Optimal VWAP Tracking

Jędrzej Białkowski[1], Daniel Mitchell[2], and Stathis Tompaidis[3]

[1]Department of Economics and Finance, University of Canterbury,

Christchurch, New Zealand

jedrzej.bialkowski@canterbury.ac.nz

[2]McCombs School of Business, University of Texas at Austin,

Austin, TX, USA

dan.mitchell@utexas.edu

[3]McCombs School of Business, University of Texas at Austin,

Austin, TX, USA

stathis.tompaidis@mccombs.utexas.edu

We thank Glenn Boyle, and seminar participants at Northwestern University, the University of Chicago, the University of Texas at Austin, and the INFORMS 2012 meetings in Phoenix for helpful comments and suggestions.

Optimal VWAP Tracking

Abstract

We consider the problem of finding a strategy that tracks the volume weighted average price (VWAP) of a stock, a key measure of execution quality for large orders used by institutional investors. We obtain the optimal, dynamic, VWAP tracking strategy in closed form in a model without trading costs with general price and volume dynamics. We build a model of intraday volume using the Trade and Quote dataset to empirically test the strategy, both without trading costs and when trading has temporary and permanent effects, and find that the implementation cost is lower than the cost charged by brokerage houses.

JEL classification: G12, G29, C61

Keywords: Volume Weighted Average Price, Algorithmic Trading, Trading Volume, Trading Costs, Dynamic Programming

1 Introduction

Algorithmic trading is an area of increasing importance in financial exchanges. Trading algorithms range from high-frequency algorithms that submit and cancel orders thousands of times per second, to algorithms that trade less frequently, whose objective varies from minimizing trading costs to tracking the performance of a benchmark. One of the most common benchmarks for judging the execution quality of individual stocks, used by institutional investors such as pension funds and mutual funds, is the daily volume-weighted average price (VWAP), originally suggested by Berkowitz, Logue, and Noser Jr. (1988). Madhavan (2002) discusses various types of VWAP strategies, their advantages and disadvantages, and challenges encountered when implementing a strategy to track VWAP. In the United States, survey data in a report by Bank of America (2007) indicate that VWAP orders represent close to 50% of all trading by institutional investors, while Almgren, Thum, Hauptmann, and Li (2005) report that VWAP orders for individual stocks represent 16% of all trades. Domowitz and Yegerman (2011) report that when trying to place an order that ranges from 0.5% to 5% of daily market volume between 23% and 27% of all algorithmic trades follow VWAP strategies, while Australian Securities Exchange (2010) estimate that 32.3% of buy-side algorithmic trades are VWAP orders. End-of-day VWAP for individual stocks is reported by exchanges, and guaranteed VWAP orders are available through most brokerage houses. For example, Interactive Brokers offer a guarranteed VWAP order for a cost of 10 basis points, while Białkowski, Darolles, and Le Fol (2008) report that in Europe the cost of VWAP orders ranges between 10 and 20 basis points. Beyond its popularity with institutional investors, VWAP has been used by regulators for assessing taxes in cases of issuance of shares to existing shareholders.[1] While alternatives such as implementation shortfall, exist,

[1] An Australian company, Woolworth's Group, established a real estate investment trust, owning 69 Woolworths Group anchored shopping centres and freestanding stores, by distributing newly issued shares to its shareholders on November 26th, 2012. The tax basis of these shares was set to the daily VWAP of the Woolworth's stock on the day of issuance by the Australian

VWAP remains a widely used measure of trade efficiency.[2]

The relatively small literature on VWAP tracking algorithms has developed along two directions. The first direction derives static, predetermined, optimal execution strategies before trading opens for the day, and the strategies are not updated during the day. The second direction employs trading strategies that try to match a dynamic volume profile without regard to changes to the price of the stock. In this paper we extend the literature by presenting a closed-form, dynamic, solution to the problem of optimally tracking VWAP. The closed-form solution is obtained in a model that allows co-dependence between stock price and volume dynamics under three assumptions: market orders are used; personal orders do not impact the market price; and the intraday stock price is a martingale conditional on the realized stock volume.

Our first result is that, under the three assumptions of our model, the difference between the VWAP achieved by an arbitrary, personal, trading strategy; i.e, the VWAP calculated based on personal trades and execution prices, and the market VWAP; i.e. the VWAP calculated based on market volume and prices, does not depend on the choice of strategy. This result implies that no VWAP tracking algorithm can, on average, outperform or underperform the market VWAP. Since the first moment of the difference between market and personal VWAP does not depend on the trading strategy chosen, consistent with an objective of mean-variance optimization, we instead try to minimize the variance of the difference of the market VWAP and the VWAP of a personal strategy.

Our closed-form solution for the optimal trading strategy illustrates that the strategy

taxing authorities.

[2]In May 2012, NASDAQ OMX sought the approval of the Securities and Exchange Commission (SEC) to offer "Benchmark Orders" that track individual stock VWAP on the NASDAQ exchange, see NASDAQ (2012b). While the request was disapproved by the SEC in January 2013, it indicates the increasing interest in VWAP orders. We note that the issues raised by an exchange offering a VWAP contract are beyond how to determine a strategy that tracks VWAP; e.g., the issue of unfair competition between the exchange and the brokerage houses was one of the objections raised, see the comments at NASDAQ (2012a), and the SEC decision at SEC (2013). These issues are outside the scope of our paper.

is independent of the stock's running VWAP, the stock's price, and the running VWAP of the trading strategy, but depends on the correlation between the stock returns and the trading volume; the expected volume; and, the variance of returns in different time intervals. Two of the results derived from our closed-form solution are in contrast to the literature and are of interest to market participants trying to implement a VWAP tracking strategy. First, we find that it is occasionally optimal to trade against the direction of the desired overall trade when realized market volume is much lower than anticipated; i.e., if personal cumulative traded amount turns out to be too high relative to the realized market volume it may be optimal to trade in the opposite direction to restore balance with the realized market volume. Our second result, that is in contrast to the existing literature, is that the optimal trading strategy may deviate from the expected relative volume of a time interval. The deviation is most pronounced when the correlation between stock return variance and trading volume is positive. The intuition behind this result lies in the disproportionate impact that time intervals with large volumes have on the calculation of the VWAP. When stock return variance covaries positively with trading volume, a larger-than-expected trading volume has a larger impact on the market VWAP than a smaller-than-expected volume. The difference is largest for intervals with high expected volumes. Both results have direct implications on the implementation of VWAP tracking algorithms by institutional investors.

An additional contribution of our paper is that we propose a cross-sectional model of trading volume and stock returns, which we calibrate to intraday trading data for the 30 stocks that form the Dow Jones Industrial Average. Common features identified in the literature include a "U-shaped" pattern of volume for US stocks, with higher volume during the morning and afternoon and lower volume during the middle of the day; a positive correlation between volume and variance of stock returns; and volume autocorrelation. We incorporate all these features in a model of trading volume and stock returns. Dividing the day into intervals, we allow for an arbitrary shape in average intraday volume, and capture autocorrelation using an autoregressive model of volume. The correlation between volume and stock return variance is captured by assuming that the logarithm of the stock return variance varies linearly with the logarithm of the realized trading volume. We calibrate the model using 60-day rolling windows and find that the performance of the theoretically optimal VWAP strategy improves when the variance of daily logarithmic returns is lower.[3] Performance also improves, in the absence of trading costs, as the number of shares traded increases relative to the market volume for the stock — a result due to the increased influence of personal trades to the market VWAP as the number of shares traded increases. Finally, we find that the theoretically optimal strategy exhibits larger variability early in the day. The performance of the VWAP achieved by the theoretical, personal, trading strategy can be measured by the standard deviation between the personal and the market VWAPs: in our empirical study for the 30 stocks in the Dow Jones Industrial Average this standard deviation ranges between 6 and 27 basis points.

While we do not derive the optimal VWAP tracking strategy in a model where personal trades do impact market prices, we extend our empirical study of the strategy's performance to account for trading costs. We calibrate a model for trading costs that is commonly used in the literature and in practice, which separates trading costs into market impact, meaning the change to the market price of a stock due to personal trades, and execution slippage, which captures the difference between the market price of a stock and the execution price of the personal trade.[4] Our study shows that the standard deviation of the difference between personal and market VWAPs is similar to the case without trading costs, while the cost of executing the trading strategy, expressed as the difference between personal and market VWAPs, increases with the variance of the returns of the stock and the size of the personal

[3] In line with our results, Domowitz and Yegerman (2011) find that the execution cost of a VWAP order increases in periods characterized by high market volatility.

[4] Market impact is permanent and affects the market price of the stock following the trade, while execution slippage is temporary and only affects the price received for a particular trade.

volume. Demonstrating the usefulness of the theoretically optimal VWAP tracking strategy, the average difference between personal and market VWAPs remains in all cases below the typical 10 basis points charged by brokerage houses. To account for both return and variance, we consider the annualized Sharpe ratio achieved by the trading strategy we propose, when the implementation of a VWAP trade costs 10 basis points. We find that, when we trade up to 5% of expected daily market volume, the annualized Sharpe ratios range from 2.4 — for volatile stocks — to over 22, for the 30 stocks in the Dow Jones Industrial Average.

