Optimal Trade Execution of Equities in a Limit Order Market

Richard Coggins, Adam Blazejewski,
School of Electrical and Information Engineering,
The University of Sydney, NSW, 2006.
Australia

Michael Aitken
Chair of Capital Market Technologies,
University of New South Wales, NSW, 2052.
Australia

Emails: richardc@ee.usyd.edu.au, adamb@ee.usyd.edu.au, mai@smarts.com.au

Abstract: *This paper describes an approach for optimising trade execution in a limit order market. The way a trade is executed becomes important when the trade is a significant proportion of the days turnover in a particular security. Under these circumstances limited liquidity leads to a significant transaction cost referred to as trade shortfall. We describe a method for calculating a trade execution plan which balances intra-day variations in the supply of liquidity against the risk of adverse future price movements. Our trade execution plans correspond to solutions of discrete time dynamic programming problems. This formulation admits the specification of transaction costs within a value at risk framework. The trade execution plans are derived and tested for three popular stocks on the Australian Stock Exchange (ASX). The performance of the plans is evaluated on an out of sample test set of the limit order book for each security and compared to three simpler trade execution strategies.*

Keywords: *Algorithms, trade execution, dynamic programming, transaction costs*

1 INTRODUCTION

When a large equity fund manager makes a decision to purchase or liquidate a holding in a particular security the realisation of the transaction or its execution is not straight forward, as is the case for the small investor. The quality of execution is measured by transaction costs. There is no formal definition of the best execution [9], but in practice it means execution at the most favourable prices. Treynor suggested that a performance difference between a real portfolio and a paper portfolio be used as a measure of transaction costs [12]. A paper portfolio is a theoretical portfolio which could be traded without transaction costs. Perold called the above difference an implementation shortfall [11]. The shortfall can be divided into two components. The explicit component represents broker fees and taxes. The implicit component captures market impact cost, opportunity cost, and bid-ask spread. To correctly calculate transaction costs, however, one also needs to include other factors like investment manager type, investment style, risk aversion, trade difficulty, choice of a pre-trade benchmark, broker reputation, market sensitive announcements, seasonality, and others.

Best execution is difficult to achieve in practice due to limited liquidity and price volatility in the market over the time period that the transaction is to be completed. Limited liquidity implies that as a large trade is attempted, the trader becomes a large demander of liquidity and as a result will incur a premium to complete the trade. This premium represents the cost of immediacy, a transaction cost which we will refer to as the trade execution shortfall (the difference in the price of the security when the trade commences and the average price achieved for the completed trade for a sell and visa versa for a buy) [4]. The goal of a trade execution strategy is to minimise the trade execution shortfall subject to time and specified risk level constraints. Intuitively, one observes that by distributing trading volume across time and choosing peaks in liquidity, to the extent that they are known, trading shortfall can be reduced at the risk that the price of the security (as measured by the midpoint of the bid-ask spread) moves adversely in relation to the type of transaction to be completed. Hence, we are seeking to balance different sources of risk, namely a risk characterising our uncertain knowledge of future liquidity versus a risk corresponding to our uncertain knowldege of future prices. The realisation of a particular strategy is then in turn, determined by our propensity to take risks to minimise the trade execution shortfall.

In this paper, we consider the problem of trade execution in the context of a fully electronic limit order market, namely, the Australian Stock Exchange (ASX). A survey of equity trading transaction costs in the US context is provided in [7]. The method we propose is an extension of a framework derived by Almgren and Chriss [1, 2] which in turn builds on work by [3]. The orginal framework and our extensions of it are described in section 2. Huberman and Stanzl [6] minimise the mean and variance of the cost of buying a block of shares. Their objective function is similar to the one in [1] but it refers to the total cost of trading, not only transaction costs. A solution for a non-stationary market impact function is provided as a recursive formula. This model does not include liquidity risk, however, when market impact itself is stochastic. Hisata and Yamai [5] adopted the transaction costs model developed by Almgren and Chriss [2], but changed the objective function to match the Value at Risk (VaR) framework. They point out that the traditional VaR framework does not include market impact cost and liquidity constraints. By incorporating these factors they formulated a Liquidity-adjusted Value at Risk framework (L-VaR). They obtain closed-form solutions for optimal execution time with linear and non-linear market impact costs. The main limitation of their model, however, is the assumption that the trading speed is constant. Numerical examples with L-VaR and traditional VaR calculations were presented, suggesting that the traditional approach overestimates the risk for very liquid stocks, while understimating it for illiquid stocks. A stochastic market impact model was also considered, with solutions found through numerical methods.

Obtaining high performance trade execution plans also depends on our ability to predict any patterns in the security's behaviour. Section 3 describes the methods we use to characterise a security's liquidity and price behaviour. In section 4 we will compare our results with three other trade execution benchmarks which we denote as one-interval, uniform and VWAP execution. One interval execution is intended as a risk averse strategy, where we take the currently available liquidity and avoid the risk of future volatility and liquidity risk. A uniform trade execution strategy recognises a need to reduce the peak demand for liquidity, but is agnostic with respect to the intra-day levels and patterns in the security's liquidity and price. A VWAP strategy seeks to minimise the difference in volume weighted average price of the transaction and the volume weighted average price of the entire market for the security and typically uses a historical trading volume as a proxy for liquidity [8] and may incorporate technical rules to take into account any knowledge of price behaviour. Hence, a VWAP strategy makes use of intra-day liquidity and price patterns, however, it differs from our optimal approach in the sense that it has an implicit approach to the transaction risk and may not necessarily involve the optimisation of a trade execution objective. We compare the four trade execution techniques by measuring the trade shortfalls on an out of sample interval of the limit order book for the respective securities, National Australia Bank (NAB), BHP-Billiton (BHP) and Telstra (TLS). The next section describes our general approach to trade execution shortfall minimisation.

