0.73	0.97	1.34
	Trade size	3	17	42	76	110	163	299	1126
STO	Price impact	1.34	1.51	1.54	1.46	1.20	1.70	2.43	4.55
	Trade size	9	36	84	203	443	893	2091	9542
SUN	Price impact	0.80	0.85	0.90	0.74	0.77	0.97	1.27	2.32
	Trade size	6	32	81	162	319	588	1293	5680
TLS	Price impact	2.80	2.69	3.03	2.14	2.45	2.36	2.67	3.54
	Trade size	14	89	225	480	1145	2423	5938	41728
WBC	Price impact	0.36	0.43	0.40	0.45	0.67	0.73	0.89	1.41
	Trade size	8	36	83	167	275	460	891	3403
WOW	Price impact	0.36	0.44	0.48	0.42	0.58	0.74	0.99	1.59
	Trade size	5	24	65	123	208	348	715	2845
WPL	Price impact	0.45	0.51	0.59	0.58	0.69	0.91	1.17	1.83
	Trade size	4	20	51	99	170	278	548	1948

Table 3: Reports the permanent price impact for trades of different size quantiles ($q1, \ldots, q8$). The average trade size for each quantile is reported under the estimated price impact. Estimates are made for all sample stocks using using Equation 4 with x lags.

	RL SELL	RL BUY	VAR	VAR 95% LOWER	VAR 95% UPPER
AMC	1.73	2.07	1.14	0.51	1.74
AMP	1.70	2.04	1.54	-0.64	3.66
ANZ	1.04	0.97	0.94	0.70	1.17
BHP	1.06	1.04	0.78	0.45	1.08
BXB	1.92	1.87	1.64	0.64	2.65
CBA	0.65	0.65	0.67	0.56	0.77
CSL	1.16	1.40	1.20	1.01	1.39
IAG	1.50	1.66	1.41	-0.42	3.14
MQG	0.86	1.00	0.94	0.75	1.13
NAB	0.87	0.87	0.68	-0.27	1.58
NCM	2.17	2.18	1.73	1.13	2.28
ORG	2.09	2.23	2.10	1.18	3.02
QBE	2.17	2.12	1.52	0.73	2.29
RIO	0.98	1.01	0.96	0.78	1.12
STO	3.04	3.19	2.07	1.00	3.14
SUN	1.35	1.58	1.00	0.40	1.60
TLS	1.10	1.08	1.00	-0.47	2.48
WBC	1.06	0.99	0.98	0.74	1.21
WOW	1.07	1.08	0.98	0.75	1.22
WPL	1.31	1.39	1.26	0.99	1.52

Table 4: Columns 1 and 2 reports the estimated permanent price impact for a sell and buy respectively when the RL models is used for estimation. Column 3 reports the permanent price impact that we estimate when using the trade sign VAR. Columns 4 and 5 report the corresponding lower and upper confidence bounds, respectively. The available action for the RL model is buy (+1) and sell (-1) making the specification similar to the VAR model used for columns 3 to 5.

	I (Illiquid)	II	III	IV	V (Liquid)	Liquid/Illiquid
1 (Large order)	0.003	0.003	0.037	0.041	0.041	**13.727**
2	0.003	0.048	0.030	0.022	0.021	**6.984**
3	0.013	0.045	0.026	0.024	0.022	**1.697**
4	0.032	0.022	0.021	0.020	0.019	**0.606**
5	0.044	0.023	0.022	0.021	0.021	**0.475**
6	0.042	0.020	0.020	0.021	0.023	**0.542**
7	0.035	0.019	0.022	0.025	0.027	**0.769**
8 (Small order)	0.029	0.020	0.021	0.026	0.027	**0.920**
Small/Large	**9.895**	**6.955**	**0.552**	**0.627**	**0.663**	

Table 5: Reports the probability a trader will submit a trade of certain size under five different states of market liquidity. The Table reports ratio's of the probability of a trader submitting an order under liquid versus illiquid market conditions in the bold column. Similarly, the Table reports ratios of the probability of a trader submitting a small order versus a large order under different states of liquidity in the bold row