1.1 Literature review

The quality of equity trade execution is economically important. The empirical study by Keim and Madhavan (1997) shows that transaction costs are systematically related to complexity of trade and market liquidity. Moreover, the paper provides evidence for variation in trading cost across financial institutions. The performance of institutional trades in the bearish and bullish phase of the markets was subject of analysis by Chiyachantana, Jain, Jiang, and Wood (2004). The study shows that in bullish markets, institutional purchases have a bigger price impact than sells; however, in the bearish markets, sells have a higher price impact. The price impact varies depending on order characteristics and firm-specific factors. The very recent study by Anand, Irvine, Puckett, and Venkataraman (2012) scrutinizes propriety dataset of institutional investors' equity transactions. An empirical analysis allows authors to conclude that trade implementation process contributes to relative portfolio performance.

Literature on optimal VWAP tracking strategies is relatively recent. Konishi (2002) derives a static optimal execution strategy for tracking VWAP, where the trading strategy is determined at the beginning of the day and no revision is allowed in the middle of the day. McCulloch and Kazakov (2012) extends the model of Konishi (2002) to the case where the stock price is not a martingale and show that the mean-variance optimal VWAP trading strategy is the sum of a minimum variance strategy and a directional pricing strategy, but do not derive these strategies. Fuh, Teng, and Wang (2010) investigate the performance of heuristic trading strategies to track VWAP both in a simulation study and in an empirical study for 20 stocks in the Taiwan Stock Exchange. Białkowski et al. (2008) and Humphery-Jenner (2011) suggest using an adaptive, fixed-schedule, trading strategy, with trades adjusted to match the expected market volume over the remaining time intervals in two models of volume that allow for auto-regression and jumps, but do not allow for correlation between price return variance and volume and do not optimize the trading strategy. Kakade, Kearns, Mansour, and Ortiz (2004) provide results on the existence of dynamic algorithms for tracking VWAP and discuss the performance of static algorithms relative to dynamic algorithms, but do not offer a way to construct such algorithms. Frei and Westray (2012) consider a continuous time model and offer an algorithm to construct the optimal trading strategy that minimizes a weighted sum of the difference between personal and market VWAP and the variance of the difference, in a model where interval volume as a percentage of daily volume is assumed to be known and where the trader faces execution slippage costs. Compared to the rest of the literature, we employ a model that allows general dynamics for the trading volume and the price of the stock, and derive a closed-form solution for a dynamic trading strategy that minimizes the standard deviation between the personal and market VWAPs.

There is a large literature on modeling market volume and its link with the dynamics of stock price. Papers in the literature include Clark (1973), Epps and Epps (1976), Tauchen and Pitts (1983), Karpoff and Boyd (1987), Easley and O'Hara (1987), Admati and Pfleiderer (1988), Jain and Joh (1988), Foster and Viswanathan (1990), Gallant, Rossi, and Tauchen (1992), Chan, Christie, and Schultz (1995), Kaastra and Boyd (1995), Andersen (1996), Gouriéroux, Jasiak, and Fol (1999), Chan and Fong (2000), Lo and Wang (2000), Manganelli (2005), Darat, Rahman, and Zhong (2003), Giot, Laurent, and Petitjean (2010), Manchaldore, Palit, and Soloviev (2010), and Brownlees, Cipollini, and Gallo (2011). Our model of volume incorporates the common features identified in the literature and is calibrated based on the Trade and Quote dataset over 18 months of trading for the 30 stocks of the Dow Jones Industrial Average.

Estimating market impact from trades empirically is discussed in the papers by Bouchaud, Gefen, Potters, and Wyart (2004), Almgren et al. (2005), Engle, Furstenberg, and Russell (2008), Obizhaeva and Wang (2006), Gatheral, Schied, and Slynko (2012), and reviewed in Gatheral and Schied (2013). A theoretical foundation for the permanent impact from a trade is provided in the paper by Kyle (1985), who derives a linear equilibrium from fundamental principles. Consistent with the theoretical model of Kyle (1985), our model of permanent impact is linear to the number of shares traded. Our model for temporary impact follows the model in Almgren et al. (2005) and assumes a nonlinear impact function.

2 Model

To develop a trading strategy that tracks VWAP in a realistic setup, we build a model of stock returns and volume dynamics that is flexible enough to accommodate the observed regularities in trading volume and stock returns. The one assumption that underlies the model is that the intraday stock price is a martingale, conditional on the realized stock volume.

We formulate the model in discrete time: the trader wants to trade a fixed number of shares of stock, which we normalize to one, over the course of a day and is evaluated at the end of the day based on how well the VWAP of the trading strategy approximates the market VWAP. Trading occurs at equally spaced discrete times during the day, called time buckets. At time bucket t the trader makes a decision about what fraction of his position to trade based on his information up to, and including, time bucket t. At time bucket $t+1$ his order is executed at the market price at time bucket $t+1$. The timeline of our model is illustrated in Fig. 1.[5]

Uncertainty is described by $(\Omega, \mathcal{F}, \mathbb{P})$, where Ω is the set of possible outcomes, $\mathcal{F}$ is a σ-algebra describing the set of all possible events and $\mathbb{P}$ describes the probability measure of each set in $\mathcal{F}$. Time is indexed by $t = -1, 0, 1, ..., T$, and the state of information available to the trader is described by a filtration, $\mathcal{F} = (\mathcal{F}_{-1}, \mathcal{F}_0, \mathcal{F}_1, ..., \mathcal{F}_T)$, which is a set of increasing σ-algebras. Time $t = -1$, corresponds to the end of trading during the previous trading day, while $t = 0$ is a time after the end of trading on the previous day and before the beginning of trading, and $t = 1$ is the first trading bucket of the day.

At time bucket t the state variables are denoted by X_t. Based on this information the trader makes a decision to trade a fraction, $c(t, X_t)$, of his position, a trade which is executed at bucket $t+1, t = 0, \ldots, T-1$; i.e., the personal amount to be traded at time $t = 1$ is decided at time $t = 0$, the amount to be traded at time $t = 2$ is decided at time $t = 1$, etc. In order to enforce non-anticipativity we require that $c(t, X_t)$ is measurable with respect to $\mathcal{F}_t$. At the terminal time bucket, T, all uncertainty is resolved.

The relevant information to the trader is $X_t = \left(\kappa_t, \{\nu_i\}_{i=1}^t, V_t, S_t, M_t, P_t\right)$, where κ_t is the trader's cumulative traded volume up to and including time bucket t; ν_t is the volume traded in the market at time bucket t; V_t is the cumulative volume traded in the market up to and including time bucket t; S_t is the price of the stock at time bucket t; M_t is the running VWAP of the market; and P_t is the trader's running VWAP. Given the values of all these variables the trader makes a decision to trade a fraction $c(t, X_t)$ shares, which we abbreviate by c_t. Before trading begins, we initialize the values of cumulative personal volume and cumulative market volume to zero; $\kappa_0 = V_0 = 0$.

[5] An alternative framework would be to consider trading at time intervals that are random, for example determined by the realized volume during the day. While there may be benefits to such an approach, it is unclear how to guarantee executing a certain amount of shares during a fixed time interval like one day.

2.1 Volume dynamics

The volume at time bucket t, exclusive of personal trades, ν_t, is given by an auto-regressive process

$$\nu_{t+1} = f(t+1, \{X_i\}_{i=1}^{t}, V_{-1}, R_{t+1}), \quad t = 0, \ldots, B-1 \tag{2.1}$$

where f is a general function that depends on the previous periods' information, $\{X_i\}_{i=1}^{t}$, the cumulative volume of the previous day, V_{-1}, and an $\mathcal{F}_{t+1}$ measurable random variable, R_{t+1}. The function f is designed to reflect the "U" shape of volume that is observed during the day in the United States.[6] Fig. 2 illustrates both the U-shaped average daily volume and the deviations that may occur from the average on a single day for McDonalds stock.

The cumulative volume, up to and including bucket t, and personal trades, is given by V_t

$$\begin{aligned} V_0 &= 0 \\ V_t &= \sum_{i=1}^{t}(\nu_i + c_{i-1}) \\ V_{t+1} &= V_t + \nu_{t+1} + c_t, \end{aligned} \tag{2.2}$$

where c_t is the volume of personal trades decided at time bucket t and executed during time bucket $t+1$.

The cumulative personal volume of the trader is simply the sum of all previous trades, so that

$$\begin{aligned} \kappa_0 &= 0 \\ \kappa_t &= \sum_{i=0}^{t-1} c_i \\ \kappa_{t+1} &= \kappa_t + c_t. \end{aligned} \tag{2.3}$$

[6]The shape of the traded volume is different, for example, in European stock markets, where trading during the middle of the day may be higher than trading in the morning, due to the opening of the United States exchanges.

2.2 Stock return dynamics

We assume that the stock price is a martingale, conditional on realized volume, over the time the VWAP is calculated. Without loss of generality, we express the dynamics of the stock price in a geometric form. We assume that the bucket return variance is a function of the realized bucket trading volume, while the mean is chosen to guarantee that the stock price is a martingale,

$$\begin{aligned} S_{t+1} &= S_t \cdot \varepsilon_{t+1}, \\ E\left[\varepsilon_{t+1} \middle| \nu_{t+1}\right] &= 1, \\ E\left[\varepsilon_{t+1}^2 \middle| \nu_{t+1}\right] &= \exp\left(\sigma^2(\nu_{t+1})\right). \end{aligned}$$

The function $\sigma^2(\nu_{t+1})$ is a general, positive, function of volume and will be abbreviated as σ_{t+1}^2. In Section 4 we will adopt a functional form that implies that the logarithm of the stock return variance varies linearly with the logarithm of the realized trading volume.

