Journal of Banking & Finance 35 (2011) 2319–2329

Contents lists available at ScienceDirect

Journal of Banking & Finance

journal homepage: www.elsevier.com/locate/jbf

Optimal VWAP trading under noisy conditions

Mark L. Humphery-Jenner *

School of Banking and Finance, Australian School of Business, University of New South Wales, Sydney NSW 2052, Australia
Tilburg University, 5000LE Tilburg, The Netherlands

ARTICLE INFO

Article history:
Received 2 July 2009
Accepted 24 January 2011
Available online 1 February 2011

JEL classification:
C53
G12
G29

Keywords:
VWAP strategies
Algorithmic trading
Intra-day volume

ABSTRACT

This article proposes an empirically tractable way to incorporate intra-day noise into a VWAP trading rule. In volatile markets, news arrives unexpectedly and rapidly. This should influence a trader's trading decisions. However, the literature has not incorporated such information into an algorithmic trading framework. Subsequently, this paper presents a Dynamic VWAP (DVWAP) framework that allows informed traders to utilize random news; and thus, improve trade-execution.

© 2011 Elsevier B.V. All rights reserved.

1. Introduction

This paper analyzes how traders can trade on randomly arriving noise. Volatile markets often feature unexpected or random news. The theoretical literature indicates that such noise should influence traders' trading decisions (see e.g. Foster and Viswanathan, 1993; Foster and Viswanathan, 1996). While Bialkowski et al. (2008) suggest a way to utilize market-wide information, the literature has not suggested a tractable trading framework in which to incorporate stock-specific data. Thus, existing models may not fully exploit intra-day noise. Subsequently, this paper proposes a Dynamic VWAP (DVWAP) framework that allows traders to utilize observable, intra-day, stock-specific information; and thus improve execution.

This paper investigates the following situation. There is an 'large' trader and other market participants. The 'large' trader is so called because he/she wants to shift a large block of stock on a particular day. This can be for liquidity reasons, or for information reasons (as per Inci et al., 2010; Kang, 2010). The trader wants to minimize the transactions costs involved (following Barron and Karpoff, 2004). A key transaction cost is market impact. A way to minimize market impact is to break the large order into smaller blocks and trade the smaller blocks throughout the day. The other market participants trade throughout the day generating noise. This noise may, but need not, reflect information arising throughout the day.

The large trader's goal is optimal execution. If the large trader acts strategically, then optimal execution includes two key components. First, the trade size literature suggests that the large trader should disguise his/her trades by mimicking the actions of the other market participants (or noise traders). The typical way to do this is to break the order into parcels whose size equals the average order size for the stock (following Barclay and Warner, 1993; Chan and Fong, 2000; Cooney and Sias, 2004; Thapa and Poshakwale, 2010). Second, optimal execution should minimize transaction costs (Bauer et al., 2009), operationalized by minimizing the market-impact of a trade (Kissell et al., 2004, Kissell and Malamut, 2006).

Historical VWAP (HVWAP) algorithms attempt to achieve optimal execution. HVWAPs do this by distributing an order throughout the day so as to mimic the historical-average intra-day distribution of volume (see Madhavan, 2002; Kakade et al., 2004; Bialkowski et al., 2008; Kearns, 2008). HVWAP does this as follows. First, over some historical period it obtains intra-day volume data. Second, for each given interval ('bin') b within the day, it calculates the proportion of total daily volume traded in bin b. Thus, for stock i for each day t, this obtains a set of proportions: $\{v_{i,t_0},\ldots,v_{i,t_T}\}$. The proportion in an arbitrary bin b is v_{i,t_b}. Third, it breaks the order into smaller parcels, where the proportion traded in bin b is $v^h_{i,t_b}\%$ of the total order.

The historical approach is problematic. The historical approach does not consider 'random' news. This is problematic since random news can increase (or decrease) the amount of volume traded in

* Address: School of Banking and Finance, Australian School of Business, University of New South Wales, Sydney NSW 2052, Australia. Tel.: +61 2 9385 1000.
E-mail address: M.HumpheryJenner@unsw.edu.au

0378-4266/$ - see front matter © 2011 Elsevier B.V. All rights reserved.
doi:10.1016/j.jbankfin.2011.01.028

any interval b. Thus, if the large trader ignores this volume spike (or lull), then he/she may imperfectly mimic other traders. Further, if there is a volume spike (lull), then the market-impact of an order decreases (increases). Thus, if the large trader ignores the volume spike (or lull) then he/she may trade less (more) than is optimal.

The literature has not resolved this problem. Kissell et al. (2004) and Kissell and Malamut (2006) discuss an Implementation Shortfall ('IS') framework that utilizes contemporaneous market data. The IS approach attempts to predict price movements. It accelerates (decelerates) trade if it predicts that the price will become less (more) favourable in the future. However, the IS framework has two key problems.

First, the IS framework does not constrain the algorithm to distribute trade throughout the day. This ignores the strategic possibility of waiting for more favourable market conditions or future unexpected news releases (as raised in Fong and Liu, 2010). This conflicts with theoretical predictions that, in markets with high levels of unexpected news, large traders should 'strategically' wait for future news arrival, and this should partially reduce the amount traded in a given trading period (see e.g. Foster and Viswanathan, 1996; Back et al., 2000). Second, the IS framework does not attempt to mimic average trade sizes; and thus, may not mimic other traders or minimize transactions costs.

Bialkowski et al. (2008) integrate some dynamic market movements into a VWAP model. Bialkowski et al. (2008) posit that the turnover for stock i, on day t, in time interval b (x_{i,t_b}) is the historical market-average turnover in interval b over the prior month, plus lagged intra-day deviations from the market-average. They examine the stocks in the CAC40 between the start of September 2003 and the end of August 2004.

The Bialkowski et al. (2008) model can help predict departures from historical market-average turnover, but may not identify stock-specific volume spikes. This is primarily because it restricts itself to analysis of turnover. Instead, other intra-day factors, such as raw-volume, spreads, or order-book depth, may influence deviations from historical profiles. DVWAP proposes a general model that can incorporate these variables. And, analysis of turnover or volume is a special case. Second, it focuses on stock-specific deviations from a common market-turnover-factor. Implicitly, this assumes that, on average, a common market factor is a good predictor of stock-specific turnover. Instead, a more stock-specific model may be more appropriate and may improve the ability to detect stock-specific news.

The problems with the historical or traditional approach induce two issues. First, how can the large trader utilize this random intra-day information; and, second, does utilizing the random information improve execution?

This paper proposes a Dynamic VWAP framework (DVWAP) that allows VWAP algorithms to incorporate intra-day stock-specific data. DVWAP does this by both (a) basing itself on the historical intra-day distribution of trade, as in HVWAP, (b) optimizing the use of contemporaneous intra-day information so as to exploit 'random news'. This enhances the approach in Bialkowski et al. (2008) to allow traders to utilize intra-day stock-specific information; and thus, ameliorate information asymmetries. Empirical results for 124 trading days across 130 stocks in the ASX 200 confirm that the DVWAP framework can improve execution over a purely historical approach and has a lower tracking error (and variance thereof) than the Bialkowski et al. (2008) approach does.

This paper contributes to the literature on optimal trading, and especially optimal algorithmic trading. Algorithmic strategies have become increasingly popular (Chew, 2007), with up to 67% of investment managers indicating that they use trading algorithims (Grossman, 2005). VWAP algorithms are one of the most popular algorithms. Thus, improvements to VWAP algorithms contribute to a growing area of financial practice. This paper builds on the insights in Bialkowski et al. (2008) in order to incorporporate stock-specific information; and thus, enhance VWAP algorithms.

The remainder of the paper proceeds as follows. Section 2 discusses a framework in which an informed trader might utilize 'random news' and information asymmetry. Section 3 analyzes the framework. The purpose is to show that the model can improve execution. It examines a model that incorporates 'unexpected' volume data into the framework, and contrasts this approach with that in Bialkowski et al. (2008). Section 4 concludes.

2. What is the DVWAP framework?

The goal is to improve execution. DVWAP induces better execution than HVWAP does if it trades more (less) when the market has more (less) volume. Thus, DVWAP has better execution than HVWAP does if it is more able to track the intra-day dispersion of volume. Therefore, the goal of DVWAP is to minimize tracking error. For a given stock i, the tracking error is the difference between (a) the percentage of total daily volume traded on day t, in bin b (v^a_{i,t_b}), and (b) the percentage of the total order traded in bin b (v^d_{i,t_b}). This induces two issues: First, what optimization problem does DVWAP confront; and second, how does DVWAP resolve this optimization problem.

2.1. What is the optimization problem?

DVWAP aims to reduce tracking error. The basic idea behind DVWAP is to multiply the amount traded under a purely historical approach by a parameter, $\hat{\theta}_{i,t_b}$, in order to make the algorithm track intra-day volume-movements. This parameter derives from observable market-characteristics in the current day. Formally, DVWAP functions by optimizing the amount traded in each bin subject to two constraints. First, DVWAP must spread the order throughout the day. Second, DVWAP must ensure completion of the order and must not trade more than the total order size.

To set ideas, define the following variables. Let the trader have $X_{i,t}$ units of stock i to trade on day t. Let the proportion of $X_{i,t}$ traded in bin b be v^d_{i,t_b}, where d denotes dynamic. Let the total volume for stock i on day t be $V_{i,t}$. The actual proportion of $V_{i,t}$ traded in bin b is v^a_{i,t_b}, where a stands for actual. This is only observable at the end of the day. Let the historical average proportion of $V_{i,t}$ traded in bin b be v^h_{i,t_b} in bin b, where h denotes historical.

HVWAP functions as follows. Take the historical average proportion of volume traded in bin b: v^h_{i,t_b}. Now, the proportion of the total order, $X_{i,t}$, traded in bin b equals the historical average proportion of volume traded in bin b; that is, the amount traded in bin b is $v^h_{i,t_b} \times X_{i,t}$.