	Market Order		Submit limit order		Cancel limit order	
	Buy	Sell	Buy	Sell	Buy	Sell
Nobs.	64093	57050	171782	154590	103757	97769
Price impact (bp)	0.67	-0.71	0.13	-0.14	-0.10	0.07
Contribution (%)	29.5	27.7	15.1	14.9	7.3	5.2

Table 6: Reports metrics for six different actions a trader makes. The actions are to submit a market order to buy or sell, to submit a limit order to the best bid or offer or cancel an order from the best bid or offer. Metrics reported are the number of observations each action occurs. The permanent price impact an actions causes and the contribution each action makes to total price movement

10 Appendix

A Polynomial VAR

We explain the difficulties that arise from a VAR specification which contains quadratic terms by providing a simple example. Suppose we have a VAR defined as

$$\begin{bmatrix} x_t \\ x_t^2 \end{bmatrix} = \begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix} \begin{bmatrix} x_{t-1} \\ x_{t-1}^2 \end{bmatrix} + \begin{bmatrix} \epsilon_{1,t} \\ \epsilon_{2,t} \end{bmatrix} \tag{21}$$

it follows that $a_{21}x_{t-1} + a_{22}x_{t-1}^2 + \epsilon_{2,t} = (a_{11}x_{t-1} + a_{12}x_{t-1}^2 + \epsilon_{1,t})^2$. For some large negative realization of $\epsilon_{2,t}$ such that $\epsilon_{2_t} < -(a_{21}x_{t-1} + a_{22}x_{t-1}^2)$ there will be no solution. While Hasbrouck (1991a) addresses this problem by replacing x_t^2 with $sign(x_t)|x_t|^2$ difficulties still arise with estimation of the impulse response function. Once again, lets assume a simple VAR defined by

$$\tilde{x}(t) = \mathbf{A}\tilde{x}(t-1) + \tilde{\epsilon}(t) \tag{22}$$

where $\tilde{x}(t)$ is a 2×1 vector at time t, $\mathbf{A}$ is a 2×2 matrix of coefficients and $\tilde{\epsilon}(t)$ is a 2×1 vector of residuals. Equation 22 takes the VMA representation of

$$\tilde{x}(t) = \sum_{j=0}^{\infty} \mathbf{A}^j \tilde{\epsilon}(t-j) \tag{23}$$

To obtain the cumulative impulse response function we now consider moving $\tilde{\epsilon}(t-K)$ to $\tilde{\epsilon}(t-K) + \Delta$ and consider $\tilde{x}(t+\Delta) - \tilde{x}(t)$ where $\tilde{x}(t+\Delta)$ refers to the realization of the model when $\tilde{\epsilon}(t-K)$ is replaced by $\tilde{\epsilon}(t-K) + \Delta$. It is clear that $\tilde{x}(t+\Delta) - \tilde{x}(t) = \mathbf{A}^K\Delta$.

Consider the case when

$$\tilde{x}(t) = \begin{bmatrix} x_t \\ x_t^2 \end{bmatrix} \qquad \Delta = \begin{bmatrix} \Delta_1 \\ \Delta_2 \end{bmatrix} \tag{24}$$

To implement an impulse response perhaps one may set $\Delta_2 = 0$ which yields the following

$$\begin{bmatrix} x_t \\ x_t^2 \end{bmatrix} = \begin{bmatrix} \epsilon_{1,t} \\ \epsilon_{2,t} \end{bmatrix} + \mathbf{A} \begin{bmatrix} \epsilon_{1,t-1} + \Delta_1 \\ \epsilon_{2,t-1} \end{bmatrix} + \mathbf{A}^2 \begin{bmatrix} x_{t-2} \\ x_{t-2}^2 \end{bmatrix} \tag{25}$$

But this unconstrained answer would not reflect the fact that $(x_{t+\Delta})^2 = x_{t+\Delta}^2$. A potential solution for this problem is we choose appropriate values for Δ_1 and Δ_2 such that $(x_{t+\Delta})^2 = x_{t+\Delta}^2$.