2.3 VWAP definition

The running market VWAP, M_t, is defined as the market-volume-weighted average price, including personal trades,

$$M_t = \sum_{i=1}^{t} \frac{\nu_i + c_{i-1}}{V_t} S_i, \quad t = 1, \ldots, T. \tag{2.4}$$

The running personal VWAP, P_t, is defined as the personal, volume-weighted, average price

$$P_t = \sum_{i=1}^{t} \frac{c_{i-1}}{\kappa_t} S_i, \quad t = 1, \ldots, T,$$

where

$$\kappa_t = \sum_{i=1}^{t} c_{i-1}, \quad t = 1, \ldots, T-1,$$
$$\kappa_T = 1,$$

and where c_{i-1} is the personal volume traded in time bucket i, and κ_i is the cumulative personal volume. Volume is normalized so that the total personal volume traded during the day equals 1.

The market and personal running VWAP evolution equations are given by

$$M_{t+1} = \frac{M_t V_t + S_{t+1}(\nu_{t+1} + c_t)}{V_{t+1}},$$
$$P_{t+1} = \frac{P_t \kappa_t + S_{t+1} c_t}{\kappa_{t+1}}.$$

The terminal values of the market and personal VWAP are then random variables given by

$$\begin{aligned} M_T &= M_t \frac{V_t}{V_T} + S_t \sum_{i=1}^{T-t} \frac{(\nu_{t+i} + c_{t+i-1})}{V_T} \prod_{j=1}^{i} \varepsilon_{t+j}, \\ P_T &= P_t \kappa_t + S_t \sum_{i=1}^{T-t} c_{t+i-1} \prod_{j=1}^{i} \varepsilon_{t+j}. \end{aligned} \tag{2.5}$$

3 Optimization Problem and Optimal Trading Strategy

The goal of the trader is to match his personal VWAP at the end of the day to the market VWAP. We first show that, when personal trades do not impact the market price and there are no trading costs, the expected difference between personal and market VWAPs, $E[P_T -$

$M_T]$, is independent of the trading strategy c.[7]

Lemma 3.1. *If the stock price process, conditional on realized trading volume, is a martingale, and the amount of shares that has to be traded by the end of the day is known in advance, then, the expected difference between personal and market VWAPs, $E[P_T - M_T]$, is independent of the trading strategy c.*

The proof is provided in the Appendix. Lemma 3.1 holds under minimal assumptions on the price and volume evolution and implies that, under the assumptions stated, no trading strategy can either outperform or under-perform the market VWAP.

3.1 Objective

We choose an objective consistent with mean-variance optimization. Specifically, given that the expected difference between personal and market VWAPs is independent of the trading strategy, we seek to minimize the expectation of the second moment of the difference of the personal and market VWAPs at the terminal time bucket T, $E[(P_T - M_T)^2]$. We define the set of all measurable functions, with respect to $\mathcal{F}_t$, as the action set A_t. We define the trader's optimization problem at time bucket t in terms of a value function G_t where

$$G_T = (P_T - M_T)^2 ,$$
$$G_t = \min_{c_t \in A_t} E\left[G_{t+1} \middle| \mathcal{F}_t\right] , \quad t < T.$$

To enforce the constraint that the trader trades exactly one share over the course of the day we set

$$c_{T-1} = 1 - \kappa_{T-1}. \tag{3.1}$$

[7] The assumption that personal trades do not influence the market price is consistent with much of the literature on VWAP tracking, e.g., Konishi (2002), Kakade et al. (2004), Białkowski et al. (2008), Fuh et al. (2010), Humphery-Jenner (2011), and McCulloch and Kazakov (2012).

3.2 Optimal Trading Strategy

Theorem 3.1 describes, in closed-form, the strategy that minimizes the expected variance of the difference between the personal and market VWAPs at bucket T.

Theorem 3.1. *Under the stock price dynamics described by Eq. 2.4, the optimal trading strategy for the trader at time bucket t is given by*

$$c_t = 1 - \kappa_t + E\left[\frac{\frac{1-V_T}{V_T}\left(e^{\sigma^2_{t+1}+\sigma^2_{t+2}} - e^{\sigma^2_{t+1}}\right)}{E\left[\left(\frac{1-V_T}{V_T}\right)^2 \left(e^{\sigma^2_{t+1}+\sigma^2_{t+2}} - e^{\sigma^2_{t+1}}\right)\middle| \mathcal{F}_t\right]} \sum_{i=2}^{T-t} \frac{\nu_{t+i}}{V_T} \middle| \mathcal{F}_t \right]. \tag{3.2}$$

The proof, based on dynamic programming and induction, is provided in the Appendix. Eq. 3.2 is complicated due to two factors: the influence of the personal trades on the determination of the market VWAP, giving rise to the expression $(1-V_T)/V_T$; and the correlation between the variance of stock returns and the realized bucket trading volume, giving rise to the expression $\left(e^{\sigma^2_{t+1}+\sigma^2_{t+2}} - e^{\sigma^2_{t+1}}\right)$. The intuition behind the optimal strategy is significantly simpler in the case when there is no dependence between the variance of stock returns and trading volume, and when the personal volume is an insignificant percentage of market volume, $V_T \gg 1$. In that case the optimal trading strategy simplifies, and is described in the following Lemma:

Lemma 3.2. *When the variance of stock returns is independent of bucket trading volume, under the stock price dynamics described by Eq. 2.4, and in the limit when the personal trading volume is an insignificant percentage of market trading volume, the optimal trading strategy is given by*

$$c_t = E\left[\frac{V_{t+1}}{V_T}\middle| \mathcal{F}_t\right] - \kappa_t. \tag{3.3}$$

The proof of this lemma, following from Theorem 3.1, is provided in the Appendix. The optimal trade, c_t, is such that the cumulative personal volume traded at bucket $t+1$, equals the expected cumulative market volume, relative to the total volume in the day, $\kappa_{t+1} = \kappa_t + c_t = E[V_{t+1}/V_t]$.

An advantage of obtaining the optimal trading strategy in closed-form is that we can easily derive properties of the optimal trading strategy. From Eq. 3.3, we note that it is possible to optimally trade against the overall direction of trade. For example, if at time bucket t, we have already traded 60% of the total amount we have committed to trade by the end of the day, and the expected cumulative market volume by the end of the next time bucket is smaller than 60% of the daily volume; e.g., 57%, then we would optimally trade in the opposite direction during time bucket $t+1$, to reduce our personal traded amount; e.g., trade 3% in the opposite direction to reduce our position to 57% of the total amount we have committed to trade.

When the bucket volume and the stock return variance are correlated it is possible for the optimal strategy to deviate from the expected bucket volume. While it is difficult to derive a general result from Eq. 3.2, we have found that positive correlation between bucket volume and stock return variance leads to a trading strategy that overweights periods with high expected trading volume and underweights periods with low expected trading volume. The intuition behind this result is striking: since positive correlation implies that expected trading volume is, on average, positively linked with stock return variance, larger than expected trading volume in buckets with high expected trading volume is likely to have a disproportionately higher effect on the determination of VWAP than lower than expected trading volume. To anticipate this influence, the optimal trading strategy trades more during these periods.

We note that the optimal solution does not depend on the values of the running market VWAP, M_t, the running personal VWAP, P_t, and the current stock price, S_t, since these quantities are absent from the optimal solutions. This simplification is due to the assumption that the price process is a martingale conditional on realized volume. However, we note that, the value function does depend on the values of P_t, M_t, S_t. In addition, in the case where we consider trading costs, the price process will not be a martingale and the optimal solution will depend on both the market running VWAP and the personal running VWAP.

4 Empirical Analysis

To illustrate the efficiency of the trading strategy we propose we choose a specific model for the bucket volumes, Eq. 2.1, and calibrate it to market data. Early literature on models of trading volume focuses on a daily frequency, e.g., Clark (1973) describes a model for daily volume of cotton futures. More recent work describes intraday models of volume that account for time-of-day effects and autocorrelation; e.g., Białkowski et al. (2008). Compared to the recent literature we offer a model of the shape of the intraday trading volume that allows for correlation between realized trading volume and stock return variance.

4.1 Trading volume and stock return model

We describe a model of normalized intraday trading volume, i.e., the intraday shape of volume, that includes common components across stocks, such as shape factors, day-of-the-week effects, intraday and overnight autoregressive terms and stock-specific residual uncertainty.