2 MINIMISING EXECUTION SHORTFALL

The formulation of the optimal trade execution problem described in this section builds on the work of Almgren and Chriss [1, 2]. The formulation can be summarised as follows. Firstly, we will define a price process for the security. Here we adopt a discrete time arithmetic random walk parameterised by the volatility over a time interval. This is a reasonable approximation to the more plausible geometric random walk over the intra-day time scales we are considering. Next we will introduce a temporary market impact function which measures the average price concession when we sell a given quantity of shares (from this point onwards we will consider the liquidation of a security for the purpose of explanation, however the buy case follows directly). Associated with the market impact function is a measure of uncertainty of the market impact, which we will call impact risk. Given normally distributed random shocks for price and market impact we can determine the expected execution shortfall and the variance of the shortfall. A utility function is then constructed from the mean and variance which is a risk constrained shortfall to be optimised. For certain forms of the volatilty and market impact functions trading trajectories can be derived analytically [1, 2]. Here, we extend the formulation to include non-stationary market impact, impact risk and volatility which require numerical solution, for which we employ deterministic discrete time dynamic programming [10]. Finally, we will show the connection between the mean/variance utility function for shortfall and the value at risk objective.

Let X denote the quantity of shares which we wish to sell in a time T. We wish to derive trade execution plans which specify the number of shares to sell at N discrete time intervals. Let $\tau = T/N$ denote the length of the time intervals. Let k be an index over the N time intervals. A trading trajectory is then defined by a list $x_0, ..., x_N$ where x_k is the number of shares we still hold after time interval k, $x_0 = X$ and we require $X_N = 0$. Equivalently we can specify a trade by a list of shares to be traded at each time interval specified by $n_k = x_k - x_{k-1}$. The price process we will assume for the security is,

$$S_k = S_{k-1} + \sigma \tau^{1/2} \xi_k, \tag{1}$$

where S_k is the security price at time interval k, σ is the volatility and ξ_k is a set of IID zero mean unit variance normal random variables. In contrast to [1], for the purposes of our empirical work presented in this paper, we do not attempt to characterise long term market impact. However, we will consider temporary market impact. By temporary market impact, we mean that as we trade we may exhaust the liquidity available at a series of price levels leading to a lower price as we sell during a given time interval k. We assume that the effect of our liquidity consumption is limited to each individual trading interval, with a new equilibrium price being established at the start of the next interval as if we hadn't traded. The price that we achieve as a consequence of this temporary market impact is given by,

$$\tilde{S_k} = S_{k-1} - h(\frac{n_k}{\tau}) + \tau^{-1/2} f(\frac{n_k}{\tau}) \tilde{\xi_k}, \tag{2}$$

where $\tilde{\xi_k}$ are zero mean, unit variance normal random variables independent of the ξ_k. In section 3 we will indicate how $h()$ and $f()$ can be determined empirically. Hence, the total trading revenue on completion of all trades is $\sum_{k=1}^{N} \tilde{S}_k n_k$. The trade execution shortfall is therefore given by,

$$XS_0 - \sum_{k=1}^{N} n_k \tilde{S}_k \tag{3}$$

From this we can determine the expectation and variance of the shortfall with respect to the random innovations of the price process ξ,

$$E(x) = \tau \sum_{k=1}^{N} v_k h(v_k) \tag{4}$$

$$V(x) = \sigma^2 \tau \sum_{k=1}^{N} x_k^2 + \tau \sum_{k=1}^{N} v_k^2 f(v_k)^2 \tag{5}$$

where we have used $v_k = n_k/\tau$. We now construct a utility function,

$$U(x) = E(x) + \lambda V(x) \tag{6}$$

For $h(v)$ of the form $h(v) = \eta v^i$ where i is an integer, $U(x)$ can be minimised analytically [1]. In this paper we consider two practical extensions of this model. To account for the bid-ask spread and the fact that market impact is fixed below a critcal volume h_{v0} (see also Figure 1) we consider $h(v)$ of the form,

$$h(v) = \begin{cases} \eta(v - h_{v0}) + \epsilon & v > h_{v0} \\ \epsilon & v \leq h_{v0} \end{cases} \tag{7}$$

Further, we condider the case where the market impact function $h(v)$ has a volume dependent uncertainty associated with it. Following Almgren's approach and extending it to take take into account the bid-ask spread and critical volume we define the impact risk by a function,

$$f(v) = \begin{cases} \beta(v - h_{v0}) + \alpha & v > f_{v0} \\ \alpha & v \leq f_{v0} \end{cases} \tag{8}$$