$$\begin{bmatrix} x_t \\ x_t^2 \end{bmatrix} = \begin{bmatrix} \epsilon_{1,t} \\ \epsilon_{2,t} \end{bmatrix} + \mathbf{A} \begin{bmatrix} \epsilon_{1,t-1} + \Delta_1 \\ \epsilon_{2,t-1} + \Delta_2 \end{bmatrix} + \mathbf{A}^2 \begin{bmatrix} x_{t-2} \\ x_{t-2}^2 \end{bmatrix} \tag{26}$$

Now setting $(x_{t+\Delta})^2 = x_{t+\Delta}^2$ yields

$$(\epsilon_{1,t} + a_{11}(\epsilon_{1,t-1} + \Delta_1) + a_{12}(\epsilon_{2,t-1} + \Delta_2) + \phi_1)^2 = \epsilon_{2,t} + a_{21}(\epsilon_{1,t-1} + \Delta_1) + a_{22}(\epsilon_{2,t-1} + \Delta_2) + \phi_2 \tag{27}$$

where ϕ_1 is the first term of the vector $\mathbf{A}^2 \begin{bmatrix} x_{t-2} \\ x_{t-2}^2 \end{bmatrix}$ and ϕ_2 is the second term. Equation 27 can be solved for Δ_2 as a function of Δ_1. However inspecting Equation 27 we note that the choice of the residual shock, Δ_2, chosen for period $t-1$ is a function of the residuals at period t. Typically, the concept of the impulse response function calculated at time $t-k$ requires knowledge of the information available at time $t-k$, this framework fails to satisfy this condition.

B Two uninformed market participants with differing trade size

To formally test whether both participants have equivalent permanent price impacts we conduct the following regression:

$$i_k^{ls} = \beta i_k^{ss} + \epsilon \tag{28}$$

where i_k^{ls} is the permanent price impact for a trade with a size that falls into quantile k initiated by market participant ls and i_k^{ss} is the permanent price impact for a trade with a size that falls into quantile k initiated by market participant ss. In Equation 28, if $\beta > 1$ then the large trade size market participant, ls, has a larger permanent price impact than the small trade size market participant, ss, for an equivalent sized trade and we would conclude that ls has more private information. Conversely, if $\beta < 1$ we conclude that ss has more private information than ls. If $\beta = 1$ then neither participant has more private information. We find that β is insignificantly different from one for all sample stocks. The estimated coefficients range from 0.998 to 1.011 with a maximum t-statistic of 0.34 when testing if the coefficient is statistically different from 1. This result confirms the *two participant nonlinear VAR* we propose draws correct conclusions when comparing the information content of a trade for multiple market participants.

C Multi factor VAR system

Here we investigate if the conflicting results occur for all sample stocks. We estimate the *two participant trade sign VAR* system and the VAR model defined by Equation 9 for all sample stocks.

Columns 1 and 2 of Table A1 report the long term impulse responses for each market participant for the *two participant trade sign VAR*. This VAR specification does not control for the order book depth imbalance and we observe, for all sample stocks, the market participant that trades with the order book imbalance has a larger estimated long term impulse response function. In contrast, columns 4 and 5 of Table A1 report the long term impulse responses estimated via Equation 9. Here, for 18 of the 20 sample stocks we conclude the trader that trades with the order book imbalance now has a smaller estimated long term impulse response function. Column 6 reports the coefficient estimates for DI affect on log returns, ω_t. For all sample stocks the coefficient is positive and significant which demonstrates the importance the order book imbalance has on quote revisions.