DVWAP aims to optimize the amount traded in each bin b. The optimization problem is in Eq. (1). Acknowledge that there is some proportion of the total order ($X_{i,t}$) after which the algorithm should apportion the remaining amount of the total order so as to spread the order throughout the whole day. For example, if $\gamma = 80\%$, then after the algorithm has traded $0.8 \times X_{i,t}$, then it should apportion the remaining trade across the remaining bins.

The optimization problem is as follows. Break the trading day into two parts: Part (a) is when the algorithm has traded less than γ. Part (b) is when the algorithm has traded more than γ.

For Part (a); if the algorithm has traded less than γ, then DVWAP starts with the amount that HVWAP would trade in bin t. Recall that in bin b HVWAP trades v^h_{i,t_b} of the total order. DVWAP then estimates the ratio $\hat{\theta}_{i,(t-k)_b} = \left(\frac{\widehat{v^a_{i,(t-k)_b}}}{v^h_{i,(t-k)_b}}\right)$. DVWAP then multiplies v^h_{i,t_b} by a parameter $\hat{\theta}_{i,t_b}$. The parameter is estimated on historical

'training' data and operates to minimize the difference between the amount traded in b and the actual volume in bin b. The model constrains $\hat{\theta}_{i,t_b}$ between $[\underline{\theta}, \bar{\theta}]$ to ensure that the algorithm does not trade too much in a given bin. Section 2.2 discusses the estimation of $\underline{\theta}$ and $\bar{\theta}$. Section 3.3 discusses the impact on DVWAP's performance of different $\underline{\theta}$ and $\bar{\theta}$ parameters.

For Part (b); if the algorithm has traded more than γ, then DVWAP apportions the remainder of the order across the remaining bins. It functions as follows. First, determine the amount left to trade. Second, for each remaining bin, calculate the percentage that HVWAP would trade in that bin; and, divide this by the total that HVWAP would trade over the remaining bins. This creates a 'weighting parameter' for the remaining bins. Third, multiply the total left to trade by the 'weighting parameter'.

This situation induces the minimization problem:

$$\min_{\theta,\gamma} \left[\mathbb{I}_{\gamma^c} \hat{\theta}_t v^h_{i,t_b} + \mathbb{I}_{\gamma} \left(\frac{v^h_{i,t_b}}{\sum_{b=t_\gamma}^{T} v^h_{i,t_b}} \right) \right] - v^a_{i,t_b} \tag{1}$$

$$\text{subject to} \quad \begin{array}{l} \sum_{b=1}^{T} v^d_{i,t_b} = 1 \\ v^d_{i,t_b} \leqslant \bar{\theta} \times v^h_{i,t_b} \\ -v^d_{i,t_b} \leqslant -\underline{\theta} \times v^h_{i,t_b} \end{array} \tag{2}$$

Here, $\mathbb{I}_{\gamma^c}$ (and $\mathbb{I}_{\gamma}$) is an indicator that the algorithm has traded less than (more than) γ% of the total order. The bin t_γ is the bin after which the trader has traded γ% of the total order. v^d_{i,t_b} is the percentage of the total order that DVWAP trades in bin b. v^h_{i,t_b} is the percentage of the total order that HVWAP would trade in bin b; v^a_{i,t_b} is the actual percentage of total daily volume that trades in bin b.

The issue is how to estimate the parameters. There are two key parameters of interest, θ and γ. Ideally, θ and γ solve Eq. (1). If the actual percentage of total daily volume that trades in bin b, v^a_{i,t_b}, were known, then the problem is trivial. However, the total daily volume is only observable at the end of the day. Thus, v^a_{i,t_b} is not ex ante observable. Thus, the goal is to use training (or historical) data to estimate a functional form for $\hat{\theta}$. Specifically DVWAP posits that $\hat{\theta}$ is a function of observable intra-day stock characteristics, denoted $\mathbf{x}_{i,(t-k)_{b-1}}$. Thus, the parameter $\hat{\theta}_{i,t_b}$ is the parameter derived from $\hat{\theta}_{i,t_b} = f(\mathbf{x}_{i,t_{b-1}})$ such that $v^d_{i,t_b} = \theta v^h_{i,t_b}$, and $\hat{\theta}_{i,t_b} \in [\underline{\theta}, \bar{\theta}]$.

The issue is how to solve the optimization problem. The optimization problem in Eq. (1) lacks an ex-ante analytical solution since the proportion of total daily volume traded in bin b (v^a_{i,t_b}) depends on the total daily volume, which is only know after the close of trade. Thus, a trading algorithm cannot incorporate v^a_{i,t_b} into an ex ante solvable function. But, it can (a) use training data to estimate the function $(v^a_{i,t_b}/v^h_{i,t_b}) = f(\mathbf{x}_{i,t_{b-1}})$, and (b) use contemporaneous market data to predict $\widehat{v^a_{i,t_b}/v^h_{i,t_b}}$ in the present day. Section 2.2 formalizes this practicable solution.

2.2. What is a practicable solution?

DVWAP aims minimize tracking error by using intra-day market data. The solution is to generate the 'dynamic VWAP recursion'. The recursion summarizes the modeling process. Define the following notation: there are N bins during the day; v^h_{i,t_b} is the amount that HVWAP would trade in bin b; v^d_{i,t_b} is the amount that DVWAP trades in bin b; $X_{i,t}$ is the total amount of stock i to trade on day t. Preliminarily, set γ, which is a risk-tolerance parameter that tells the algorithm to revert to HVWAP after we have traded γ% of the total order. Further, set the risk-parameters $[\underline{\theta}, \bar{\theta}]$, which define the allowable deviation from the HVWAP trading profile. That is, they bind v^d_{i,t_b} to be above $\underline{\theta} \times v^d_{i,t_b}$ and below $\bar{\theta} \times v^d_{i,t_b}$. These are also risk-tolerance parameters that the trader can set to match the level of desired risk.

There are two states: (1) where the trader has traded less than γ% of the total order, and (2) where the trader has traded more than γ% of the total order. Consider situation first situation. Proceed as follows:

First, use historical intra-day volume data to find v^h_{i,t_b}, the proportion of total daily volume that ordinarily trades in bin b. Typically, this is average proportion of actual daily volume over the past month.

Second, let $\mathbf{x}_{i,t_{b-1}}$ be a vector of intra-day stock characteristics from day t that are expected to influence intra-day volume. These are characteristics observable in bin b (hence the subscript $b-1$); and thus, must logically date from momentarily before bin b. Use historical data (from k days before the current day t) to model the ratio of actual volume to 'historical' volume. That is, model $v^a_{i,t-k_b}/v^h_{i,t-k_b} = f(\mathbf{x}_{i,(t-k)_{b-1}})$, where $f : \mathbb{R}^n \to \mathbb{R}$ maps n observable stock-characteristics to the volume-ratio. This gives the functional form of f that DVWAP will uses on a vector of characteristics from the current day t.

Third, consider bin b on the current day t. Take the vector of characteristics $\mathbf{x}_{i,(t)_{b-1}}$. Apply the function f to these characteristics to yield $\hat{\theta}_{i,t_b} = f(\mathbf{x}_{i,(t-k)_{b-1}})$. Acknowledge that the parameter $\hat{\theta}$ is simply the predicted ratio of actual volume to historical volume. That is the predicted value of: $v^a_{i,t_b}/v^h_{i,t_b}$. Apply the parameter $\hat{\theta}$ to v^h_{i,t_b} to yield $v^d_{i,t_b} = \hat{\theta}_{i,t_b} v^h_{i,t_b}$.

Now, consider the situation where DVWAP has traded more than γ% of the total order. DVWAP takes the amount that HVWAP would trade in bin b and divides it by the total amount that HVWAP would trade in the rest of the day. Thus, the amount traded is $v^h_{i,t_b}/\sum_{j=0}^{N-b} v^h_{i,t_b}$. This situation induces the dynamic VWAP recursion.

$$\begin{aligned} v^d_{i,t_b} &= \mathbb{I}_{\gamma^c} v^h_{i,t_b} \hat{\theta}_{i,t_b} + \mathbb{I}_{\gamma} \frac{v^h_{i,t_b}}{\sum_{j=0}^{N-b} v^h_{i,t_b}} \\ &= \mathbb{I}_{\gamma^c} v^h_{i,t_b} f(\mathbf{x}_{i,(t-k)_{b-1}}) + \mathbb{I}_{\gamma} \frac{v^h_{i,t_b}}{\sum_{j=0}^{N-b} v^h_{i,t_b}} \end{aligned} \tag{3}$$

In words, the recursion says: First, if we have traded less than γ% of the total order, then we trade $v^h_{i,t_b} f(\mathbf{x}_{i,(t-k)_{b-1}})$ in bin b. This amount is the amount traded under HVWAP (i.e. v^h_{i,t_b}) multiplied by an adjustment parameter (i.e. $f(\mathbf{x}_{i,(t-k)_{b-1}})$). The adjustment parameter has two components, the information from the prior bin ($\mathbf{x}_{i,(t-k)_{b-1}}$) and the functional form ($f : \mathbb{R}^n \to \mathbb{R}$). The information from the prior bin is from day t and can include several lags of volume movements, turnover, or order-book information. The functional form f derives from historical modeling of the equation $v^a_{i,t_b}/v^h_{i,t_b} = f(\mathbf{x}_{i,(t-k)_{b-1}})$ on historical data.

In the second state, if we have traded more than γ% of the total order, then we default back to a HVWAP type approach in order to avoid over-trading. Here, there are $N-b$ bins left in the day. Thus, for bin b we would take the amount HVWAP would trade in bin b (i.e. v^h_{i,t_b}) and divide it by the amount that HVWAP would trade over the remainder of the day (i.e. $\sum_{j=0}^{N-b} v^h_{i,t_b}$).