The normalization of realized bucket trading volume is chosen such that all stocks have the same expected normalized volume, by modeling $\nu_{b,s}^d \times K_s$, where $\nu_{b,s}^d$ is the realized trading volume on day d, bucket b, for stock s; K_s is given by

$$\frac{1}{K_s} = \frac{1}{|W|} \sum_{d \in W} \sum_{b=1}^{B} \nu_{b,s}^d = \frac{1}{|W|} \sum_{d \in W} V_s^d$$

where V_s^d is the realized daily volume for stock s on day d; W is the window of days over which we calibrate the model and which includes day d; and $|W|$ is the length of the window

W in days.[8] This normalization allows us to use cross-sectional data to calibrate the shape of the volume curve. The model we use is:

$$\begin{aligned}
\log\left(\nu_{b,s}^{d} \times K_s\right) = & f_b + f_{\text{dow}} + \sum_{\text{AR}=1}^{N_{\text{AR}}} \Psi_{\text{AR}} \log\left(\nu_{b-\text{AR},s}^{d} \times K_s\right) \mathcal{I}(b > \text{AR}) \\
& + \Psi_{\text{on}} \log\left(V_{B,s}^{d-1} \times K_s\right) \mathcal{I}(\text{overnight}) \mathcal{I}(b = 1) \\
& + \Psi_{\text{md}} \log\left(V_{B,s}^{d-1} \times K_s\right) \mathcal{I}(\text{multiple days}) \mathcal{I}(b = 1) \\
& + \epsilon_{b,s}^{d} \\
\epsilon_{b,s}^{d} \sim & \mathrm{N}(0, \omega_s^2)
\end{aligned} \tag{4.1}$$

The dependent variable, $\log\left(\nu_{b,s}^{d} \times K_s\right)$, is the logarithm of the normalized volume for stock s, on day d, at time bucket b. The explanatory variables include a stock-independent time-of-day shape factor f_b, and a stock-independent day-of-week adjustment factor f_{dow}. The stock independence assumption for these explanatory variables leads to a parsimonious model of volume, and is reasonable for stocks for which the arrival of information is similar across the time of day and day of the week.

The additional explanatory variables capture autoregressive behavior in volume, both intraday and between market close and market open. The coefficients Ψ_{AR}, of the variables $\log\left(\nu_{b-\text{AR},s}^{d} \times K_s\right) \mathcal{I}(b > \text{AR})$ capture intraday, order AR, autocorrelation between bucket volumes in a fashion that is stock independent, as well as independent of time-of-day and day-of-week.[9]

The coefficients $\Psi_{\text{on}}, \Psi_{\text{md}}$, of the variables $\log\left(V_{B,s}^{d-1} \times K_s\right) \mathcal{I}(\text{overnight}) \mathcal{I}(b = 1)$, and $\log\left(V_{B,s}^{d-1} \times K_s\right) \mathcal{I}(\text{multiple days}) \mathcal{I}(b = 1)$, capture dependence of the opening volume of the day with respect to the cumulative volume, V_B^{d-1}, of the previous trading day, separately

[8] Białkowski et al. (2008) employed a similar normalization in a factor analysis of the shape of the volume curve.

[9] $\mathcal{I}(b > \text{AR})$ is an indicator function that guarantees that only intraday buckets are taken into account. $\mathcal{I}(b = 1)$ is an indicator function that guarantees that the trading volume from the previous day only affects the opening trading volume — at least directly.

for the case of consecutive trading days (overnight – on), and trading days that are multiple days apart (multiple days – md). The residuals, $\epsilon^d_{b,s}$, are i.i.d., normally distributed, variables with stock-dependent variance.[10]

This model is chosen based on its parsimony, as well as the examination of several other models; e.g., models that included stock-dependent autoregressive coefficients, as well as separate terms for the case where trading days are separated by weekends, holidays, or consecutive weekends and holidays. The chosen model was found to perform at least as well, in terms of tracking VWAP, as any of the other models.

The relationship between the variance of stock returns and realized trading volume has been examined in the literature; e.g., see Andersen (1996), Manganelli (2005), and Malinova and Park (2010). The models in the literature postulate a relationship between the logarithm of the return variance and the logarithm of trading volume. We specify the dynamics of the stock price to be log-normal, where the variance of logarithmic returns in each time bucket is determined by the bucket volume excluding personal trades, and the mean is chosen to maintain the stock price as a martingale,[11]

$$
\begin{aligned}
S_{t+1} &= S_t \cdot \varepsilon_{t+1}, \\
\varepsilon_{t+1} &\sim \mathrm{LN}\left(-\frac{1}{2}\sigma^2(\nu_{t+1}), \sigma^2(\nu_{t+1})\right)
\end{aligned}
$$

where ε_{t+1} is a log-normal random variable with mean and variance parameters determined

[10] Given the cross-sectional heteroskedasticity of the residuals, we employ an iterative method of weighted least squares to estimate the model coefficients and standard errors.

[11] To test the assumption that the stock price, conditional on volume, is a martingale we calculated, for each of the 30 stocks in the Dow Jones Industrial Average, the correlation between daily traded volume and daily returns over the 18 months in the data. The assumption that the correlation was equal to zero could not be rejected, at the 95% level of confidence, for 20 out of the 30 stocks. For the remaining stocks, the correlation was found to be significantly negative, with values between -11% and -30%. The negative values of the correlation indicate that negative returns are linked to higher market volume compared to positive returns. However, in our empirical results, we did not find a significant deterioration in the performance of the VWAP strategy we propose for the stocks for which the assumption that the stock price is a martingale, conditional on volume, is violated.

by $\sigma^2(\nu_{t+1})$, which is a non-negative function of the market's volume at time bucket $t+1$ and is abbreviated by σ^2_{t+1}. In line with the literature, we assume that the relationship between the logarithm of the market trading volume, exclusive of personal trades, and logarithmic stock return variance is given by a linear function

$$\log(\sigma^2) = \alpha_s + \beta_s \log(\nu) \tag{4.2}$$

The stock-dependent parameters α_s, β_s for the relationship between bucket volume and bucket variance of logarithmic stock returns are estimated by maximizing the likelihood function for logarithmic stock returns. The details are provided in the Appendix.

4.2 Model Calibration

We apply our calibration procedure to data from the Trade and Quote (TAQ) database. Restricting to the 30 stocks in the Dow Jones Industrial Average, listed in Table 1, and given all the trades for these stocks reported in TAQ during the 18 months between July 1, 2010 and December 30, 2011, we construct 26 time buckets per day of 15 minutes each. For each time bucket we calculate the volume-weighted average price and the bucket volume, normalized according to the procedure in the previous subsection. The procedure is repeated for each day in the data using a rolling 60-day window.

The results of the calibration of the volume model are presented in Table 2 and Fig. 3. The number of autoregressive terms used in Eq. 4.1 is set to one so that $N_{AR} = 1$, following the results in Białkowski et al. (2008). Fig. 3 presents the average shape of the daily normalized volume over all the 60-day rolling windows in the data for each day of the week, as well as the 10% and 90% deciles, while Table 2 presents the corresponding average values and 10% and 90% deciles for each of the mean reversion coefficients.

We note that the coefficients for the overnight and multiple day variables take similar values, indicating that the shape of the volume is not influenced heavily, on average, by the length of the time between market close and market open.

Table 3 presents the results for the calibrated values of the coefficient β in Eq. 4.2, which determines the relationship between realized volume and stock return variance. From the table we notice that the values range between 1 and 1.8 for most stocks, where a value of 1 indicates that if bucket volume doubles then the standard deviation of returns is expected to increase by a factor of 41%.

4.3 VWAP tracking: Empirical results

To evaluate the empirical performance of the optimal strategy given in Eq. 3.2, we first calibrate our volume model each day using a weighted-least-squares regression and then maximize the likelihood function for the logarithmic return model, using the previous 60 days of data. Then the trading strategy in Eq. 3.2 is estimated, at each time bucket, by sampling 100,000 paths from the volume-return model and performance is evaluated against the actual volume path on that day.

As an example, the top panel of Fig. 4 illustrates the bucket volume and the trading strategy for McDonalds stock on September 28, 2010. From the figure we note that the trading strategy lags deviations in volume by one time bucket. While the fluctuations in the trading strategy are more pronounced during the beginning of the day, once much of the daily volume has been realized, the trading strategy becomes smoother towards the end of the day — even though bucket volume deviations may be equally large during the end as during the beginning of the day.

The bottom panel of Fig. 4 illustrates the distribution of the bucket volume and of the trading strategy for McDonalds stock over all the 18 months in the data. From the figure we note that, in line with our theoretical results, the optimal trading strategy changes direction in a small percentage of cases; this reversal is evident for example from the second bucket in the figure. Another striking feature is the large width of the 95% interval for the optimal trading strategy during the second time bucket, due to the large variability of volume during the first time bucket.

Figure 5, summarizes the performance of the optimal strategy across all thirty stocks included in the Dow Jones Industrial Average when the personal volume is 1% (top panel) and 5% (bottom panel) of the expected market volume. The figure plots the standard deviation of the difference between personal and market VWAP for each stock over the 18 month period in the data against the standard deviation of the stock's daily logarithmic returns. The figure illustrates that as the variance of stock returns increases, the performance of the optimal trading strategy deteriorates. Beyond the increased variance of returns, this deterioration is due to increased fluctuations in the shape of the daily volume as the returns become more volatile.

Comparing the top and bottom panels of Figure 5, suggests that, without trading costs, the optimal strategy performs better as the percentage of personal volume increases. This a result due to the fact that as the personal volume increases, it influences market VWAP more.

Tables 4 and 5 display the first four moments of the distribution of $(M_T - P_T)$ in the cases of trading 1% and 5% of expected market volume, respectively, as well as the daily standard deviation of returns and the Sharpe ratio that the VWAP strategy would have achieved assuming the industry practice cost of 10 basis points for implementing the strategy. Both tables illustrate that the distribution has heavy tails for all stocks.[12] In addition, the difference in performance between trading 1% and 5% is relatively minor. While it is surprising that the mean for 28 out of 30 stocks is negative, the values are relatively small, with the largest deviation being less than 3 basis points. Standard deviations of the VWAP

[12] We report the histogram of the distribution for each stock in the Appendix. The histograms illustrate the existence of large outliers.

difference vary little as the volume of personal trades changes, and range from less than 6 basis points, for stocks that with low daily standard deviation of logarithmic returns such as Waste Management and Procter & Gamble, to 27 basis points for Bank of America, which also has the highest daily standard deviation of logarithmic returns.