	VAR model no exogenous variable			VAR model with exogenous variable			
	wi impact	ai impact	$wi < ai$	wi impact	ai impact	ω_t	$wi > ai$
AMC	6.94	5.75	F	3.12	6.79	2.63	T
AMP	2.41	2.24	F	1.77	2.40	2.26	T
ANZ	4.90	3.77	F	2.75	5.13	0.87	T
BHP	5.19	3.48	F	2.76	4.48	1.13	T
BXB	6.53	4.78	F	3.44	5.42	2.50	T
CBA	10.51	6.76	F	6.14	7.81	0.58	T
CSL	6.09	0.84	F	3.52	0.91	1.25	F
IAG	1.98	1.79	F	1.44	1.97	2.11	T
MQG	23.67	18.67	F	13.09	20.95	0.93	T
NAB	5.10	3.19	F	2.67	4.44	0.77	T
NCM	10.84	3.20	F	4.64	4.24	2.44	F
ORG	10.65	9.36	F	5.80	12.11	3.16	T
QBE	6.00	5.58	F	3.81	6.63	2.23	T
RIO	15.55	10.97	F	6.94	13.13	1.00	T
STO	6.01	4.68	F	4.60	6.05	1.88	T
SUN	4.93	4.36	F	2.85	5.44	1.93	T
TLS	0.73	0.60	F	0.64	0.67	1.32	T
WBC	5.42	4.37	F	2.24	6.12	1.04	T
WOW	7.46	5.39	F	3.50	6.93	1.15	T
WPL	14.71	10.68	F	8.72	12.66	1.49	T

Table A1: Columns 1 and 2 report the estimated permanent price impact for the market particpant that trades with the order book imbalance and against the order book imbalance respectively. The permanent price impacts are estimated using the VAR model defined by Equation 6. Column 3 reports true or false if Column 1 is statistically greater than Column 2. Columns 4, 5 and 7 report the same measures as Columns 1,2 and 3 respectively, but instead use the VAR model defined by Equation 9 for estimation. This VAR model includes the depth imbalance is an explanatory variable for estimation. The coefficient for the depth imbalance, ω_t, is reported in Column 6.

D One informed and one uninformed market participant

In this appendix we investigate the *two participant nonlinear VAR's* performance when one market participant is more informed about the future price than the other. To form the two market participants we allocate all seller (buyer) initiated trades which occur at the top (bottom) quantile of midpoint prices to the informed market participant (i), and allocate all remaining trades to the uninformed market participant (u). This trade allocation method ensures the informed participant has knowledge of future price movements and buys (sells) prior to a price increase (decrease).

For our representative stock, we estimate the permanent price impact of the two participants using the *two participant nonlinear VAR*. Figure A1 displays the price impact of different trade quantiles for the two market participants and shows that equivalent buy orders for the informed market participant cause a larger price impact than buy orders for the uninformed market participant. Similarly for sells, the informed market participant has a larger negative permanent price impact for than the uninformed market participant. This figure demonstrates the *two participant nonlinear VAR* correctly captures the difference in private information between the two market participants.