This non-linearity in $h(v)$ requires us to minimise $U(x)$ numerically. Secondly we consider the case where $\eta(k)$, $\sigma(k)$ are functions of time as measured by index k. When $h(v)$ has a linear form $h(v) = \eta(k)v$ and $f(v)$ is constant with respect to volume traded $f(v) = \alpha$, we can obtain an analytical solution to the trading trajectory as follows. By writing $v_k = x_{k-1} - x_k$ and differentiating $U(x)$ with respect to each x_k we obtain,

$$\frac{\partial U(x_k)}{\partial x_k} = 2\{\lambda \sigma_k^2 x_k - (\eta_k + \lambda \alpha_k^2)(x_{k-1} - 2x_k + x_{k+1})\} \tag{9}$$

where we have shown the subscript k to indicate the parameters as a function of time. The set of equations in x_k are second order difference equations. To solve the equations we impose the boundary conditions, $x_0 = X$ and $x_N = 0$, implying an initial holding of X shares and complete liquidation of the holding by time step N. Given these constraints the ratio $r(k)$ defined by $x_k = r(k)x_{k-1}$ can be derived,

$$r(k) = \frac{A_{N-k}}{A_{N-k} + A_{N-k+1} + \lambda \sigma_{N-k}^2 - A_{N-k+1} r(k-1)} \tag{10}$$

where A_k is given by,

$$A(k) = \eta_k + \lambda \alpha_k^2 \tag{11}$$

Although (10) does not take into account non-linearities in the impact function, it does none the less take into account the non-stationarities in the parameters, and hence is useful for rapidly calculating approximate strategies for a large number of securities.

In order to compute more accurate optimal trading trajectories we employ the following reward function and two dimensional state space model (the model is defined in the context of a sell program, a buy formulation follows straightforwardly). Denote the state variable for the problem to be $S = (s_1, s_2)$ where s_1 represents the stock being held and s_2 measures time intervals (half hour bins). Let a_k denote the action variable, specifying the number of shares to be sold at each time step with the constraint $0 \leq a_k \leq s_1(k)$ implying that we never buy shares and we can never sell more shares than we currently hold. Let $R(S, a)$ denote the reward function which is a function of the current state S and the action taken by the agent a_k ($a_k = n_k$). Then,

$$R = E(S, a) + \lambda V(S, a) \tag{12}$$

where $E()$ and $V()$ have the general form of (4) and (5) respectively and,

$$\begin{aligned} s_1(k) &= s_1(k-1) - a(k) \\ s_2(k) &= s_2(k-1) + 1 \end{aligned} \tag{13}$$

For a sell program, the states are initialised as $s_1(1) = X$ and $s_2(1) = 1$. Note, the time interval $s_2(k)$ is modelled explicitly as part of the state space due to the reward function R depending on the non-stationary impact coefficient $\eta(k)$ and the non-stationary volatility $\sigma(k)$ and the other parameters used to model market impact. Numerical solutions were computed using a Matlab discrete dynamic programming toolbox [10] incorporating the non-linear impact function (7).

3 MEASURING SECURITY CHARACTERISTICS

In our experiments we consider a three week trading period for each security in order to parameterise our models. This provides us with fifteen days of trading data which we treat identically (we have chosen to ignore possible weekly seasonal effects). We divide each trading day into half hour intervals, giving us 12 intervals, since the ASX opens at 10am and closes at 4pm. For the data we consider we ignore off-market trades, crossings and undisclosed orders. For each security we then collect the following data:

- total daily trading volume for each of the 15 days.

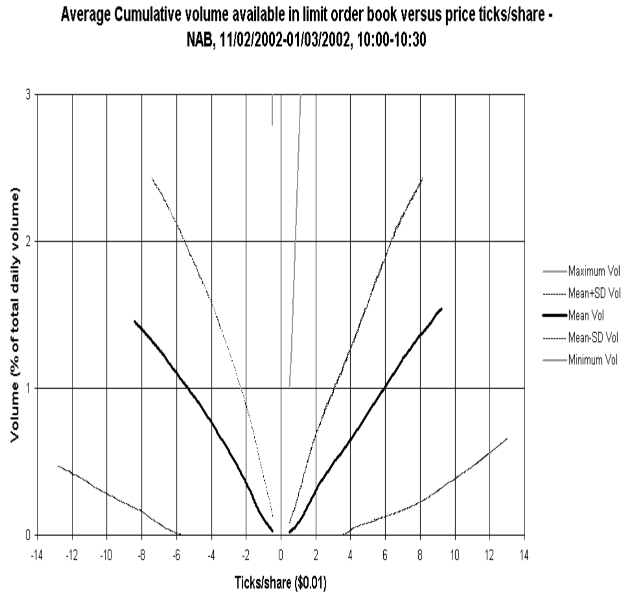

Figure 1: Transformed normalised cumulative order volume for NAB stock averaged over a 3 week period. This plot is derived from the order book by taking the time weighted average of the order volume, normalising it by the daily traded volume (y-axis), and then calculating the volume weighted price impact per share (x-axis) for respective volumes at each price tick.

- the difference in midpoint quote at the start and end of each half hour interval.

- the average, normalised, time weighted cumulative order volume available at each price tick in the order book relative to the midpoint price, for each half hour interval. Normalisation is with respect to the daily traded volume.