4.4 Robustness

To test the robustness of our results we consider two tests. In the first test, we increase the number of daily time buckets from 26 time buckets of 15 minutes to 78 time buckets of 5 minutes each. To maintain consistency with the autoregressive model of volume between the original and the new time discretization, we expand the model of market volume from an autoregressive model of order 1, to an autoregressive model of order 3. The first four moments of the distribution of $(M_T - P_T)$ are displayed in Table 6 in the case of 5 minute trading buckets. The results are very similar to the case of the 26, 15 minute, daily time buckets, with slight improvements for most stocks, and a slight deterioration of the performance of the strategy for a few stocks, such as Bank of America.

The second robustness test of the model is to consider a volume model that is stock-specific. The volume for each stock is modeled by an autoregressive Inverse Gaussian distribution. Specifically we model

$$\nu_{1,s}^{d} \sim \mathrm{IG}(f_{1,s} + \psi_{0,s} V_{B,s}^{d-1}, \lambda_s)$$
$$\nu_{b+1,s}^{d} \sim \mathrm{IG}(f_{b+1,s} + \psi_s \nu_{b,s}^{d}, \lambda_s), b = 1, \ldots, B-1$$

where IG is the inverse Gaussian/Wald distribution, $f_{b,s}$ are stock specific parameters that capture changes in the expected volume for different time buckets during the day, ψ_s is a stock specific autoregressive coefficient, and λ_s is the shape parameter of the IG distribution. The volume during the first bucket of the day, $\nu_{1,s}^{d}$ is modeled with separate parameters $f_{1,s}$

and $\psi_{0,s}$, based on the cumulative volume in the stock during the previous trading day, $V_{B,s}^{d-1}$. The remaining intraday bucket volumes, for buckets two to B, depend on the immediately earlier bucket volume. The calibration procedure of this model can be found in the Appendix.

In Table 7 we display the first four moments and the Sharpe ratio of the distribution of $(M_T - P_T)$ under this model when we trade 1% of expected volume of a day making trades every 15 minutes. The IG model performs worse in all cases than the regression model in Section 4.1, in terms of VWAP tracking. The IG model is also much less parsimonious since it has stock-specific intraday volume parameters, and there is much more time variability in the calibrated parameters.

5 Market Impact and Slippage

While the optimal strategy was derived in the context where trades do not impact market prices and there is no execution slippage between market prices and the price at which we are able to execute personal trades, in reality trades are known to both move market prices, and to execute at worse prices than the posted market price. To demonstrate the relevance of the algorithm we propose in this paper, in this section we consider a study of the performance of the strategy when the act of trading has an impact on the price of the stock and, when personal trades may execute at a price different than the market price. The first effect is permanent and affects all market participants: personal trades are considered by market participants to reveal fundamental information about the stock price and are incorporated in the price of the stock. The second effect is transitory and only affects personal trades. The slippage between the price achieved for a personal trade vs. the market price can be due, for example, to the dynamics of the limit order book. Since we are assuming that personal trades are executed through market orders, an order with a large size may exhaust the liquidity at more than one tick levels of the limit order book.[13]

We note that in the case where personal trades affect the market price permanently, trading strategies in which personal VWAP outperforms market VWAP, without any trading skill, are possible. For example, one such strategy is to transact the entire order at the beginning of the day, influencing market price for the entire day, and skewing market VWAP to the benefit of the personal VWAP. These strategies have been recognized in the literature and by practitioners, see Johnsey (2006), and are the basis of the main criticism of using VWAP tracking orders. We note that by avoiding optimization when permanent market impact is possible, we are not subject to this criticism.

5.1 Model

Our framework for modeling permanent impact and slippage is based on the model described by Almgren et al. (2005), which is popular both in the literature and in practice. While the model in Almgren et al. (2005) considers an arithmetic Brownian motion to drive the price, our model is based on geometric Brownian motion and we modify the impact and slippage functions in Almgren et al. (2005) accordingly.

Assuming that the market price of a stock at a time bucket t, is described by a variable S_t, we introduce the notation $\hat{S}_t$ for the price at which a order of c_{t-1} shares is executed during time bucket t. We note that both the market price S_t and the execution price $\hat{S}_t$ are influenced by the trade: the market price of the stock is affected only by the permanent impact of the trade, while the execution price that the trader gets, is affected by both the permanent and temporary impact.

For the permanent impact of a trade, following Almgren et al. (2005), we assume that

[13] We do not consider costs paid to the exchange for removing market liquidity by submitting market orders, or benefits obtained if we were to increase liquidity by submitting limit orders.

permanent impact is captured by a function ξ_1

$$\begin{aligned} S_{t+1} &= S_t \xi_1(c_t) \varepsilon_{t+1}, \\ \xi_1(c_t) &= \exp(\gamma c_t). \end{aligned} \tag{5.1}$$

This formulation has the benefit that the permanent impact of trades to the logarithm of the stock price is linear in the number of shares transacted. Linearity implies that the overall permanent impact of trading a given amount of shares over a time interval, is independent of the way the total number of shares are split into separate trades. This property is in line with the intuition that the information imputed from a trade only depends on the total size of the trade, rather than the way the trade is divided before submitting it to the market. In addition when the permanent impact function is linear, investors cannot manipulate the market by, for example, buying small quantities at one time and then liquidating at a future time.

Execution slippage between the market price and the execution price received for an order is captured by a second function, ξ_2

$$\begin{aligned} \hat{S}_{t+1} &= S_{t+1} \xi_2(c_t, \nu_{t+1}), \\ \xi_2(c_t, \nu_{t+1}) &= 1 + \eta \cdot \operatorname{sign}(c_t) \left| \frac{c_t}{\nu_{t+1}} \right|^{3/5} \end{aligned} \tag{5.2}$$

The functional form of the slippage, expressed in terms of the size of the personal trade as a percentage of the market volume, is common in the literature and is referred to at the three-fifths' rule, see Almgren et al. (2005) and references therein.

Given the price dynamics for the market price and the execution price, the dynamics of the state variables are given by

$$
\begin{aligned}
V_{t+1} &= V_t + \nu_{t+1} + c_t, \\
S_{t+1} &= S_t e^{\gamma c_t} \varepsilon_{t+1}, \\
\hat{S}_{t+1} &= S_{t+1} \left(1 + \eta \cdot \operatorname{sign}(c_t) \left| \frac{c_t}{\nu_{t+1}} \right|^{3/5} \right), \\
M_{t+1} &= \frac{M_t V_t + S_{t+1} \nu_{t+1} + \hat{S}_{t+1} c_t}{V_{t+1}}, \\
P_{t+1} &= \frac{P_t \kappa_t + \hat{S}_{t+1} c_t}{\kappa_{t+1}},
\end{aligned}
\tag{5.3}
$$

where ν_t is the bucket volume exclusive of personal trades, while V_t is the total volume including personal trades. The market VWAP, M_t, accounts for all transactions, including market transaction that occur at the market price, S_t, and personal transactions that occur at the execution price $\hat{S}_t$.

Given the evolution of the state variables, neither the market price nor the execution price are martingales and we are not able to reduce the optimization problem to a problem without the market and personal VWAP state variables.

5.2 Numerical Results

While it is impossible to establish performance of the strategy described in Eq. 3.2 empirically, when trades have both a permanent and a transitory impact, unless we were to implement the strategy in the marketplace, we have tried to use as much empirical information as possible. First, we estimate the trading strategy in Eq. 3.2 similarly to the case without market impact and slippage, by calibrating the parameters for the model for volume and stock price using a 60-day rolling window. We also estimate the parameters that determine permanent market impact and execution slippage; i.e., η and γ in Eqs. 5.1, 5.2, according to the model described in Almgren et al. (2005). We then use the actual bucket volume and value-weighted bucket prices calculated for the TAQ data and modify the market price according to Eq. 5.1, and the execution price according to Eq. 5.2. We calculate market and personal VWAPs using the modified market and execution prices according to Eq. 5.3.

Tables 8 and 9 report the first four moments and Sharpe ratios of the distribution of $(M_T - P_T)$ when market impact and slippage are incorporated into the price model. The tables show that the standard deviation of the difference $(M_T - P_T)$ is similar to the case without trading costs, while the average difference deteriorates as the size of the personal trades increases relative to market volume.

Rather than plotting the standard deviation of the difference between the market and personal VWAPs as a function of the variance of daily logarithmic returns, Figure 6 plots the average difference between market and personal VWAPs for personal volume of 1% (top panel) and 5% (bottom panel). The figure illustrates that, as expected, the under-performance of the personal VWAP relative to the market VWAP deteriorates as the standard deviation of logarithmic returns increases. However, we note that following the strategy given in Eq. 3.2 results in an average difference between the market and personal VWAPs that is below the 10 basis points cost charged by brokerage houses to guarantee VWAP execution. Even when trading up to 5% of expected market volume, the strategy achieves annualized Sharpe ratios between 2.4 and 22.3, indicating that the strategy is viable and competitive compared to the costs imposed by brokerage houses.