This section defined the DVWAP relationship. The issue is then whether it works to improve trade-execution. Section 3 applies it to a simple example and examines it in the context of the approach taken in Bialkowski et al. (2008).

3. Application

This section applies DVWAP within a practical setting. The purpose is to test if DVWAP improves execution, and to contrast the DVWAP approach to the approach in Bialkowski et al. (2008). Section 3.1 contains an intentionally simplistic example. It models the ratio $\widehat{v^a_{i,t_b}/v^h_{i,t_b}}$ as a function of 'unexpected' or 'decomposed' intra-day volume-movements. Section examines the ARMA (1,1)

and SETAR approaches in Bialkowski et al. (2008). The results indicate that DVWAP has lower tracking error than HVWAP, or the approaches in Bialkowski et al. (2008) do, and DVWAP's tracking error has lower variance (or is more stable) than that the tracking errors in Bialkowski et al. (2008) models have.

3.1. The volume example

The literature suggests that past volume information may predict future volume spikes or lulls. The theoretical model in Foster and Viswanathan (1993) indicates that unanticipated news affects stock volumes, and that volumes are autocorrelated. This implies that past volume information may help predict future volume spikes/lulls. Brooks et al. (2003) study 21 unanticipated CEO deaths, and find that the unanticipated news/information can influence stock volume for up to an hour after the news release. Nofsinger and Prucyk (2003) find a two hour delay between an unanticipated news announcement, and a subsequent volume increase. They also find persistence in volume increases. These results suggest that lagged volume may help predict future volume movements.[1]

This induces three key issues. First, does contemporaneous volume convey useful information; second, how would DVWAP use this information; and third, does using the information allow DVWAP to out-perform HVWAP?

This section analyzes volume information using Reuters data for the ASX. The sample spans 4 April 2008 to 1 October 2008 (124 trading days), with the data from April being used solely to calculate parameters for May. The sample contains 197 stocks in the ASX 200. The sample loses three stocks due to data availability. Each day is split into 5 min bins. Since the ASX commences trade for stocks at different times (depending on the stock's position in the alphabet), this yields either 71 or 72 bins for each stock. In total, this yields 1,484,980 estimates. Data analysis uses *R*, SAS and STATA.

The results indicate that DVWAP should incorporate lagged volume information; and that, DVWAP performs better than HVWAP does when DVWAP incorporates contemporaneous volume information.

3.1.1. Would incorporating contemporaneous volume improve VWAP algorithms?

Incorporating contemporaneous volume is useful only if that volume conveys information that can improve the algorithm's ability to predict future volume patterns. This occurs if the autocorrelation structure of daily volume patterns differs from that of historical averages on which HVWAP bases itself, and it is possible to use this information to predict volume movements.

The results indicate that the autocorrelation structure of volume patterns differs between actual volumes and historical average volume profiles, and this deviation conveys information. Fig. 1 contains the intra-day autocorrelation for the stock AAC. Panels (a) and (b) contain the autocorrelation structure for the historical average volume profile. The panels indicate that the historical volume profile assumes significant autocorrelation (consistent with Bialkowski et al. (2008)). However, Panels (c) and (d), which contain the autocorrelation structure for AAC on 9 April 2008 indicate that the autocorrelation structure for a given stock on a given day may not follow the historical profile. This indicates that the deviation from the historical profile may contain information about future volume movements.

[1] I also note that other factors, such as access to the derivatives market (see Chang et al., 2009; Muck, 2010), could influence trading behavior; and thus, help to predict future volume movements. The DVWAP model is not limited to using only historical volume information.

The deviation from the historical profile conveys information. Panels (e) and (f) examine 'Decomposed Volume'. Decomposed volume is he difference between the actual volume in a given 5-min bin b and volume expected in bin b if trade followed a historical volume pattern. Eq. (4) summarizes this.

$$\begin{aligned}(\text{Decomposed Volume})_{i,t_b} &= n^a_{i,t_b} - \hat{n}^a_{i,t_b}\\ \text{where}&\\ n^a_{i,t_b} &= \text{Actual number of shares traded in bin } b\\ \hat{n}^a_{i,t_b} &= \text{Expected amount traded in interval } b\\ &= n^a_{i,t_{b-1}} \times \frac{v^h_{i,t_b}}{v^h_{i,t_{b-1}}}\end{aligned} \tag{4}$$

Deviations from the historical profile may convey information. Panels (e) and (f) indicate that deviations from the historical profile at time $b-1$ convey information about deviations at time b. Thus, lagged deviations may convey useful information about future volume movements. This indicates focusing on the 'unexpected' part of volume may improve the performance of VWAP algorithms.

3.1.2. How does DVWAP incorporate contemporaneous volume?

DVWAP incorporates contemporaneous volume information into a VWAP framework. Three issues obtain: First, what is the functional form of θ. Second, what is the optimal value of γ. Third, what are the values of $\underline{\theta}$ and $\bar{\theta}$.

The $\hat{\theta}_{i,t_b}$ parameter determines the amount traded in bin b. To obtain this, proceed as follows. First, the model posits (following Section 3.1.1) that a predictor of volume may be the lagged difference between actual volume and expected volume, denoted ξ_{i,t_b}. The model takes first-differences to remove intra-day seasonality of volume.

Second, using the month before day t, the model estimates the relation between lagged intra-day volume and current intra-day volume. Let, $\mathbf{x}_{i,(t-k)_{b-1}}$ be the stock-specific volume variable on the prior month. And, β_i is a vector of coefficients. Thus, it estimates Eq. (5).

$$(v^a_{i,(t-k)_b}/v^h_{i,(t-k)_b}) = \mathbf{x}_{i,(t-k)_{b-1}}\beta_i \tag{5}$$

In the specific case of decomposed volume the model is as follows. For each stock and each month, this paper estimates rolling regressions of Eq. (7). The parameters for stock j in month m derive from the estimation of Eq. (7) for stock j in month $m-1$. This yields 1370 sets of coefficients. The model uses three lagged differences in order to balance the information gained from more lags with the loss of observations due to the lagged starting period.

$$\frac{v^a_{i,(t-k)_b}}{v^h_{i,(t-k)_b}} = \beta_0 + \beta_1\Delta\xi_{i,(t-k)_{b-1}} + \beta_2\Delta\xi_{i,(t-k)_{b-2}} + \beta_3\Delta\xi_{i,(t-k)_{b-3}} + \varepsilon_t \tag{6}$$

where

$$\begin{aligned}\xi_{i,(t-k)_{b-1}} &= \frac{n_{i,(t-k)_{b-1}} - \hat{n}_{i,(t-k)_{b-1}}}{\hat{n}_{i,(t-k)_{b-1}}}\\ n_{i,(t-k)_{b-1}} &= \text{Total volume in bin } b-1\\ \hat{n}_{i,(t-k)_{b-1}} &= \text{Predicted volume in bin } b-1\\ &= n^a_{i,(t-k)_{b-2}} \times \frac{v^h_{i,(t-k)_{b-1}}}{v^h_{i,(t-k)_{b-2}}}\end{aligned} \tag{7}$$

The γ parameter derives from data over the prior month. The optimal γ minimizes the absolute deviation between (a) the percentage of the order traded in bin b, v^d_{i,t_b}, and (b) the percentage of total daily volume traded in bin b, v^a_{i,t_b}. That is, given the estimation of Eq. (7), γ should minimize $\left|v^d_{i,t_b} - v^a_{i,t_b}\right|$. Here, this is approximately $\gamma = 0.8$ on average. Thus, this paper utilizes the value $\gamma = 0.8$ for future equations.

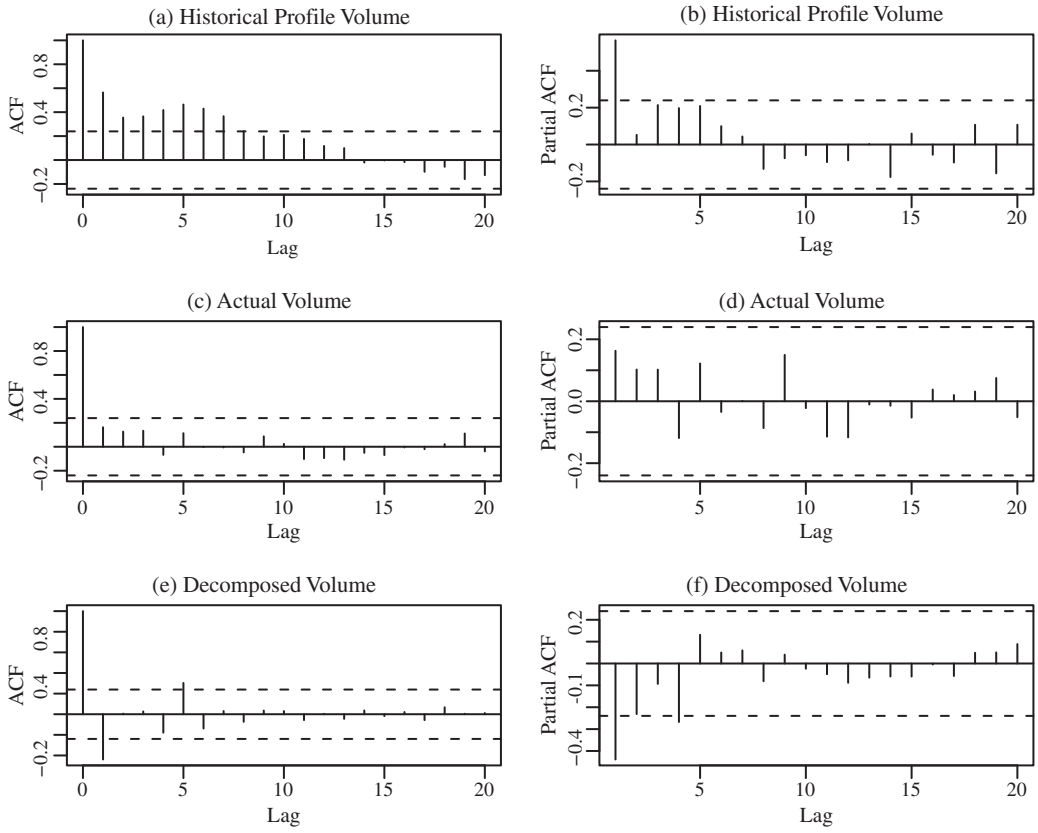


Fig. 1. Autocorrelation functions and partial autocorrelation functions for AAC.