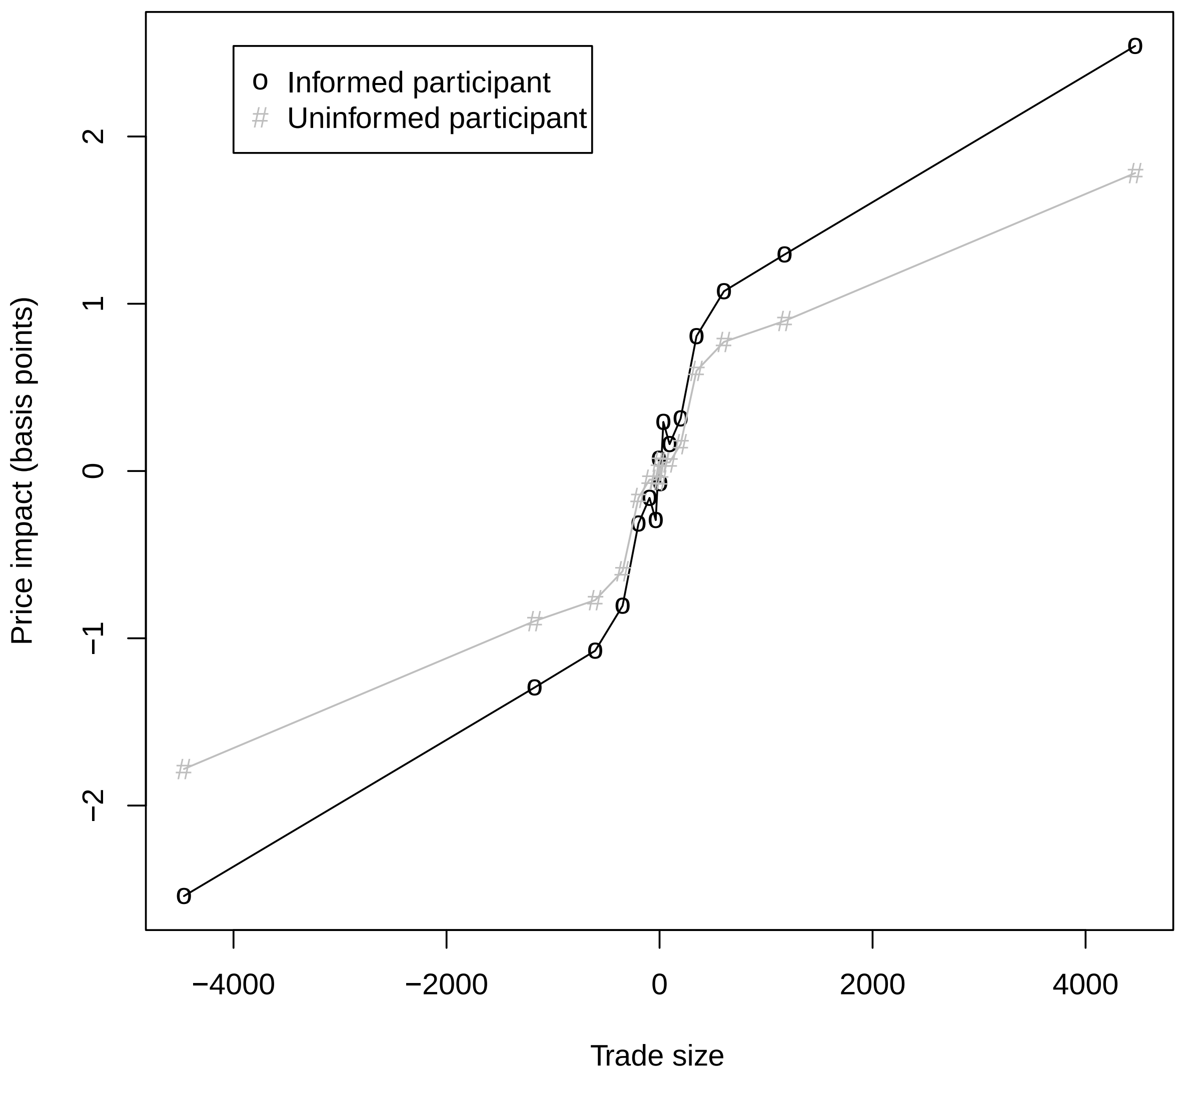

Figure A1: Plots the permanent price impact for trades of various size estimated using the VAR model defined by Equation 6. The impact is estimated for two market participants. One participant is informed and has private information about the future price of the stock, while the other is uninformed and has no information.

Next, we formally test if the *two participant nonlinear VAR* correctly concludes the informed market participant has a larger permanent price impact than the uninformed for all sample stocks. We estimate the *two participant nonlinear VAR* using 5 lags and 8 trade size quantiles for all sample stocks and obtain the corresponding impulse response function for each trade size quantile for both market participants. Using these values, we estimate Equation 28 and obtain β estimates that are statistically greater than 1 for all sample stocks. The lowest (highest) value is 1.163 (1.79) for TLS (MQG), with a sample stock average of 1.49. These results demonstrate that the *two participant nonlinear VAR* correctly concludes the informed market participants trades cause a larger permanent price impact than equivalent size trades executed by the uninformed market participant.