6 Conclusion

We have presented a model for intraday trading volume for individual stocks and an optimal trading strategy that minimizes the squared difference between the daily VWAP of the stock and the VWAP achieved by the strategy, and an empirical study of the performance of the trading strategy over an 18-month period for the stocks in the Dow Jones Industrial Average. Our contribution is to derive an optimal strategy for tracking VWAP in closed form, and to empirically determine the performance of the strategy, both with and without trading costs. Our results show that institutional investors, who frequently employ VWAP tracking strategies, can achieve VWAP tracking at costs lower than those reported by brokerage houses.

While we are able to compute the optimal trading strategy in closed-form in a general setting using dynamic programming, allowing us to provide intuition regarding the characteristics of the strategy, challenges remain. We have shown that the trading strategy is optimal when trading costs are ignored: it would be interesting to determine the optimal strategy when market impact costs and slippage are included and to study the difference between the strategies. However, due to the difficulty of solving the problem when market impact costs and slippage are included, it is likely that such a study would be numerical and limited in nature to a few representative cases. Another extension that would be of considerable practical interest would be to allow for both limit and market orders and incorporate the limit order book dynamics in the determination of the optimal trading strategy. We leave these extensions for future work.

A Proof of Lemma 3.1 and Theorem 3.1

A.1 Proof of Lemma 3.1

We have that

$$P_T - M_T = P_t \kappa_t - M_t \frac{V_t}{V_T} + S_t \sum_{i=1}^{T-t} \left(c_{t+i-1} - \frac{\nu_{t+i} + c_{t+i-1}}{V_T} \right) \prod_{j=1}^{i} \varepsilon_{t+j}$$

By iterated expectations and the martingale property of the stock price S

$$\begin{aligned}
E\left[S_t \sum_{i=1}^{T-t} c_{t+i-1} \left(1 - \frac{1}{V_T}\right) \prod_{j=1}^{i} \varepsilon_{t+j} \middle| \mathcal{F}_t \right] &= E\left[\sum_{i=1}^{T-t} c_{t+i-1} \left(1 - \frac{1}{V_T}\right) E\left[S_t \prod_{j=1}^{i} \varepsilon_{t+j} \middle| c_{t+i-1}, V_T \right] \middle| \mathcal{F}_t \right] \\
&= E\left[S_t \sum_{i=1}^{T-t} c_{t+i-1} \left(1 - \frac{1}{V_T}\right) \middle| \mathcal{F}_t \right] \\
&= S_t(1-\kappa_t) - S_t E\left[\frac{\sum_{i=1}^{T-t} c_{t+i-1}}{V_t + \sum_{i=1}^{T-t} (\nu_{t+i} + c_{t+i-1})} \middle| \mathcal{F}_t \right] \\
&= S_t(1-\kappa_t) - S_t(1-\kappa_t) E\left[\frac{1}{V_t + \sum_{i=1}^{T-t} \nu_{t+i} + (1-\kappa_t)} \middle| \mathcal{F}_t \right] \\
&= S_t(1-\kappa_t) \left(1 - E\left[\frac{1}{V_t + \sum_{i=1}^{T-t} \nu_{t+i} + (1-\kappa_t)} \middle| \mathcal{F}_t \right] \right) \\
&= S_t(1-\kappa_t) \left(1 - E\left[\frac{1}{V_T} \middle| \mathcal{F}_t \right] \right)
\end{aligned}$$

where the last equation follows from

$$V_T = V_t + \sum_{i=1}^{T-t} \nu_{t+i} + (1-\kappa_t)$$

This calculation shows that the difference between the personal and market VWAP is independent of the values of the personal trades.

If we start at time bucket zero, then the expectation becomes

$$\begin{aligned}
E[P_T - M_T|\mathcal{F}_0] &= S_0 E\left[1 - \frac{1}{V_T}\middle|\mathcal{F}_0\right] - E\left[\sum_{i=1}^{T}\frac{\nu_i}{V_T}\middle|\mathcal{F}_0\right] \\
&= S_0\left(1 - E\left[\frac{1}{V_T}\left(1 + \sum_{i=1}^{T}\nu_i\right)\middle|\mathcal{F}_0\right]\right) \\
&= S_0\left(1 - E\left[\frac{V_T}{V_T}\middle|\mathcal{F}_0\right]\right) \\
&= 0
\end{aligned}$$

A.2 Proof of Theorem 3.1

The first step is to examine V_T and see that

$$\begin{aligned}
V_T &= V_t + \sum_{i=1}^{T-t}(\nu_{t+i} + c_{t+i-1}) \\
&= V_t + 1 - \kappa_t + \sum_{i=1}^{T-t}\nu_{t+i}.
\end{aligned}$$

This means that V_T does not depend on any individual trade, rather it only depends on the number of shares the trader has left, which is a known quantity and thus the trades do not contribute to the randomness of V_T.

With this result we can next expand G_T using Eq. 2.5 and single out the terms that depend on the trades. We find that

$$
\begin{aligned}
(P_T - M_T)^2 &= \left(P_t\kappa_t - M_t\frac{V_t}{V_T} - S_t\sum_{i=1}^{T-t}\frac{\nu_{t+i}+c_{t+i-1}}{V_T}\prod_{j=1}^{i}\varepsilon_{t+j} + S_t\sum_{i=1}^{T-t}c_{t+i-1}\prod_{j=1}^{i}\varepsilon_{t+j}\right)^2 \quad \text{(A.1)}\\
&= \left(S_t\left(1-\frac{1}{V_T}\right)\sum_{i=1}^{T-t}c_{t+i-1}\prod_{j=1}^{i}\varepsilon_{t+j}\right)^2\\
&+ 2\left(S_t\left(1-\frac{1}{V_T}\right)\sum_{i=1}^{T-t}c_{t+i-1}\prod_{j=1}^{i}\varepsilon_{t+j}\right)\cdot\left(P_t\kappa_t - M_t\frac{V_t}{V_T} - S_t\sum_{i=1}^{T-t}\frac{\nu_{t+i}}{V_T}\prod_{j=1}^{i}\varepsilon_{t+j}\right)\\
&+ \left(P_t\kappa_t - M_t\frac{V_t}{V_T} - S_t\sum_{i=1}^{T-t}\frac{\nu_{t+i}}{V_T}\prod_{j=1}^{i}\varepsilon_{t+j}\right)^2.
\end{aligned}
$$

It is important to remember that c_{t+i-1} is a random variable with respect to $\mathcal{F}_t$ if $i > 1$, because it is a function of state variables that are not yet observed. We can immediately disregard the third summand in the above equation because it has no dependence on the c's. Next we take expectations and consider part of the second term, and find

$$
\begin{aligned}
&E\left[\left(S_t\left(1-\frac{1}{V_T}\right)\sum_{i=1}^{T-t}c_{t+i-1}\prod_{j=1}^{i}\varepsilon_{t+j}\right)\cdot\left(P_t\kappa_t - M_t\frac{V_t}{V_T}\right)\middle|\mathcal{F}_t\right] =\\
&E\left[\left(S_t\left(1-\frac{1}{V_T}\right)\sum_{i=1}^{T-t}c_{t+i-1}\right)\cdot\left(P_t\kappa_t - M_t\frac{V_t}{V_T}\right)\middle|\mathcal{F}_t\right]
\end{aligned}
$$

because

$$
\begin{aligned}
E\left[\left(1-\frac{1}{V_T}\right)\frac{c_{t+i-1}}{V_T}\prod_{j=1}^{i}\varepsilon_{t+j}\middle|\mathcal{F}_t\right] &= E\left[\left(1-\frac{1}{V_T}\right)\frac{c_{t+i-1}}{V_T}E\left[\prod_{j=1}^{i}\varepsilon_{t+j}\middle|V_T, c_{t+i-1}\right]\middle|\mathcal{F}_t\right]\\
&= E\left[\left(1-\frac{1}{V_T}\right)\frac{c_{t+i-1}}{V_T}\cdot 1\middle|\mathcal{F}_t\right]
\end{aligned}
$$

since S_t is a martingale. Now considering the constraint that the c's must sum to 1, we find that

$$E\left[\left(S_t\left(1-\frac{1}{V_T}\right)\sum_{i=1}^{T-t}c_{t+i-1}\prod_{j=1}^{i}\varepsilon_{t+j}\right)\cdot\left(P_t\kappa_t-M_t\frac{V_t}{V_T}\right)\middle|\mathcal{F}_t\right]= \\ E\left[S_t\cdot\left(1-\frac{1}{V_T}\right)\cdot(1-\kappa_t)\cdot\left(P_t\kappa_t-M_t\frac{V_t}{V_T}\right)\middle|\mathcal{F}_t\right] \tag{A.2}$$

which does not involve the c's so this term can be ignored for the optimization, meaning the only part of the value function left is

$$E\Bigg[\left(S_t\left(1-\frac{1}{V_T}\right)\sum_{i=1}^{T-t}c_{t+i-1}\prod_{j=1}^{i}\varepsilon_{t+j}\right)^2- \\ 2\left(S_t\left(1-\frac{1}{V_T}\right)\sum_{i=1}^{T-t}c_{t+i-1}\prod_{j=1}^{i}\varepsilon_{t+j}\right)\cdot\left(S_t\sum_{i=1}^{T-t}\frac{\nu_{t+i}}{V_T}\prod_{j=1}^{i}\varepsilon_{t+j}\right)\middle|\mathcal{F}_t\Bigg].$$