The limits $\underline{\theta}$ and $\bar{\theta}$ derive from data over the prior month. The optimal $[\underline{\theta}, \bar{\theta}]$ minimize the absolute deviation between $\theta_t v^h_{i,t_b}$ and v^a_{i,t_b} $\forall \theta \in [\underline{\theta}, \bar{\theta}]$. Here, the approximate average values are $\underline{\theta} = 0.9$ and $\bar{\theta} = 1.1$.

This motivates examination of whether DVWAP tracks actual daily volume better than HVWAP does. DVWAP does this if it has a better tracking performance on average (examined in Section 3.1.3), and across a variety of market and stock conditions (examined in Section 3.1.4).

3.1.3. How do the changes perform on average?

The results indicate that DVWAP performs better than HVWAP does. The trading goal is to trade more (less) if overall market volume is higher (lower). This occurs if the algorithm tracks the intra-day dispersion of trade. Thus, DVWAP achieves better execution than HVWAP does if it reduces the absolute average difference between (a) the proportion traded in bin b and (b) the proportion of daily volume that trades in bin b. Thus, DVWAP performs better than HVWAP does if both (a) it has lower tracking error on average, and (b) it's tracking error is lower in the majority of bins.

This paper compares DVWAP's performance with HVWAP's performance by considering its tracking error. It focuses on out-of-sample results, although reports soem in-sample results for robustness. The tracking errors are defined as the absolute difference between the percentage of the order executed in bin b versus the percentage of market volume traded in bin b. The absolute difference is more appropriate than the straight difference since absolute differences penalize positive and negative differences whereas straight mean differences allow negative differences to cancel positive differences. Thus, the tracking errors are:

$$\epsilon^h_{i,t_b} = \left| v^h_{i,t_b} - v^a_{i,t_b} \right| \tag{8}$$

$$\epsilon^d_{i,t_b} = \left| v^d_{i,t_b} - v^a_{i,t_b} \right| \tag{9}$$

The first issue is if DVWAP reduces tracking error on average. The results indicate that, on average, DVWAP tracks actual daily volume better than HVWAP does. Table 1 contains mean and

Table 1
Univariate statistics for the whole sample.

	ϵ^d_{i,t_b} (1)	ϵ^h_{i,t_b} (2)	$\epsilon^d_{i,t_b} - \epsilon^h_{i,t_b}$ (1) − (2)
Panel A: 5 min bins			
Mean	1.068[a]	1.084[a]	−0.015[a]
Median	0.708[a]	0.726[a]	−0.001[a]
Panel B: 15 min bins			
Mean	2.177[a]	2.253[a]	−0.076[a]
Median	1.512[a]	1.588[a]	−0.033[a]
Panel D: 60 min bins			
Mean	3.227[a]	3.398[a]	−0.171[a]
Median	2.344[a]	2.505[a]	−0.172[a]

This Table contains univariate statistics for the whole sample. The statistics are for the tracking error for DVWAP and HVWAP. Tracking error for DVWAP $\left(\epsilon^d_{i,t_b}\right)$ or HVWAP $\left(\epsilon^h_{i,t_b}\right)$ is the absolute difference between the proportion of the order that DVWAP or HVWAP trades in bin t, and the proportion of total daily volume that trades in bin b. Thus, for DVWAP, tracking error is: $\epsilon^d_{i,t_b} = \left| v^d_{i,t_b} - v^a_{i,t_b} \right|$, where v^d_{i,t_b} is the percentage of total that DVWAP trades in bin b and v^a_{i,t_b} is the percent of total daily volume that trades in bin b. Similarly, tracking error for HVWAP is $\epsilon^h_{i,t_b} = \left| v^h_{i,t_b} - v^a_{i,t_b} \right|$, where v^h_{i,t_b} is the percent of total that HVWAP trades in bin b. The superscript a denotes significance at 1% on a difference in mean test (for means) and on a wilcoxon sign rank test (for medians).

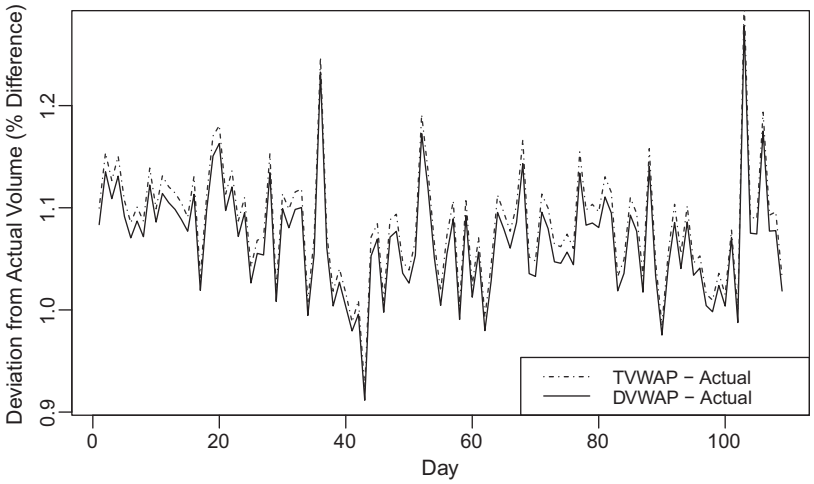


Fig. 2. Average daily deviations from actual volume. This figure plots the daily values of (HVWAP-Actual) and (DVWAP-Actual).

median tracking errors. Panel A indicates that ϵ^{d}_{i,t_b} is lower than is ϵ^{h}_{i,t_b} (at 1% significance) for 5-min bins. Panels B and C confirm these results for 15-min and 60-min bins, respectively.

Fig. 2 illustrates the average deviation of DVWAP from actual and HVWAP from actual for each individual trading day. The large deviation spikes coincide with significant corporate announcements common to many stocks (such as common earnings announcements and annual report releases). In these situations, neither HVWAP nor DVWAP can anticipate ex-ante the large volume spike or the precise timing thereof. Overall, the results indicate that DVWAP tracks actual market volume better than HVWAP does.

The second issue is if DVWAP reduces tracking error in the majority of bins. For bin *b*, DVWAP predicts volume movements better than HVWAP does if DVWAP's tracking error in bin *b* is lower than HVWAP's is. Thus, DVWAP predicts volume movements better than HVWAP does if there are more bins for which DVWAP has lower tracking error. Thus, Table 4 examines the percentage of bins for which DVWAP has a lower tracking error.

The results in Table 4 indicate that DVWAP in the majority of bins. The results indicate that DVWAP has lower tracking error than HVWAP does in approximately 65% of all bins. This result is robust to the 'spikiness' of the day (the intra-day variance of volume between bins), the historical intra-day variance of volume ('VolVar'), or the historical inter-day price variance ('PriceVar').

The results overall indicate that DVWAP performs better than HVWAP does.

3.1.4. Is DVWAP's performance robust to stock and market characteristics?

DVWAP has lower tracking error than HVWAP does on average. However, DVWAP performs better than HVWAP does only if it can out-perform HVWAP across a variety of market conditions and stock characteristics. This paper assesses this by first sorting the data by factors that might affect DVWAP's performance, and then through multivariate regression analysis.

Univariate analysis indicates that DVWAP tracks actual market volume better than HVWAP does. To assess the performance on volatile days, this paper sorts the sample into groups based on contemporaneous volume volatility (the 'spikiness' of a day), historical intra-day volume volatility ('VolVar'), and historical inter-day price-volatility ('PriceVar'). This is important due to the reported volume-volatility relationship (see for example Xu et al., 2006; Darrat et al., 2007).

Table 2 contains the out-of-sample tracking error sorted by the measures of variance. The results indicate that, after sorting by all types of volatility, DVWAP has lower tracking error than HVWAP does (at 1% significance). Table 3 contains in-sample tracking errors. The in-sample results are consistent with the out-of-sample results. Unsurprisingly, DVWAP performs better (has a lower tracking error) in in-sample tests.

Multivariate analysis confirms that DVWAP's tracking performance is robust to stock and market characteristics. Four key historical may influence DVWAP's performance ability.

First, if the stock's intra-day volume is more volatile or if liquidity is more unpredictable, then DVWAP is more likely to make a prediction-error, and performance should worsen. The intra-day variance in the level of volume between five-minute bins ('VolVar') proxies for this (consistent with Johnson, 2008).

Second, the average trade size may influence DVWAP's performance. Average trade size ('lnTradeSize') can indicate the level of activity in a stock and the potential for informed trade.[2] However, it may also indicate additional volatility. If average trade size represents additional activity or informed trade, then it may provide DVWAP with more information on which to predict volume movements; and thus, may improve DVWAP's performance. However, if average trade size indicates noise, then it may reduce the informativeness of past volume information.

Third, if trade in the stock is more active, then a volume spike is less likely to persist over time. This may harm DVWAP's ability to trade upon predicted volume spikes. Thus, the models include 'lnDailyTrades': the daily-average number of trades. Fourth, if the stock's price is more volatile, then the level of volume may be more volatile. And, if the volume is more volatile, then DVWAP may be less able to predict future volume. Thus, the models include 'PriceVar', the historical inter-day variance of the stock price.