Figure 1 shows an example of the transformed normalised cumulative order volume as a function of price ticks from the midquote for a particular half hour bin. This provides a volume normalised, price invariant measure of market impact as a function of traded volume. Figure 2 shows average intraday impact coefficient $\eta(k)$, volatility $\sigma(k)$ and impact risk $\alpha(k)$. Note, η is derived from the inverse of the slope from the corresponding data exemplified in Figure 1.

Given these measurements the expected shortfall and variance of the shortfall may be written:

$$E(S, a) = \tau \sum_{k=1}^{N} a_k h(a_k, \eta_k, \epsilon_k, h_{kv0}) \tag{14}$$

$$V(S, a) = \tau \sum_{k=1}^{N} \sigma_k^2 (s_1(k) - a_k)^2 + \tau \sum_{k=1}^{N} a_k^2 f^2(a_k, \beta_k, \alpha_k, f_{kv0}) \tag{15}$$

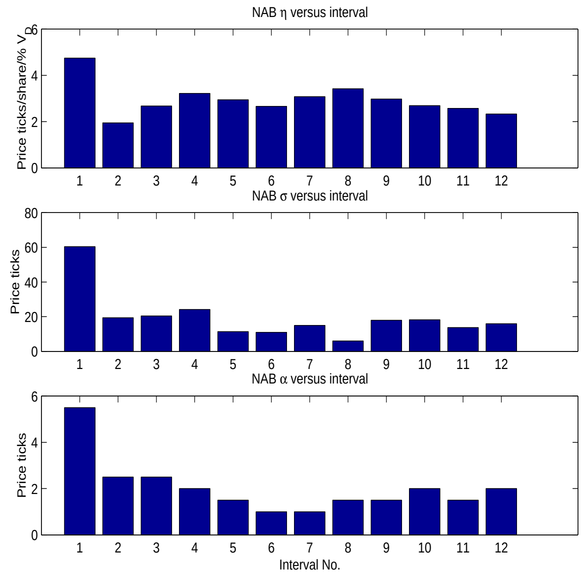

Figure 2: Intraday impact coefficient η, volatility σ and impact risk α for NAB stock averaged over a 3 week period.

4 SIMULATED TRADING PERFORMANCE

By parameterising the reward function R as described in the previous section we are then able to compute optimal trading trajectories for a given stock. Three typical trading trajectories are shown in Figure 3.

Given a trading trajectory x it then remains to test its performance. We do this by considering 63 days of trading immediately after the three weeks training period which was used to parameterise our models. Each of the trading trajectories we calculate sets half hour trading targets during the day. We then simulate trading over the half hour by trading a fixed fraction of the trading target against the orders available in the schedule every five minutes. This entails the following difficulties. It is possible that we match against the same orders a number of times. It also assumes that orders we consume, don't ultimately influence future order flows, i.e., we only take account of instantaneous market impact. To assess the performance of the optimal approach that we have presented in this paper we also evaluated a number of simpler strategies:

- One interval (ONEINT): we liquidate all the stock during the first half hour of trading.

- Uniform (UNIFORM): we liquidate our holding uniformly throughout the day.

- (VWAP) trader: we use historical fractional trading volumes averaged over the 3 week training period as targets for the half hour intervals.

- Optimal (Lambda1, Lambda2): the algorithm presented in this paper, constrained to trade all the holding in one day, at two different risk settings λ (Lambda1=10^{-6}, Lambda2=10^{-1}).

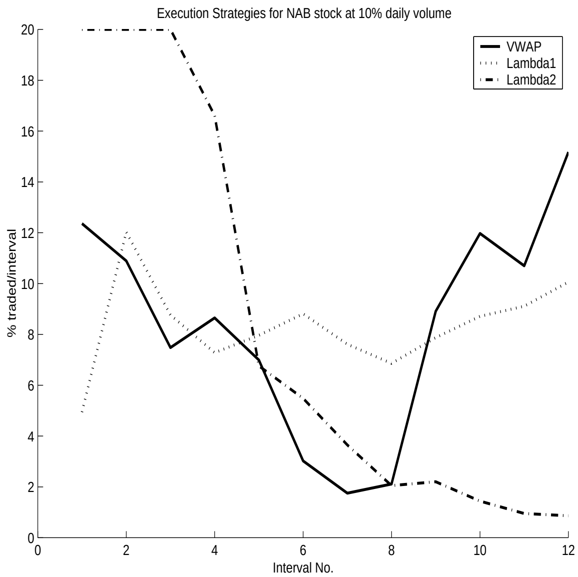

Figure 3: Two trading trajectories computed for the NAB stock with the risk parameter $\lambda = 10^{-6}$ and $\lambda = 0.1$ respectively. $\lambda = 10^{-6}$ corresponds to the case where the risk of adverse price movements and our uncertainty about future liquidity later in the day are negligible and is therefore a risk neutral strategy. $\lambda = 0.1$ corresponds to a shortfall value at risk $Var(p) = 66$c/share, with $p = 0.99$ implying that this shortfall is predicted to be exceeded 1% of the time. Note that for this second trajectory, the first two intervals trade 20%, which is due to a rule we introduced that limits trading in a half hour period to a maximum of 2% of forecast daily volume. We did this to ensure our impact was temporary. The VWAP trajectory is based on the average fraction of trading volume, traded in the stock over the three week training period at that time of day.