Here each term is multiplied by S_t^2, which is a constant with respect to $\mathcal{F}_t$, so we can eliminate this also and finally we find that we want to minimize

$$E\Bigg[\left(\left(1-\frac{1}{V_T}\right)\sum_{i=1}^{T-t}c_{t+i-1}\prod_{j=1}^{i}\varepsilon_{t+j}\right)^2-2\left(\left(1-\frac{1}{V_T}\right)\sum_{i=1}^{T-t}c_{t+i-1}\prod_{j=1}^{i}\varepsilon_{t+j}\right)\cdot\left(\sum_{i=1}^{T-t}\frac{\nu_{t+i}}{V_T}\prod_{j=1}^{i}\varepsilon_{t+j}\right)\middle|\mathcal{F}_t\Bigg].$$

If we define

$$\hat{c}_t\equiv\{c_i\}_{i=t}^{T-1}, \\ \hat{A}_t\equiv\{A_i\}_{i=t}^{T-1}, \\ \phi\equiv 1-\frac{1}{V_T}$$

then we can rephrase the optimization problem by reducing the expected values of the ε's to functions of the ν's using the law of iterated expectations to find that

$$\hat{G}_t = \min_{\hat{c}_t \in \hat{A}_t} E\Bigg[\phi^2 \sum_{i=1}^{T-t}\sum_{j=1}^{T-t} c_{t+i-1}c_{t+j-1}\text{exp}\left(\sum_{\ell=1}^{\min(i,j)} \sigma_{t+\ell}^2\right) - 2\phi \sum_{i=1}^{T-t}\sum_{j=1}^{T-t} c_{t+i-1}\frac{\nu_{t+j}}{V_T}\text{exp}\left(\sum_{\ell=1}^{\min(i,j)} \sigma_{t+\ell}^2\right)\Bigg|\mathcal{F}_t\Bigg]$$

which simplifies to

$$\hat{G}_t = \min_{\hat{c}_t \in \hat{A}_t} E\left[\phi \sum_{i=1}^{T-t}\sum_{j=1}^{T-t} c_{t+i-1}\left(\phi c_{t+j-1} - 2\frac{\nu_{t+j}}{V_T}\right)\text{exp}\left(\sum_{\ell=1}^{\min(i,j)} \sigma_{t+\ell}^2\right)\Bigg|\mathcal{F}_t\right]. \tag{A.3}$$

Next, we proceed by induction. First we will find the base case, at $T-2$, and then we will prove the induction step. From Eq. A.3, the objective function at $T-2$ is

$$\hat{G}_{T-2} = \min_{c_{T-2}\in A_{T-2}} E\Bigg[\phi c_{T-2}\left(\phi c_{T-2} - 2\frac{\nu_{T-1}}{V_T}\right)e^{\sigma_{T-1}^2} + \phi c_{T-2}\left(\phi c_{T-1} - 2\frac{\nu_T}{V_T}\right)e^{\sigma_{T-1}^2} + \phi c_{T-1}\left(\phi c_{T-2} - 2\frac{\nu_{T-1}}{V_T}\right)e^{\sigma_{T-1}^2} + \phi c_{T-1}\left(\phi c_{T-1} - 2\frac{\nu_T}{V_T}\right)e^{\sigma_{T-1}^2+\sigma_T^2}\Bigg|\mathcal{F}_{T-2}\Bigg].$$

We substitute $c_{T-1} = 1-\kappa_{T-1} = 1-\kappa_{T-2}-c_{T-2}$, set the derivative with respect to c_{T-2} equal to zero, and solve for c_{T-2} and we find that

$$c_{T-2} = 1-\kappa_{T-2} - \frac{E\left[\phi\left(e^{\sigma_{T-1}^2} - e^{\sigma_{T-1}^2+\sigma_T^2}\right)\frac{\nu_T}{V_T}\Big|\mathcal{F}_{T-2}\right]}{E\left[\phi^2\left(e^{\sigma_{T-1}^2} - e^{\sigma_{T-1}^2+\sigma_T^2}\right)\Big|\mathcal{F}_{T-2}\right]}. \tag{A.4}$$

For the induction step we assume that c_j takes the form of Eq. 3.2 for $j = t+1, t+2, ..., T-2$, and $c_{T-1} = 1-\kappa_{T-1}$, we now search for c_t. For notational convenience we will call

$$
\begin{aligned}
e_t &\equiv \frac{\phi\left(e^{\sigma_{t+1}^2} - e^{\sigma_{t+1}^2+\sigma_{t+2}^2}\right)}{E\left[\phi^2\left(e^{\sigma_{t+1}^2} - e^{\sigma_{t+1}^2+\sigma_{t+2}^2}\right)\Big|\mathcal{F}_t\right]}, \\
\gamma_t &\equiv \frac{\nu_t}{V_T}, \\
c_t &= 1 - \kappa_t - E\left[e_t \sum_{i=2}^{T-t} \gamma_{t+i}\Big|\mathcal{F}_t\right].
\end{aligned}
$$

We next consider what happens to the progression of the c's.

$$
\begin{aligned}
c_{T-1} &= 1 - \kappa_{T-1} = 1 - \kappa_{T-2} - c_{T-2} = \frac{E\left[\phi\left(e^{\sigma_{T-1}^2} - e^{\sigma_{T-1}^2+\sigma_T^2}\right)\frac{\nu_T}{V_T}\Big|\mathcal{F}_{T-2}\right]}{E\left[\phi^2\left(e^{\sigma_{T-1}^2} - e^{\sigma_{T-1}^2+\sigma_T^2}\right)\Big|\mathcal{F}_{T-2}\right]} = E\left[e_{T-2}\gamma_T\big|\mathcal{F}_{T-2}\right] \\
c_{T-2} &= 1 - \kappa_{T-2} - \frac{E\left[\phi\left(e^{\sigma_{T-1}^2} - e^{\sigma_{T-1}^2+\sigma_T^2}\right)\frac{\nu_T}{V_T}\Big|\mathcal{F}_{T-2}\right]}{E\left[\phi^2\left(e^{\sigma_{T-1}^2} - e^{\sigma_{T-1}^2+\sigma_T^2}\right)\Big|\mathcal{F}_{T-2}\right]} \\
&= E\left[e_{T-3}\left(\gamma_T + \gamma_{T-1}\right)\big|\mathcal{F}_{T-3}\right] - E\left[e_{T-2}\gamma_T\big|\mathcal{F}_{T-2}\right] \\
&\vdots \\
c_{t+j-1} &= E\left[e_{t+j-2}\sum_{i=j}^{T-t}\gamma_{t+i}\Big|\mathcal{F}_{t+j-2}\right] - E\left[e_{t+j-1}\sum_{i=j+1}^{T-t}\gamma_{t+i}\Big|\mathcal{F}_{t+j-1}\right] \\
&\vdots \\
c_{t+1} &= 1 - \kappa_t - c_t - E\left[e_{t+1}\sum_{j=3}^{T-t}\gamma_{t+j}\Big|\mathcal{F}_{t+1}\right]
\end{aligned}
$$

The last equality on each line is achieved by plugging in the first equality from the line below. We now plug these equations into Eq. A.3 and look at the terms that involve c_t. Eq. A.3

becomes

$$
\begin{aligned}
\hat{G}_t =& E\Bigg[\phi c_t \left(\phi c_t - 2\gamma_{t+1}\right) e^{\sigma^2_{t+1}} + \phi c_t \left(\phi c_{t+1} - 2\gamma_{t+2}\right) e^{\sigma^2_{t+1}} + \phi c_t \sum_{j=3}^{T-t} \left(\phi c_{t+j-1} - 2\gamma_{t+j}\right) e^{\sigma^2_{t+1}} \\
&+ \phi c_{t+1} \left(\phi c_t - 2\gamma_{t+1}\right) e^{\sigma^2_{t+1}} + \phi c_{t+1} \left(\phi c_{t+1} - 2\gamma_{t+2}\right) e^{\sigma^2_{t+1}+\sigma^2_{t+2}} + \phi c_{t+1} \sum_{j=3}^{T-t} \left(\phi c_{t+j-1} - 2\gamma_{t+j}\right) e^{\sigma^2_{t+1}+\sigma^2_{t+2}} \\
&+ \sum_{j=3}^{T-t} \phi c_{t+j-1} \left(\phi c_t - 2\gamma_{t+j}\right) e^{\sigma^2_{t+1}} + \sum_{j=3}^{T-t} \phi c_{t+j-1} \left(\phi c_{t+1} - 2\gamma_{t+j}\right) e^{\sigma^2_{t+1}+\sigma^2_{t+2}} + \Delta_9 \Bigg| \mathcal{F}_t \Bigg].
\end{aligned}
\tag{A.5}
$$