Fifth and similarly, if the stock is more liquid, then market-prices are likely more efficient (Comerton-Forde, 1999); and thus, volume spikes are less likely to persist. This may reduce DVWAP's performance. This paper proxies liquidity by examining the average number of trades at the best bid or ask price scaled by the average daily volume ('BaRatio'): If there are relatively more trades at

[2] Barclay and Warner (1993, **?**) discuss the relation between volume and informed trade in more detail.

Table 2
Univariate results.

Group	Spikiness			VolVar			PriceVar		
	ϵ^d_{i,t_b}	ϵ^h_{i,t_b}	$\epsilon^d_{i,t_b} - \epsilon^h_{i,t_b}$	ϵ^d_{i,t_b}	ϵ^h_{i,t_b}	$\epsilon^d_{i,t_b} - \epsilon^h_{i,t_b}$	ϵ^d_{i,t_b}	ϵ^h_{i,t_b}	$\epsilon^d_{i,t_b} - \epsilon^h_{i,t_b}$
1	1.073[a]	1.091[a]	−0.018[a]	0.72[a]	0.743[a]	−0.023[a]	0.987[a]	0.998[a]	−0.011[a]
	0.751[a]	*0.774*[a]	*−0.023*[a]	*0.508*[a]	*0.533*[a]	*−0.024*[a]	*0.682*[a]	*0.693*[a]	*−0.011*[a]
2	1.036[a]	1.052[a]	−0.016[a]	0.821[a]	0.838[a]	−0.017[a]	0.966[a]	0.984[a]	−0.018[a]
	0.72[a]	*0.739*[a]	*−0.019*[a]	*0.576*[a]	*0.593*[a]	*−0.016*[a]	*0.657*[a]	*0.676*[a]	*−0.019*[a]
3	1.035[a]	1.052[a]	−0.017[a]	0.997[a]	1.006[a]	−0.009[a]	0.992[a]	1.004[a]	−0.012[a]
	0.712[a]	*0.732*[a]	*−0.02*[a]	*0.698*[a]	*0.709*[a]	*−0.011*[a]	*0.653*[a]	*0.669*[a]	*−0.016*[a]
4	1.029[a]	1.045[a]	−0.016[a]	1.046[a]	1.057[a]	−0.011[a]	1.047[a]	1.06[a]	−0.013[a]
	0.702[a]	*0.72*[a]	*−0.019*[a]	*0.698*[a]	*0.709*[a]	*−0.012*[a]	*0.709*[a]	*0.725*[a]	*−0.016*[a]
5	0.69[a]	0.708[a]	−0.018[a]	0.749[a]	0.765[a]	−0.016[a]	0.655[a]	0.676[a]	−0.021[a]
	0.69[a]	*0.708*[a]	*−0.018*[a]	*0.749*[a]	*0.765*[a]	*−0.016*[a]	*0.655*[a]	*0.676*[a]	*−0.021*[a]
6	1.024[a]	1.039[a]	−0.015[a]	1.118[a]	1.132[a]	−0.013[a]	1.094[a]	1.106[a]	−0.013[a]
	0.683[a]	*0.7*[a]	*−0.017*[a]	*0.75*[a]	*0.767*[a]	*−0.017*[a]	*0.736*[a]	*0.752*[a]	*−0.016*[a]
7	1.033[a]	1.048[a]	−0.015[a]	1.075[a]	1.088[a]	−0.013[a]	1.017[a]	1.032[a]	−0.016[a]
	0.684[a]	*0.701*[a]	*−0.017*[a]	*0.732*[a]	*0.749*[a]	*−0.017*[a]	*0.674*[a]	*0.691*[a]	*−0.017*[a]
8	1.04[a]	1.055[a]	−0.015[a]	1.23[a]	1.246[a]	−0.016[a]	1.083[a]	1.097[a]	−0.015[a]
	0.671[a]	*0.688*[a]	*−0.017*[a]	*0.798*[a]	*0.819*[a]	*−0.021*[a]	*0.729*[a]	*0.747*[a]	*−0.018*[a]
9	1.113[a]	1.126[a]	−0.013[a]	1.113[a]	1.129[a]	−0.015[a]	1.109[a]	1.122[a]	−0.013[a]
	0.703[a]	*0.718*[a]	*−0.015*[a]	*0.755*[a]	*0.774*[a]	*−0.019*[a]	*0.738*[a]	*0.755*[a]	*−0.016*[a]
10	1.28[a]	1.293[a]	−0.014[a]	1.182[a]	1.199[a]	−0.017[a]	1.154[a]	1.173[a]	−0.019[a]
	0.773[a]	*0.787*[a]	*−0.014*[a]	*0.738*[a]	*0.756*[a]	*−0.018*[a]	*0.73*[a]	*0.753*[a]	*−0.023*[a]

This table contains univariate statistics for the whole sample sorted by contemporaneous intra-day volume variance (spikiness), historical intra-day volume variance (VolVar) and historical inter-day price variance (PriceVar). Group 1 (10) has the lowest (highest) spikiness, VolVar, or PriceVar. Intra-day volume variance is the variance of the number of shares traded in each bin $b \in [0,T]$. ϵ^d_{i,t_b} is the difference between the percent of the total order traded in bin b and the percent of total daily volume traded in bin b: $\epsilon^d_{i,t_b} = \left|v^d_{i,t_b} - v^a_{i,t_b}\right|$. Similarly, $\epsilon^h_{i,t_b} = \left|v^h_{i,t_b} - v^a_{i,t_b}\right|$. Numbers in normal type are means. Numbers in italics are medians. The superscript *a* denotes significance at 1% in *t*-test (for means) and a wilcoxon sign rank test (for medians).

Table 3
In sample univariate results.

Group	Spikiness			VolVar			PriceVar		
	ϵ^d_{i,t_b}	ϵ^h_{i,t_b}	$\epsilon^d_{i,t_b} - \epsilon^h_{i,t_b}$	ϵ^d_{i,t_b}	ϵ^h_{i,t_b}	$\epsilon^d_{i,t_b} - \epsilon^h_{i,t_b}$	ϵ^d_{i,t_b}	ϵ^h_{i,t_b}	$\epsilon^d_{i,t_b} - \epsilon^h_{i,t_b}$
1	1.070[a]	1.091[a]	−0.021[a]	0.713[a]	0.743[a]	−0.030[a]	0.983[a]	0.998[a]	−0.014[a]
2	1.029[a]	1.052[a]	−0.023[a]	0.811[a]	0.838[a]	−0.027[a]	0.957[a]	0.984[a]	−0.027[a]
3	1.030[a]	1.052[a]	−0.022[a]	0.992[a]	1.006[a]	−0.014[a]	0.987[a]	1.004[a]	−0.017[a]
4	1.025[a]	1.045[a]	−0.019[a]	1.042[a]	1.057[a]	−0.015[a]	1.046[a]	1.060[a]	−0.015[a]
5	1.018[a]	1.034[a]	−0.016[a]	1.149[a]	1.162[a]	−0.014[a]	1.006[a]	1.031[a]	−0.026[a]
6	1.020[a]	1.039[a]	−0.019[a]	1.110[a]	1.132[a]	−0.022[a]	1.083[a]	1.106[a]	−0.023[a]
7	1.026[a]	1.048[a]	−0.022[a]	1.070[a]	1.088[a]	−0.018[a]	1.007[a]	1.032[a]	−0.026[a]
8	1.035[a]	1.055[a]	−0.020[a]	1.230[a]	1.246[a]	−0.016[a]	1.078[a]	1.097[a]	−0.019[a]
9	1.105[a]	1.126[a]	−0.021[a]	1.108[a]	1.129[a]	−0.021[a]	1.108[a]	1.122[a]	−0.015[a]
10	1.271[a]	1.293[a]	−0.023[a]	1.173[a]	1.199[a]	−0.025[a]	1.144[a]	1.173[a]	−0.029[a]

This table contains in sample univariate statistics for the whole sample sorted by contemporaneous intra-day volume variance (spikiness), historical intra-day volume variance (VolVar) and historical inter-day price variance (PriceVar). Group 1 (10) has the lowest (highest) spikiness, VolVar, or PriceVar. Intra-day volume variance is the variance of the number of shares traded in each bin $b \in [0,T]$. ϵ^d_{i,t_b} is the difference between the percent of the total order traded in bin b and the percent of total daily volume traded in bin b: $\epsilon^d_{i,t_b} = \left|v^d_{i,t_b} - v^a_{i,t_b}\right|$. Similarly, $\epsilon^h_{i,t_b} = \left|v^h_{i,t_b} - v^a_{i,t_b}\right|$. Numbers in normal type are means. Numbers in italics are medians. The superscript *a* denotes significance at 1% in *t*-test (for means) and a wilcoxon sign rank test (for medians).

the best bid (or ask), then the stock is likely more liquid (Cao et al., 2004). Table 5 defines these variables in detail.