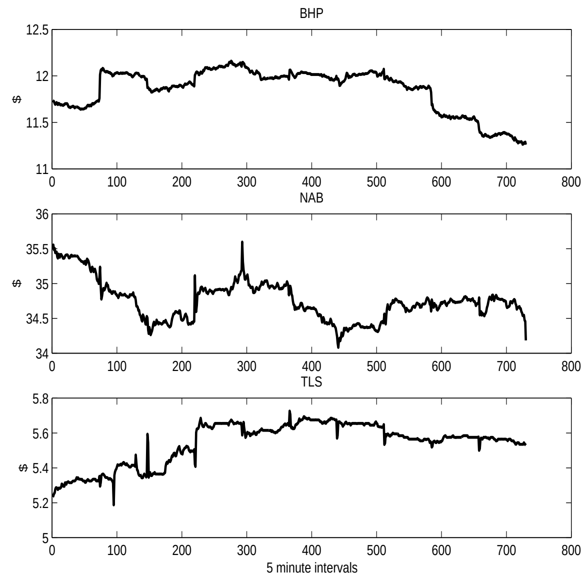

Figure 4: Intraday midquote prices at 5 minute intervals for 10 trading days of the out of sample test period. Note the sudden price jumps, some of which are intraday, others overnight. In this paper we do not estimate overnight volatility.

In order to quantify trading performance we consider two metrics. The first is execution shortfall as given in (3). The second is a volume weighted average price metric,

$$\Delta VWAP = \{\frac{VWAP_{trade}}{VWAP_{market}} - 1\} * 10^4 \qquad (16)$$

where $VWAP_{trade} = \sum_k \tilde{S}_k n_k$ and the units are basis points ($\frac{1}{100}$th of a percent). Execution shortfall directly measures the transaction costs resulting from trading. $\Delta VWAP$ on the other hand measures the trading performance against the rest of the market. For the execution shortfall measure we are also able to calculate a value at risk for any given trade execution strategy. We do this as follows. Firstly, we determine,

$$\lambda_v = -\frac{\partial E}{\partial V^{\frac{1}{2}}} \qquad (17)$$

for any trading strategy x. We then calculate,

$$p = \int_{\infty}^{\lambda_v} N(z)dz \qquad (18)$$

where $N(z)$ is the standard normal density function and the value at risk is given by $Var(p) = E(x) + \lambda_v V(x)^{\frac{1}{2}}$.

The results are shown in tables 1, 2 and 3. We note that the variance of percentage volume traded is not taken into account in forecasting the variance of the shortfall. We will take this into account in later work. We test our strategies for both the buy and sell cases so as to reveal any bias in our results due to the specific out of sample test period used. We note that the one interval strategies for each stock correspond to high risk strategies as indicated by low or zero values of p, with correspondingly low or negative transaction costs. (We note that in a small number of instances our value at risk calculation appears spurious, which we believe is due to the non-differentiable nature of our market impact function shown in (7). These are indicated in the tables by a '-'). We see that for each of the three stocks the UNIFORM, VWAP and Lambda1 strategies produce similar shortfall forecasts and measured shortfalls and the value at risk calculation shows that they are often close to risk neutral strategies ($p = 0.5$). The Lambda2 strategies correspond to more risk averse strategies as indicated by the value at risk forecasts and standard deviations of shortfall forecasts. We note that the measured transaction costs do show reduced standard deviations, but not to the extent that our simple model predicts. Comparing the three stocks, we observe that TLS is very liquid, in the sense that at all trading levels, the forecast shortfalls are in most cases 0.5 ticks, indicating that there is large order volume in the first price tick of the order book. We subsequently observe that the resulting measured shortfalls are independent of the volume traded.

We note that volatility effects are dominant in our results. This suggests, that in order to improve our trading strategies, we would need an improved price process model which exploits intraday serial correlation in prices. Figure 4 gives an indication of the level of volatility of the three stocks over 10 trading days of the test period.

Table 1: Performance of Trade Execution Strategies for BHP stock. Results are provided for the two naive strategies ONEINT and UNIFORM, the VWAP heuristic and the optimal approach with $\lambda = 10^{-6}$ and $\lambda = 10^{-1}$. Each of the strategies are tested at nominally 1%, 5% and 10% of average daily traded volume ($\hat{V_D}$). V_D is actual average daily traded volume for the test period. VaR(p) is the shortfall upper bound that can be achieved with probability p. $\hat{sf}$ and sf are the forecast shortfall and measured shortfall in cents per share traded respectively. $\Delta VWAP$ indicates the fractional difference in volume weighted price achieved for the whole trade against the entire market for the day expressed in percentage basis points. Standard deviations over the 63 trading day test period appear in brackets.