This paper uses three models to examine the relation between DVWAP's out-performance and stock characteristics. First, Eq. (10) regresses the above variables on the absolute difference between DVWAP's tracking error $\left(\epsilon^d_{i,t_b} = \left|v^d_{i,t_b} - v^a_{i,t_b}\right|\right)$ and HVWAP's tracking error $\left(\epsilon^h_{i,t_b} = \left|v^h_{i,t_b} - v^a_{i,t_b}\right|\right)$. The model uses an OLS regression with standard errors clustered by stock and by date (consistent with Petersen, 2009). Second, Eq. (11) estimates a probit model where the dependent variable (IMPROVE) equals 1 if $\left|v^h_{i,t_b} - v^a_{i,t_b}\right|$ exceeds $\left|v^d_{i,t_b} - v^a_{i,t_b}\right|$ and equals 0 otherwise. Third, Eq. (12) estimates a probit model where the dependent variable (CORRECT) equals 1 if the model correctly predicts that volume will be above or below the historical profile. Table 5 defines the variables in detail.

$$\text{DIFF} = \beta_0 + \beta_1 \text{VolVar} + \beta_2 \text{lnTradeSize} + \beta_3 \text{lnDailyTrades} + \beta_4 \text{BaRatio} + \beta_5 \text{PriceVar} + \varepsilon \quad (10)$$

$$\text{IMPROVE} = \beta_0 + \beta_1 \text{VolVar} + \beta_2\ \text{lnTradeSize} + \beta_3 \text{lnDailyTrades} + \beta_4 \text{BaRatio} + \beta_5 \text{PriceVar} + \varepsilon \quad (11)$$

$$\text{CORRECT} = \beta_0 + \beta_1 \text{VolVar} + \beta_2\ \text{lnTradeSize} + \beta_3 \text{lnDailyTrades} + \beta_4 \text{BaRatio} + \beta_5 \text{PriceVar} + \varepsilon \quad (12)$$

The results indicate that DVWAP's performance is robust to stock-level characteristics. Table 6 Column 1 indicates that DVWAP's performance decreases with the level of sophisticated trade (proxied by BaRatio), at 1% significance, and inter-day price-volatility (PriceVar). However, DVWAP's performance increases with intra-day volume-volatility (VolVar), indicating that

Table 4
Improvements probabilities.

Group	Spikiness		VolVar		PriceVar	
	$\epsilon^d_{i,t_b} > \epsilon^h_{i,t_b}$	$\epsilon^d_{i,t_b} \leqslant \epsilon^h_{i,t_b}$	$\epsilon^d_{i,t_b} > \epsilon^h_{i,t_b}$	$\epsilon^d_{i,t_b} \leqslant \epsilon^h_{i,t_b}$	$\epsilon^d_{i,t_b} > \epsilon^h_{i,t_b}$	$\epsilon^d_{i,t_b} \leqslant \epsilon^h_{i,t_b}$
1	32.71[a]	67.29[a]	34.99[a]	65.01[a]	34.43[a]	65.57[a]
2	33.39[a]	66.61[a]	34.70[a]	65.30[a]	33.25[a]	66.75[a]
3	33.35[a]	66.65[a]	34.37[a]	65.63[a]	33.84[a]	66.16[a]
4	33.81[a]	66.19[a]	34.94[a]	65.06[a]	33.87[a]	66.13[a]
5	34.20[a]	65.80[a]	33.09[a]	66.91[a]	34.13[a]	65.87[a]
6	34.16[a]	65.84[a]	33.36[a]	66.64[a]	33.92[a]	66.08[a]
7	34.23[a]	65.77[a]	33.76[a]	66.24[a]	34.46[a]	65.54[a]
8	34.50[a]	65.50[a]	32.77[a]	67.23[a]	33.57[a]	66.43[a]
9	34.50[a]	65.50[a]	33.98[a]	66.02[a]	33.52[a]	66.48[a]
10	34.11[a]	65.89[a]	32.58[a]	67.42[a]	33.54[a]	66.46[a]

This table examines the proportion percentage of bins for which DWVWAP has lower (or higher) tracking error. Tracking error for DVWAP. It sorts the sample by contemporaneous intra-day volume variance (spikiness), historical intra-day volume variance (VolVar) and historical inter-day price variance (PriceVar). Group 1 (10) has the lowest (highest) spikiness, VolVar, or PriceVar. Tracking error for DVWAP $\left(\epsilon^d_{i,t_b}\right)$ or HVWAP $\left(\epsilon^h_{i,t_b}\right)$ is the absolute difference between the proportion of the order that DVWAP or HVWAP trades in bin b, and the proportion of total daily volume that trades in bin b. Thus, for DVWAP, tracking error is: $\epsilon^d_{i,t_b} = \left|v^d_{i,t_b} - v^a_{i,t_b}\right|$, where v^d_{i,t_b} is the percentage of total that DVWAP trades in bin b and v^a_{i,t_b} is the percent of total daily volume that trades in bin b. Similarly, tracking error for HVWAP is $\epsilon^h_{i,t_b} = \left|v^h_{i,t_b} - v^a_{i,t_b}\right|$, where v^h_{i,t_b} is the percent of total that HVWAP trades in bin b. The superscript a indicates significance at 1% in a Chi-Square test.

Table 5
Model variable definitions.

Variable	Definition
DIFF	The absolute value of the difference between DVWAP's tracking error HVWAP's tracking error $\left(\epsilon^d_{i,t_b} = \left\|v^d_{i,t_b} - v^a_{i,t_b}\right\|\right)$ and HVWAP's tracking error $\left(\epsilon^h_{i,t_b} = \left\|v^h_{i,t_b} - v^a_{i,t_b}\right\|\right)$
IMPROVE	Dummy variable that equals 1 if $\left\|v^h_{i,t_b} - v^a_{i,t_b}\right\|$ exceeds $\left\|v^d_{i,t_b} - v^a_{i,t_b}\right\|$ and equals 0 otherwise
CORRECT	Dummy variable that equals 1 if DVWAPA correctly predicts that volume will be higher or lower than the historical profile. That is, CORRECT equals 1 if either $\left(v^d_{i,t_b} > v^h_{i,t_b} \& v^a_{i,t_b} > v^h_{i,t_b}\right)$ or $\left(v^d_{i,t_b} < v^h_{i,t_b} \& v^a_{i,t_b} < v^h_{i,t_b}\right)$
VolVar	Historical intra-day variance of the volume traded
lnTradeSize	Natural log of the average trade size
lnDailyTrades	Natural log of the historical average number of trades per day for the stock
BaRatio	Historical average ratio of the number of shares at the bid or ask to the average daily volume
PriceVar	Historical inter-day variance of the stock price

DVWAP is especially able to extract useful information from past volume movements. While the results in Columns 2 and 3 indicate that BaRatio reduces DVWAP's performance, the result is weakly significant at best. And, VolVar continues to improve DVWAP's performance in Columns 2 and 3. Further, the low R^2 in all models indicates that these variables do not completely explain DVWAP's ability to track actual daily volume. Thus, while stock-level characteristics may influence DVWAP's performance, they do not completely explain it.

3.2. DVWAP and the Bialkowski et al. (2008) approach

This section examines DVWAP and the approach in Bialkowski et al. (2008). Bialkowski et al. (2008) propose one way to improve VWAP algorithms. They suggest a model that decomposes a stock's intra-day turnover into a 'common' component and a 'specific' component. The 'common' component reflects market-wide factors. The 'specific' component reflects stock-specific factors. The Bialkowski et al. (2008) model significantly improves upon a purely historical approach, such as that in HVWAP. This motivates examination of how Bialkowski et al. (2008) model intra-day value, and how DVWAP compares to the Bialkowski et al. (2008) approach.

Table 6
DVWAP performance and volatility.

	DIFF (1)	IMPROVE (2)	CORRECT (3)
lnDailyTrades	−0.001	−0.024[a]	−0.013[c]
	[0.558]	[0.001]	[0.097]
lnTradeSize	0.003[a]	−0.018[a]	−0.022[a]
	[0.001]	[0.007]	[0.002]
baRatio	−0.006[a]	0.019	0.024[c]
	[0.002]	[0.188]	[0.085]
volVar	0.001[b]	0.016[a]	0.015[a]
	[0.020]	[0.001]	[0.009]
volatility	−0.002[b]	0.004	0.009
	[0.045]	[0.645]	[0.292]
Constant	−0.026[c]	0.700[a]	0.651[a]
	[0.053]	[0.000]	[0.000]
Observations	983,640	983,640	983,640
R-squared	0.003	0.004	0.003
Test-statistic	4.12[a]	39.06[a]	24.11[a]

This table contains estimates Eqs. (10)–(12). The dependent variable in Column 1 is $\epsilon^d_{i,t_b} - \epsilon^h_{i,t_b}$. The dependent variable in Column 2 equals 1 if $\left|v^h_{i,t_b} - v^a_{i,t_b}\right|$ exceeds $\left|v^d_{i,t_b} - v^a_{i,t_b}\right|$ and equals 0 otherwise. The dependent variable in Column 3 equals 1 if DVWAP correctly predicts that volume will be above or below the historical profile. Test-statistic denotes the regression F-statistic for the OLS model (Column 1), and the Wald χ^2 statistic for the Probit models (Columns 2 and 3). All models use standard errors clustered by stock and by date. Numbers in brackets are p-values. Superscripts a, b, and c denote significance at 1%, 5% and 10%, respectively.

The DVWAP method and the Bialkowski et al. (2008) method function in different ways. The major difference is that DVWAP uses contemporaneous market data to predict the ratio $v^a_{i,t_b}/v^h_{i,t_b}$. The DVWAP model in Section 3.1.2 asserts that the ratio is a function of 'unexpected volume'. And, unexpected volume is the difference between the number of shares traded in bin $b-1$ and the number that would be predicted trade if trade followed a historical pattern. The historical pattern is stock-specific. By contrast, the Bialkowski et al. (2008) approach posits that turnover contains a stock-specific component and a common component.

This section uses a similar sample to Section 3.1. However, due to additional modeling requirements, the number of estimates reduces to 871,718. The day is split into 5 min bins. The results indicate that DVWAP has lower tracking error than the models in Bialkowski et al. (2008) do.

3.2.1. The Bialkowski et al. (2008) model

This section defines the Bialkowski et al. (2008) approach. The approach builds upon the approach in Lo and Wang (2000). For present purposes, it proceeds as follows.