Trade Type (Units)	$\hat{V_D}$ (%)	V_D (%)	p	VaR(p) (c/share)	$\hat{sf}$ (c/share)	sf (c/share)	$\Delta VWAP$ (basis pts.)
ONEINT							
Buy	1	1.3	0.00	-5.8	2.0 (1.7)	-4.6 (19.1)	11.3 (50.3)
Sell	1	1.3	0.23	-5.6	1.4 (9.5)	6.3 (19.1)	-3.5 (50.3)
Buy	5	6.7	0.00	-6.1	9.5 (3.1)	-3.1 (19.2)	23.9 (50.2)
Sell	5	6.7	0.22	-5.7	7.3 (17.2)	7.6 (19.0)	-15.4 (50.7)
Buy	10	13.5	0.00	-6.2	18.8 (4.8)	-1.0 (19.6)	42.7 (53.4)
Sell	10	13.5	0.22	-5.7	14.7 (26.7)	9.8 (18.7)	-34.5 (53.1)
UNIFORM							
Buy	1	1.3	0.60	2.9	0.5 (9.5)	-5.2 (20.4)	5.6 (9.2)
Sell	1	1.3	0.60	2.9	0.5 (9.5)	6.3 (20.4)	-3.9 (9.3)
Buy	5	6.7	0.69	5.4	0.6 (9.5)	-5.1 (20.4)	6.2 (9.2)
Sell	5	6.7	0.64	3.9	0.5 (9.5)	6.4 (20.4)	-4.5 (9.3)
Buy	10	13.5	0.86	11.3	1.1 (9.5)	-5.1 (20.4)	6.8 (9.2)
Sell	10	13.5	0.83	10.1	0.9 (9.5)	6.4 (20.4)	-5.1 (9.3)
VWAP							
Buy	1	1.3	0.59	2.4	0.5 (8.6)	-5.2 (20.3)	5.8 (6.2)
Sell	1	1.3	0.59	2.5	0.5 (8.7)	6.3 (20.3)	-3.8 (6.2)
Buy	5	6.7	0.69	5.1	0.8 (8.6)	-5.1 (20.3)	6.5 (6.2)
Sell	5	6.7	0.66	4.3	0.7 (8.7)	6.4 (20.3)	-4.5 (6.3)
Buy	10	13.5	0.83	9.8	1.5 (8.6)	-5.0 (20.3)	7.4 (6.4)
Sell	10	13.5	0.83	9.7	1.3 (8.8)	6.5 (20.3)	-5.5 (6.2)
Lambda1							
Buy	1	1.3	0.60	2.9	0.5 (9.9)	-4.8 (19.7)	9.6 (27.7)
Sell	1	1.3	0.60	3.0	0.5 (9.8)	6.0 (19.4)	-1.0 (36.1)
Buy	5	6.7	0.69	5.5	0.6 (9.9)	-5.2 (20.6)	5.8 (9.7)
Sell	5	6.7	0.64	4.2	0.5 (9.9)	6.4 (20.5)	-4.6 (8.7)
Buy	10	13.5	0.84	11.0	1.1 (9.9)	-5.1 (20.6)	6.5 (9.7)
Sell	10	13.5	0.82	10.0	0.9 (9.8)	6.4 (20.5)	-5.1 (9.8)
Lambda2							
Buy	1	1.3	0.86	3.8	1.4 (2.2)	-4.7 (19.1)	10.6 (46.3)
Sell	1	1.3	-	-	0.6 (6.2)	6.0 (19.3)	-1.3 (38.4)
Buy	5	6.7	0.86	8.3	2.1 (5.7)	-4.6 (19.5)	10.8 (28.8)
Sell	5	6.7	1.00	33.6	1.6 (7.3)	6.2 (19.5)	-3.3 (27.5)
Buy	10	13.5	0.93	13.4	1.9 (7.8)	-4.8 (19.9)	9.3 (15.9)
Sell	10	13.5	0.96	15.7	1.6 (8.0)	6.4 (19.9)	-4.5 (17.6)

5 CONCLUSIONS

This study has shown that a relatively simple model forecasts the variance of trading shortfalls of the static strategies tested in this paper well, and shows that without more information on future prices, that standard deviations of trading shortfalls are many multiples of the expected shortfalls up to moderate levels of trading (10% of daily traded volume). We also see that price trends in the test period are having a strong influence on results. We plan to investigate the modelling of these trends. The same model may also be used to compute optimal strategies taking into account non-stationary volatilility and liquidity effects and non-linearities in the market impact function. We note that high price volatility makes transaction costs difficult to forecast without a directional view of prices. The models presented in this paper are static models. In the future, we will investigate state space models which also include state variables for price and liquidity. Investigation of how often the models should be recomputed is also important. We envisage extending our models to include improved daily trading volume forecasts, improved price process models encompassing non-gaussian innovations, price jumps and serial correlation, medium term and permanent trading impacts and effects of undisclosed orders.

ACKNOWLEDGMENTS

We acknowledge the provision of the order book data and the SMARTS software for preprocessing the order book data via the Australian Capital Markets Co-operative Research Centre (CMCRC) and its industry partners. Mr. Adam Blazejewski is supported by an Australian postgraduate award and a CMCRC scholarship.

REFERENCES

[1] Robert F. Almgren. Optimal execution with nonlinear impact functions and trading-enhanced risk, Working paper, October 2001.

Table 2: Performance of Trade Execution Strategies for NAB stock. Results are provided for the two naive strategies ONEINT and UNIFORM, the VWAP heuristic and the optimal approach with $\lambda = 10^{-6}$ and $\lambda = 10^{-1}$. Each of the strategies are tested at nominally 1%, 5% and 10% of average daily traded volume ($\hat{V_D}$). V_D is actual average daily traded volume for the test period. VaR(p) is the shortfall upper bound that can be achieved with probability p. $\hat{sf}$ and sf are the forecast shortfall and measured shortfall in cents per share traded respectively. $\Delta VWAP$ indicates the fractional difference in volume weighted price achieved for the whole trade against the entire market for the day expressed in percentage basis points. Standard deviations over the 63 trading day test period appear in brackets.

Trade Type (Units)	$\hat{V_D}$ (%)	V_D (%)	p	VaR(p) (c/share)	$\hat{sf}$ (c/share)	sf (c/share)	$\Delta VWAP$ (basis pts.)
ONEINT							
Buy	1	1.1	0.30	-3.9	5.1 (17.4)	4.7 (26.0)	1.2 (48.2)
Sell	1	1.1	0.28	-3.5	6.0 (16.2)	-0.4 (22.7)	-11.1 (55.1)
Buy	5	5.4	0.34	-3.6	24.1 (66.3)	10.8 (26.4)	18.7 (50.9)
Sell	5	5.4	0.32	-3.3	28.2 (68.3)	7.4 (27.4)	-33.6 (59.4)
Buy	10	10.9	0.34	-3.6	47.9 (127.5)	21.0 (27.8)	48.2 (57.1)
Sell	10	10.9	0.33	-3.2	56.1 (133.5)	24.9 (38.3)	-83.9 (84.1)
UNIFORM							
Buy	1	1.1	0.54	4.0	0.5 (31.5)	5.5 (32.2)	3.3 (9.1)
Sell	1	1.1	0.55	4.5	0.6 (31.5)	-3.5 (32.0)	-2.3 (9.0)
Buy	5	5.4	0.67	15.3	1.5 (31.5)	5.7 (32.2)	3.9 (9.2)
See Table 1 Sell	5	5.4	0.68	16.2	1.6 (31.5)	-3.3 (32.1)	-3.0 (9.2)
Buy	10	10.9	0.80	29.3	2.7 (31.5)	5.9 (32.2)	4.6 (9.2)
Sell	10	10.9	0.81	30.9	2.9 (31.5)	-3.1 (32.0)	-3.6 (9.1)
VWAP							
Buy	1	1.1	0.54	3.8	0.6 (30.2)	5.2 (32.0)	2.5 (6.5)
Sell	1	1.1	0.55	4.3	0.6 (30.2)	-3.2 (31.6)	-3.1 (6.9)
Buy	5	5.4	0.65	13.7	1.8 (30.2)	5.5 (32.0)	3.3 (6.6)
Sell	5	5.4	0.67	15.0	1.9 (30.2)	-3.0 (31.6)	-3.9 (7.0)
Buy	10	10.9	0.78	26.8	3.2 (30.3)	5.8 (32.0)	4.1 (6.7)
Sell	10	10.9	0.80	29.1	3.6 (30.3)	-2.7 (31.7)	-4.7 (7.0)
Lambda1							
Buy	1	1.1	0.54	3.9	0.5 (32.4)	5.5 (33.0)	3.4 (9.4)
Sell	1	1.1	0.55	4.4	0.6 (33.3)	-3.6 (33.0)	-2.3 (9.2)
Buy	5	5.4	0.66	14.9	1.4 (32.4)	5.8 (32.4)	4.0 (8.9)
Sell	5	5.4	0.67	16.2	1.6 (32.4)	-3.4 (33.0)	-2.8 (9.8)
Buy	10	10.9	0.79	28.4	2.6 (32.4)	6.0 (32.4)	4.7 (8.9)
Sell	10	10.9	0.79	29.5	2.7 (33.3)	-3.2 (33.1)	-3.3 (10.0)
Lambda2							
Buy	1	1.1	-	-	4.4 (15.3)	4.9 (25.9)	1.8 (41.4)
Sell	1	1.1	-	-	5.2 (14.4)	-1.5 (23.4)	-8.0 (46.3)
Buy	5	5.4	-	-	5.8 (21.9)	6.4 (27.0)	6.0 (30.8)
Sell	5	5.4	-	-	7.9 (21.8)	-1.9 (25.9)	-6.9 (30.9)
Buy	10	10.9	0.97	55.0	5.2 (26.1)	6.6 (28.9)	6.7 (20.5)
Sell	10	10.9	0.99	66.6	6.6 (26.1)	-2.4 (28.3)	-5.6 (20.8)

[2] Robert F. Almgren and Neil Chriss. Optimal execution of portfolio transactions. *Journal of Risk*, 3:5–39, 2000.

[3] Dimitris Bertsimas and Andrew W. Lo. Optimal control of execution costs. *Journal of Financial Markets*, 1(1):1–50, 1998.

[4] Lawrence E. Harris. Optimal dynamic order submission strategies in some stylized trading problems. *Financial Markets, Institutions and Instruments*, 7(2):1–76, 1998.

[5] Yoshifumi Hisata and Yasuhiro Yamai. Research toward the practical application of liquidity risk evaluation methods. *Monetary and Economics Studies*, (December):83–127, 2000.

[6] Gur Huberman and Werner Stanzl. Optimal liquidity trading, Working Paper, September, 2001.

[7] Donald B. Keim and Ananth Madhavan. The cost of institutional equity trades. *Financial Analysts Journal*, 54(4):50–69, 1998.

[8] Hizuru Konishi. Optimal slice of a vwap trade. *Journal of Financial Markets*, 5(2):197–221, 2002.