Step 1, define the turnover for stock i, on day t, and in bin $b+1$ as $x_{i,t_{b+1}}$. Let this be a function of a 'common' component and a 'specific' component. Denote the predicted 'common' component as $\hat{c}_{i,t_{b+1}}$; denote the predicted 'specific' component as $\hat{e}_{i,t_{b+1}}$. Thus, define the stock's turnover as:

$$x_{i,t_{b+1}} = \hat{c}_{i,t_{b+1}} + \hat{e}_{i,t_{b+1}} \tag{13}$$

Step 2, define a suitable 'common' component. Bialkowski et al. (2008) suggest, in their Eq. (9), that a suitable 'common' component is 'historical average of the common components of intra-day volume'. Thus, the 'common' component for bin b on day t derives from the value-weighted average turnover for all stocks in the market in bin b across days $t-1$ to $t-25$. This yields the following:

$$\hat{c}_{i,t_b} = \sum_{i=1}^{25} c_{i,(t-i)_b} \quad (14)$$

Step 3, estimate the specific component. The specific component is the part of turnover that is due to stock-specific factors. A key contribution of Bialkowski et al. (2008) is to model the specific component. They propose two ways to do this. This paper replicates both ways.

The first method defines estimates $\hat{e}_{i,t_b}$ by using an ARMA (1,1) model. It defines the specific component in bin b as a function of the specific component in bin $b-1$. This yields:

$$e_{i,t_b} = \phi_1 e_{i,t_{b-1}} + \phi_2 + \epsilon_{i,t_b} \quad (15)$$

where,

$$e_{i,t_{b-1}} = x_{i,t_{b-1}} - c_{i,t_{b-1}} \quad (16)$$

The second method estimates $\hat{e}_{i,t_b}$ by using a SETAR method. This method posits that there may be a regime shift during the day. This regime shift may occur after a certain volume or turnover trades. Bialkowski et al. (2008) suggest that there may be two regimes. They denote the regime-shift-point as τ. Let $\mathbb{I}_\tau$ be an indicator variable that equals one if volume exceeds τ. Thus, the SETAR model has the form:

$$e_{i,t_b} = (\phi_{11} e_{i,t_{b-1}} + \phi_{12})(1 - \mathbb{I}_\tau) + (\phi_{21} e_{i,t_{b-1}} + \phi_{22})(\mathbb{I}_\tau) + \epsilon_{i,t_b} \quad (17)$$

This SETAR model has five parameters: ϕ_{11}, ϕ_{12}, ϕ_{21}, ϕ_{22}, and τ. This paper jointly estimates the parameters and chooses the set of parameters that yields the lowest mean squared error.

Step 4, combine the above knowledge to forecast $\hat{x}_{i,t_{b+1}}$. The issue is then how the model performs and how its performance differs from that of DVWAP.

3.2.2. How does the turnover model perform?

This section analyzes how the models perform. Three models are of interest: DVWAP, and the ARMA and SETAR models based upon Bialkowski et al. (2008).

This section analyzes a model's performance by considering the extent to which it improves over HVWAP. A model improves over HVWAP is better able to predict the percentage of total daily volume traded in bin b. Thus, a measure of a model's performance is the difference between the tracking error of the model and the tracking error of HVWAP. To formalize this, denote the proportion of actual volume traded in bin b be v^a_{i,t_b}; the historical average amount traded in bin b be v^h_{i,t_b}. The proportion that DVWAP, ARMA, and SETAR predict to trade in bin b are v^d_{i,t_b}, v^{arma}_{i,t_b}, and v^{setar}_{i,t_b}, respectively. Denote the improvements over HVWAP as δ^d, δ^{arma} and δ^{setar}, respectively. Thus, the tracking errors are:

$$\epsilon^h_{i,t_b} = \left| v^h_{i,t_b} - v^a_{i,t_b} \right|$$
$$\epsilon^d_{i,t_b} = \left| v^d_{i,t_b} - v^a_{i,t_b} \right|$$
$$\epsilon^{arma}_{i,t_b} = \left| v^{arma}_{i,t_b} - v^a_{i,t_b} \right|$$
$$\epsilon^{setar}_{i,t_b} = \left| v^{setar}_{i,t_b} - v^a_{i,t_b} \right|$$
$$\delta^d = \epsilon^d_{i,t_b} - \epsilon^h_{i,t_b}$$
$$\delta^{arma} = \epsilon^{arma}_{i,t_b} - \epsilon^h_{i,t_b}$$
$$\delta^{setar} = \epsilon^{setar}_{i,t_b} - \epsilon^h_{i,t_b}$$

The results confirm that Bialkowski et al. (2008) approaches improve execution but that DVWAP may yield superior results.

The out-of-sample results are in Table 7. The results for ARMA and SETAR are qualitatively similar to those in Table 6 Bialkowski et al. (2008). Here, the results in Table 7 Panel A indicate that DVWAP, ARMA, and SETAR models yield lower tracking errors than HVWAP does. The models yield similar results on average (the row 'All'). However, if the day is less volatile (i.e. spikiness is low), then DVWAP yields superior performance. If the day is more volatile, then ARMA yields the best results. The results in Table 7 Panel B indicate that DVWAP's performance improvement is more consistent than that of the ARMA or SETAR models. Panel B contains the variance of δ^d, δ^{arma} and δ^{setar}. It indicates that δ^d has the lowest variance. This suggests that the performance improvement in DVWAP is more consistent than the performance improvements in the ARMA or SETAR models are.

Table 7
Comparison with Bialkowski et al. (2008).

Spikiness group	ARMA (1)	SETAR (2)	DVWAP (3)	ARMA-DVWAP (4) = (1) − (3)	SETAR-DVWAP (5) = (2) − (3)
Panel A: Means					
All	−0.01371	−0.01363	−0.0138	0.00009[a]	0.00017[a]
1	−0.0137	−0.01359	−0.01654	0.00284[a]	0.00294[a]
2	−0.01372	−0.01364	−0.01332	−0.00041[a]	−0.00032[a]
3	−0.01373	−0.01363	−0.01326	−0.00046[a]	−0.00037[a]
4	−0.01374	−0.01365	−0.01327	−0.00047[a]	−0.00038[a]
5	−0.01372	−0.01363	−0.01372	−0.00001[a]	0.00009[a]
6	−0.01372	−0.01365	−0.01449	0.00077[a]	0.00084[a]
7	−0.01372	−0.01364	−0.01408	0.00036[a]	0.00044[a]
8	−0.0137	−0.01362	−0.01433	0.00064[a]	0.00071[a]
9	−0.01371	−0.01364	−0.0127	−0.00101[a]	−0.00094[a]
10	−0.01368	−0.01363	−0.01201	−0.00167[a]	−0.00162[a]
Panel B: Variances					
All	0.00564	0.00554	0.00396	0.00169[a]	0.00158[a]
1	0.0059	0.00584	0.00498	0.00092[a]	0.00086[a]
2	0.0058	0.0056	0.00392	0.00188[a]	0.00167[a]
3	0.00567	0.00547	0.00382	0.00185[a]	0.00165[a]
4	0.00552	0.0054	0.00395	0.00157[a]	0.00145[a]
5	0.00539	0.00524	0.00391	0.00148[a]	0.00133[a]
6	0.00534	0.00527	0.00419	0.00115[a]	0.00108[a]
7	0.00524	0.00519	0.00389	0.00135[a]	0.0013[a]
8	0.00536	0.0053	0.00393	0.00142[a]	0.00136[a]
9	0.00577	0.00573	0.00366	0.00211[a]	0.00208[a]
10	0.00631	0.00626	0.00317	0.00314[a]	0.00309[a]

This table examines the performance of DVWAP, and the performance of the SETAR and ARMA models. For Panel A, The statistics in Columns 1, 2, and 3, are the improvement in tracking error for the ARMA, SETAR, and DVWAP models (δ^d, δ^{arma}, and δ^{setar}, respectively). The statistics in Column 4 and Column 5 are the difference between the performance improvement of the ARMA and SETAR models versus DVWAP. For Panel B, Columns 1, 2, and 3 contain the variance of performance improvements (that is the variance of δ^d, δ^{arma}, and δ^{setar}); Columns 4 and 5 contain the difference between the variances. The superscript *a* denotes significance at 1% significance in difference in mean tests (in Panel A) and in *F*-tests for difference in variance (in Panel B).

The in-sample results are in Table 8. The in-sample results are consistent with the out-of-sample results. In particular, (1) DVWAP has a lower tracking error in some spikiness-sorts, but not in others, and (2) DVWAP's tracking error has a lower variance than do the ARMA and SETAR models.[3]

Overall the results indicate the following. First, DVWAP can improve on the ARMA and SETAR approaches in Bialkowski et al. (2008). This performance improvement is more pronounced if the day is less volatile. And, this performance improvement does not necessarily hold across all market conditions. However, second, DVWAP's performance improvement is more stable and consistent than that of the ARMA or the SETAR models. This indicates that DVWAP can improve over ARMA and SETAR in some market conditions and can yield more consistent improvements over HVWAP than the ARMA or SETAR models do.

[3] It is important to note that Columns 1–3 of Table 8 and of Table 7 focus on the differences between the classic historical approach and the DVWAP, ARMA, and SETAR approaches. The results for the ARMA and SETAR models are similar in magnitude to those reported in Bialkowski et al. (2008).

Table 8
In sample comparison with Bialkowski et al. (2008).