[9] Jonathan R. Macey and Maureen O'Hara. The law and economics of best execution. *Journal of Financial Intermediation*, 6:188–223, 1997.

[10] Mario J. Miranda and Paul L. Fackler. *Applied Computational Economics and Finance*. MIT Press, 1st edition, 2002.

[11] Andre F. Perold. The implementation shortfall: paper versus reality. *Journal of Portfolio Management*, 14:4–9, 1988.

[12] Jack L. Treynor. What does it take to win the trading game? *Financial Analysts Journal*, 37(1):55–60, 1981.

Table 3: Performance of Trade Execution Strategies for TLS stock. Results are provided for the two naive strategies ONEINT and UNIFORM, the VWAP heuristic and the optimal approach with $\lambda = 10^{-6}$ and $\lambda = 10^{-1}$. Each of the strategies are tested at nominally 1%, 5% and 10% of average daily traded volume ($\hat{V_D}$). V_D is actual average daily traded volume for the test period. VaR(p) is the shortfall upper bound that can be achieved with probability p. $\hat{sf}$ and sf are the forecast shortfall and measured shortfall in cents per share traded respectively. $\Delta VWAP$ indicates the fractional difference in volume weighted price achieved for the whole trade against the entire market for the day expressed in percentage basis points.Standard deviations over the 63 trading day test period appear in brackets.

Trade Type (Units)	$\hat{V_D}$ (%)	V_D (%)	p	VaR(p) (c/share)	$\hat{sf}$ (c/share)	sf (c/share)	$\Delta VWAP$ (basis pts.)
ONEINT							
Buy	1	1.0	0.00	-4.0	0.7 (0.8)	-2.7 (12.2)	8.0 (91.2)
Sell	1	1.0	0.00	-3.9	0.7 (0.6)	3.1 (11.7)	-1.6 (72.3)
Buy	5	4.9	0.00	-3.4	3.2 (2.1)	-2.2 (12.0)	16.8 (81.9)
Sell	5	4.9	0.00	-3.2	3.1 (1.7)	3.5 (11.6)	-9.2 (70.7)
Buy	10	9.8	0.01	-3.2	6.4 (3.9)	-1.6 (11.9)	27.6 (76.5)
Sell	10	9.8	0.00	-3.1	6.1 (3.2)	4.0 (11.7)	-17.8 (73.2)
UNIFORM							
Buy	1	1.0	0.70	4.5	0.5 (7.7)	-2.5 (11.3)	10.4 (12.5)
Sell	1	1.0	0.70	4.5	0.5 (7.7)	3.4 (11.3)	-7.8 (12.4)
Buy	5	4.9	0.70	4.5	0.5 (7.7)	-2.5 (11.3)	10.6 (12.6)
Sell	5	4.9	0.70	4.5	0.5 (7.7)	3.4 (11.3)	-8.1 (12.4)
Buy	10	9.8	0.72	5.0	0.5 (7.7)	-2.5 (11.3)	10.8 (12.6)
Sell	10	9.8	0.72	5.0	0.5 (7.7)	3.4 (11.3)	-8.3 (12.3)
VWAP							
Buy	1	1.0	0.69	4.3	0.5 (7.6)	-2.5 (11.4)	9.8 (11.8)
Sell	1	1.0	0.69	4.3	0.5 (7.6)	3.4 (11.4)	-8.2 (11.4)
Buy	5	4.9	0.69	4.3	0.5 (7.6)	-2.5 (11.4)	10.1 (11.8)
Sell	5	4.9	0.69	4.3	0.5 (7.6)	3.5 (11.4)	-8.5 (11.4)
Buy	10	9.8	0.71	4.8	0.5 (7.6)	-2.5 (11.4)	10.3 (11.8)
Sell	10	9.8	0.71	4.8	0.5 (7.6)	3.5 (11.4)	-8.8 (11.4)
Lambda1							
Buy	1	1.0	0.70	4.7	0.5 (8.1)	-2.6 (11.9)	9.2 (75.1)
Sell	1	1.0	0.70	4.6	0.5 (8.1)	3.2 (11.6)	-2.9 (59.5)
Buy	5	4.9	0.70	4.7	0.5 (8.1)	-2.5 (11.6)	11.3 (39.3)
Sell	5	4.9	0.70	4.7	0.5 (8.1)	3.4 (11.5)	-7.9 (36.5)
Buy	10	9.8	0.70	4.7	0.5 (8.1)	-2.4 (11.4)	11.9 (26.6)
Sell	10	9.8	0.70	4.6	0.5 (8.1)	3.4 (11.4)	-8.0 (25.2)
Lambda2							
Buy	1	1.0	0.88	1.8	0.6 (1.0)	-2.7 (12.1)	8.3 (87.2)
Sell	1	1.0	0.85	1.6	0.7 (0.9)	3.2 (11.7)	-1.7 (70.3)
Buy	5	4.9	0.80	5.0	0.8 (5.0)	-2.5 (11.7)	11.2 (53.6)
Sell	5	4.9	0.79	4.9	0.8 (5.0)	3.4 (11.6)	-6.5 (46.9)
Buy	10	9.8	0.80	6.1	0.7 (6.6)	-2.4 (11.5)	11.6 (34.4)
Sell	10	9.8	0.79	6.1	0.7 (6.6)	3.4 (11.4)	-7.6 (31.5)