Spikiness group	ARMA (1)	SETAR (2)	DVWAP (3)	ARMA-DVWAP (4) = (1) − (3)	SETAR-DVWAP (5) = (2) − (3)
Panel A: Means					
2	−0.01357	−0.01358	−0.01384	0.00028[a]	0.00026[a]
3	−0.01356	−0.01358	−0.01376	0.00020[a]	0.00019[a]
4	−0.01355	−0.01357	−0.01224	−0.00132[a]	−0.00133[a]
5	−0.01355	−0.01356	−0.01070	−0.00285[a]	−0.00286[a]
6	−0.01356	−0.01352	−0.01016	−0.00340[a]	−0.00337[a]
7	−0.01353	−0.01347	−0.01257	−0.00095[a]	−0.00090[a]
8	−0.01348	−0.01340	−0.01867	0.00519[a]	0.00527[a]
9	−0.01349	−0.01337	−0.01854	0.00505[a]	0.00517[a]
10	−0.01334	−0.01319	−0.01063	−0.00271[a]	−0.00256[a]
Panel B: Variances					
1	0.00502	0.00523	0.00558	−0.00056[a]	−0.00035[a]
2	0.00550	0.00568	0.00466	0.00084[a]	0.00102[a]
3	0.00596	0.00608	0.00445	0.00152[a]	0.00163[a]
4	0.00645	0.00653	0.00381	0.00264[a]	0.00272[a]
5	0.00664	0.00656	0.00302	0.00362[a]	0.00355[a]
6	0.00659	0.00644	0.00288	0.00371[a]	0.00356[a]
7	0.00610	0.00592	0.00323	0.00287[a]	0.00269[a]
8	0.00519	0.00492	0.00512	0.00007[a]	−0.00020[a]
9	0.00477	0.00476	0.00649	−0.00172[a]	−0.00174[a]
10	0.00468	0.00453	0.00350	0.00118[a]	0.00103[a]

This table examines the insample performance of DVWAP, and the performance of the SETAR and ARMA models. For Panel A, The statistics in Columns 1, 2, and 3, are the improvement in tracking error for the ARMA, SETAR, and DVWAP models (δ^{d}, δ^{arma}, and δ^{setar}, respectively). The statistics in Column 4 and Column 5 are the difference between the performance improvement of the ARMA and SETAR models versus DVWAP. For Panel B, Columns 1, 2, and 3 contain the variance of performance improvements (that is the variance of δ^{d}, δ^{arma}, and δ^{setar}); Columns 4 and 5 contain the difference between the variances. The superscript *a* denotes significance at 1% significance in difference in mean tests (in Panel A) and in *F*-tests for difference in variance (in Panel B).

3.3. Robustness and sensitivity

This section ensures robustness that the model is robust to model specification. First, the paper re-stimates Eq. (7) using different numbers of lags. However, the results are qualitatively similar. Second, the results are robust to bin length. Eq. (7) bases itself on 5-min bins. Results are similar for 15 min and 30 min and 60 min bins.

Third, the paper ensures that the results hold after removing outlying bins that have unusually high volumes, which ordinarily correspond with significant corporate announcements. DVWAP's out-performance increases after removing outlying bins. Fourth, Eq. (7) uses three lagged differences. The paper re-estimates the model with four and five lags. The results are not sensitive to the number of independent variables. Fifth, the paper estimates Eq. (7) using GMM to control for endogeneity that would arise if anticipation of a volume spike influences lagged volume. The results hold in the GMM estimation.

Sixth, the paper ensures that the results are robust to various values of $\bar{\theta}$, $\hat{\theta}$, and γ. To do this, it iterates through various combinations of the parameters. It examines values of $\underline{\theta}$ from 0.99 to 0.80, and $\bar{\theta}$ from 1.01 to 1.20. It also examines values of γ from 0.75 to 0.95. In all models, DVWAP improves on HVWAP. If the γ value is lower, then DVWAP reverts to HVWAP sooner; and thus, (a) the tracking increases, but (b) the standard deviation of the tracking error decreases. If $\underline{\theta}$ is lower and/or $\bar{\theta}$ is higher, then (a) DVWAP's mean improvement over HVWAP increases; that is, there is lower average difference from the actual volume profile, but (b) the variance of the tracking error increases. In all cases, DVWAP has lower tracking error than HVWAP does.

The multivariate results in Section 3.1.4 are robust to model specification. This paper re-estimates the models using a panel regression (with fixed stock-date effects), and finds similar results. Further, it replaces VOLVAR with the inter-day price-volatility (as a proxy for volatility) and replaces BaRatio with spread (as a proxy for liquidity) and finds similar results.

The results are robust to the choice of an in-sample test. Tables 3 and 8 report in-sample results. All other results (unreported for brevity) are qualitatively the same in in-sample tests. Unsurprisingly, DVWAP tends to perform better in in-sample tests than it does in out-of-sample tests.

4. Conclusion

This paper examines the trader's optimal trading decision if there is random news (or noise). Random intra-day information should influence a trader's trading behaviour (Foster and Viswanathan, 1993, 1996). Thus, this paper proposes a DVWAP framework that enables informed traders to incorporate stock-specific information into a trading framework. Empirical results indicate that this can yield superior execution to a purely historical approach, or to the appraoches in Bialkowski et al. (2008), when estimated using a strong empirical model for intra-day volume.

This paper makes a key contribution to the algorithmic trading literature by proposing a mechanism to incorporate intra-day noise into an algorithmic trading rule. This has become especially important due to the growing popularity of algorithmic trading.

Acknowledgements

This article benefited from comments at the 22nd Australasian Finance and Banking Conference. The author also thanks Ekkehart Boehmer, Zhi Da, David Feldman, Kingsley Fong, Terrence Hendershott, Ron Kaniel, Ronan Powell, Mark Seasholes, Jian-Xin Wang, and Bohui Zhang for their comments.

References

Back, K., Cao, C.H., Willard, G.A., 2000. Imperfect competition among informed traders. Journal of Finance 55, 2117–2155.
Barclay, M.J., Warner, J.B., 1993. Stealth trading and volatility: which trades move prices? Journal of Financial Economics 34, 281–305.
Barron, O.E., Karpoff, J.M., 2004. Information precision, transaction costs, and trading next term volume. Journal of Banking and Finance 28, 1207–1223.
Bauer, R., Cosemans, M., Eichholz, P., 2009. Option trading and individual investor performance. Journal of Banking and Finance 33, 731–746.
Bialkowski, J., Darolles, S., Le Fol, G., 2008. Improving VWAP strategies: a dynamic volume approach. Journal of Banking and Finance 32, 1709–1722.
Brooks, R.M., Patel, A., Su, T., 2003. How the equity market responds to unanticipated events. Journal of Business 76, 109–133.
Cao, C., Feild, L.C., Hanka, G., 2004. Does insider trading impair market liquidity? Evidence from ipo lockup expirations. Journal of Financial and Quantitative Analysis 39, 25–46.
Chan, K., Fong, W.M., 2000. Trade size, order imbalance, and the volatility volume relation. Journal of Financial Economics 57, 247–273.
Chang, C.C., Hseih, P.F., Lai, H.N., 2009. Do informed option investors predict stock returns? Evidence from the Taiwan stock exchange. Journal of Banking and Finance 33, 757–764.
Chew, L.M., 2007. Algorithmic trading: influx of trading technologies in Asia. Research Report, Financial Insights.
Comerton-Forde, C., 1999. Do trading rules impact on market efficiency? A comparison of opening procedures on the Australian and Jakarta stock exchanges. Pacific-Basin Finance Journal 7, 495–521.
Cooney, J.W., Sias, R.W., 2004. Informed trading and order type. Journal of Banking and Finance 28, 1711–1743.
Darrat, A.F., Zhong, M., Cheng, L.T.W., 2007. Intraday volume and volatilitynext term relations with and without public news. Journal of Banking and Finance 31, 2711–2729.
Fong, K.Y.L., Liu, W.M., 2010. Limit order revisions. Journal of Banking and Finance 34, 1873–1885.
Foster, F.D., Viswanathan, S., 1993. The effect of public information and competition on trading volume and price volatility. Review of Financial Studies 6, 23–56.
Foster, F.D., Viswanathan, S., 1996. Strategic trading when agents forecast the forecasts of others. Journal of Finance 51, 1437–1478.
Grossman, R.L., 2005. Marching up the learning curve: the first buy-side algorithmic trading survey. Research Report, Financial Insights.
Inci, A.C., Lu, B., Seyhun, H.N., 2010. Intraday behavior of stock prices and trades around insider trading. Financial Management 39, 323–363.

Johnson, T.C., 2008. Volume, liquidity, and liquidity risk. Journal of Financial Economics 87, 388–417.
Kakade, S.M., Kearns, M., Mansour, Y., Ortiz, L.E., 2004. Competitive algorithms for VWAP and limit order trading. In: Proceedings of the ACM Electronic Commerce Conference, pp. 189–198.
Kang, M., 2010. Probability of information-based trading and the January effect. Journal of Banking and Finance 34, 2985–2994.
Kearns, M., 2008. An algorithmic approach. Research Report, Banc of America.
Kissell, R., Malamut, R., 2006. Algorithmic decision-making framework. Journal of Trading 1, 12–21.
Kissell, R., Glantz, M., Malamut, R., 2004. A practical framework for estimating transaction costs and developing optimal trading strategies to achieve best execution. Finance Research Letters 1, 35–46.
Lo, A.W., Wang, J., 2000. Trading volume: definitions, data analysis, and implications of portfolio theory. Review of Financial Studies 13, 257–300.
Madhavan, A., 2002. VWAP strategies. Trading 2002, 32–38.
Muck, M., 2010. Trading strategies with partial access to the derivatives market. Journal of Banking and Finance 34, 1288–1298.
Nofsinger, J.R., Prucyk, B., 2003. Option volume and volatility response to scheduled economic news releases. Journal of Futures Markets 23, 315–345.
Petersen, M.A., 2009. Estimating standard errors in finance panel data sets: comparing approaches. Review of Financial Studies 22, 435–480.
Thapa, C., Poshakwale, S.S., 2010. International equity portfolio allocations and transaction costs. Journal of Banking and Finance 34, 2627–2638.
Xu, X.E., Chen, P., Wu, C., 2006. Time and dynamic volume-volatility next term relation. Journal of Banking and Finance 30, 1535–1558.