Here Δ_9 is a catch-all variable that is constant with respect to c_t. We re-label each term in the sum from Eq. A.5 as a Δ_i, so that

$$
\begin{aligned}
\Delta_1 &= \phi c_t \left(\phi c_t - 2\gamma_{t+1}\right) e^{\sigma^2_{t+1}} \\
&\vdots \\
\Delta_8 &= \sum_{j=3}^{T-t} \phi c_{t+j-1} \left(\phi c_{t+1} - 2\gamma_{t+j}\right) e^{\sigma^2_{t+1}+\sigma^2_{t+2}}
\end{aligned}
$$

and $\hat{G}_t = E[\Delta_1 + ... + \Delta_8 + \Delta_9 | \mathcal{F}_t]$. We next take the derivative of each term with respect to c_t and find

$$\begin{aligned}
\frac{\partial}{\partial c_t} E\left[\Delta_1 \middle| \mathcal{F}_t\right] &= 2c_t E\left[\phi^2 e^{\sigma^2_{t+1}} \middle| \mathcal{F}_t\right] - 2E\left[\phi\gamma_{t+1} e^{\sigma^2_{t+1}} \middle| \mathcal{F}_t\right] \\
\frac{\partial}{\partial c_t} E\left[\Delta_2 \middle| \mathcal{F}_t\right] &= \frac{\partial}{\partial c_t} c_t E\left[\phi\left(\phi\left(1 - \kappa_t - c_t - E\left[e_{t+1}\sum_{j=3}^{T-t}\gamma_{t+j}\middle|\mathcal{F}_{t+1}\right]\right) - 2\gamma_{t+2}\right) e^{\sigma^2_{t+1}} \middle| \mathcal{F}_t\right] \\
&= E\left[\phi\left(\phi\left(1 - \kappa_t - 2c_t - e_{t+1}\sum_{j=3}^{T-t}\gamma_{t+j}\right) - 2\gamma_{t+2}\right) e^{\sigma^2_{t+1}} \middle| \mathcal{F}_t\right] \\
\frac{\partial}{\partial c_t} E\left[\Delta_4 \middle| \mathcal{F}_t\right] &= \frac{\partial}{\partial c_t} E\Bigg[\phi\left(1 - \kappa_t - c_t - E\left[e_{t+1}\sum_{j=3}^{T-t}\gamma_{t+j}\middle|\mathcal{F}_{t+1}\right]\right) \\
&\times \left(\phi c_t - 2\gamma_{t+1}\right) e^{\sigma^2_{t+1}} \Bigg| \mathcal{F}_t\Bigg] \\
&= E\left[\phi\left(\phi\left(1 - \kappa_t - 2c_t - e_{t+1}\sum_{j=3}^{T-t}\gamma_{t+j}\right) + 2\gamma_{t+1}\right) e^{\sigma^2_{t+1}} \middle| \mathcal{F}_t\right]
\end{aligned}$$

$$\begin{aligned}
\frac{\partial}{\partial c_t} E\left[\Delta_5 \middle| \mathcal{F}_t\right] &= \frac{\partial}{\partial c_t} E\Bigg[\phi\left(1 - \kappa_t - c_t - E\left[e_{t+1}\sum_{j=3}^{T-t}\gamma_{t+j}\middle|\mathcal{F}_{t+1}\right]\right) \\
&\times \left(\phi\left(1 - \kappa_t - c_t - E\left[e_{t+1}\sum_{j=3}^{T-t}\gamma_{t+j}\middle|\mathcal{F}_{t+1}\right]\right) - 2\gamma_{t+2}\right) e^{\sigma^2_{t+2} + \sigma^2_{t+2}} \Bigg| \mathcal{F}_t\Bigg] \\
&= E\left[2\phi\left(\phi\left(\kappa_t - 1 + c_t + E\left[e_{t+1}\sum_{j=3}^{T-t}\gamma_{t+j}\middle|\mathcal{F}_{t+1}\right]\right) + \gamma_{t+2}\right) e^{\sigma^2_{t+1} + \sigma^2_{t+2}} \middle| \mathcal{F}_t\right]
\end{aligned}$$

Now $\Delta_3,\ \Delta_6,\ \Delta_7$ and Δ_8 in Eq. A.5 require a bit more analysis; for this we examine

$$\begin{aligned}
\sum_{j=3}^{T-t} c_{t+j-1} &= \sum_{j=3}^{T-t}\left(E\left[e_{t+j-2}\sum_{i=j}^{T-t}\gamma_{t+i}\middle|\mathcal{F}_{t+j-2}\right] - E\left[e_{t+j-1}\sum_{i=j+1}^{T-t}\gamma_{t+i}\middle|\mathcal{F}_{t+j-1}\right]\right) \\
&= E\left[e_{t+1}\sum_{i=3}^{T-t}\gamma_{t+i}\middle|\mathcal{F}_{t+1}\right] - E\left[e_{t+2}\sum_{i=4}^{T-t}\gamma_{t+i}\middle|\mathcal{F}_{t+2}\right] \\
&+ E\left[e_{t+2}\sum_{i=4}^{T-t}\gamma_{t+i}\middle|\mathcal{F}_{t+2}\right] - E\left[e_{t+3}\sum_{i=5}^{T-t}\gamma_{t+i}\middle|\mathcal{F}_{t+3}\right] + \dots
\end{aligned}$$

This is a telescoping sum that converges to

$$\sum_{j=3}^{T-t} c_{t+j-1} = E\left[e_{t+1}\sum_{j=3}^{T-t}\gamma_{t+j}\middle|\mathcal{F}_{t+1}\right],$$

because $c_{T-1} = E\left[e_{T-2}\gamma_T\middle|\mathcal{F}_{T-2}\right]$. Using iterated expectations the rest of the derivatives in Eq. A.5 now become

$$\begin{aligned}
\tfrac{\partial}{\partial c_t}E\left[\Delta_3\middle|\mathcal{F}_t\right] &= E\left[\phi^2 e_{t+1}\sum_{j=3}^{T-t}\gamma_{t+j}e^{\sigma_{t+1}^2}\middle|\mathcal{F}_t\right] - 2E\left[\phi\sum_{j=3}^{T-t}\gamma_{t+j}e^{\sigma_{t+1}^2}\middle|\mathcal{F}_t\right]\\
\tfrac{\partial}{\partial c_t}E\left[\Delta_6\middle|\mathcal{F}_t\right] &= \tfrac{\partial}{\partial c_t}E\Bigg[\phi\left(1-\kappa_t-c_t-E\left[e_{t+1}\sum_{j=3}^{T-t}\gamma_{t+j}\middle|\mathcal{F}_{t+1}\right]\right)\\
&\qquad\times\sum_{j=3}^{T-t}\left(\phi E\left[e_{t+1}\gamma_{t+j}\middle|\mathcal{F}_{t+1}\right]-2\gamma_{t+j}\right)e^{\sigma_{t+1}^2+\sigma_{t+2}^2}\Bigg|\mathcal{F}_t\Bigg]\\
&= -E\left[\phi\left(\phi E\left[e_{t+1}\sum_{j=3}^{T-t}\gamma_{t+j}\middle|\mathcal{F}_{t+1}\right]-2\sum_{j=3}^{T-t}\gamma_{t+j}\right)e^{\sigma_{t+1}^2+\sigma_{t+2}^2}\middle|\mathcal{F}_t\right]\\
\tfrac{\partial}{\partial c_t}E\left[\Delta_7\middle|\mathcal{F}_t\right] &= E\left[\phi^2 e_{t+1}\sum_{j=3}^{T-t}\gamma_{t+j}e^{\sigma_{t+1}^2}\middle|\mathcal{F}_t\right]\\
\tfrac{\partial}{\partial c_t}E\left[\Delta_8\middle|\mathcal{F}_t\right] &= -E\left[\phi^2 E\left[e_{t+1}\sum_{j=3}^{T-t}\gamma_{t+j}\middle|\mathcal{F}_{t+1}\right]e^{\sigma_{t+1}^2+\sigma_{t+2}^2}\middle|\mathcal{F}_t\right]
\end{aligned}$$

After combining all of these derivatives we find that

$$\begin{aligned}
\tfrac{\partial}{\partial c_t}\hat{G}_t &= E\left[2\phi\left(e^{\sigma_{t+1}^2}-e^{\sigma_{t+1}^2+\sigma_{t+2}^2}\right)\cdot\left(\phi(1-\kappa_t-c_t)-\gamma_{t+2}-\sum_{j=3}^{T-t}\gamma_{t+j}\right)\middle|\mathcal{F}_t\right]\\
&= E\left[2\phi\left(e^{\sigma_{t+1}^2}-e^{\sigma_{t+1}^2+\sigma_{t+2}^2}\right)\cdot\left(\phi(1-\kappa_t-c_t)-\sum_{j=2}^{T-t}\gamma_{t+j}\right)\middle|\mathcal{F}_t\right].
\end{aligned}$$

Now we set this derivative equal to zero to find the minimum, so that

$$E\left[2\phi\left(e^{\sigma_{t+1}^2}-e^{\sigma_{t+1}^2+\sigma_{t+2}^2}\right)\cdot\left(\phi(1-\kappa_t-c_t)-\sum_{j=2}^{T-t}\gamma_{t+j}\right)\Bigg|\mathcal{F}_t\right]=0$$

$$(1-\kappa_t-c_t)E\left[\phi^2\left(e^{\sigma_{t+1}^2}-e^{\sigma_{t+1}^2+\sigma_{t+2}^2}\right)\Bigg|\mathcal{F}_t\right]-E\left[\phi\left(e^{\sigma_{t+1}^2}-e^{\sigma_{t+1}^2+\sigma_{t+2}^2}\right)\sum_{j=2}^{T-t}\gamma_{t+j}\Bigg|\mathcal{F}_t\right]=0$$

and therefore Eq. 3.2 is indeed the optimal choice for c_t.

A.3 Proof of Lemma 3.2

If the variance of stock returns is not dependent on volume, then stock returns and volume are independent. In that case the term $\left(e^{\sigma_{t+1}^2+\sigma_{t+2}^2}-e^{\sigma_{t+1}^2}\right)$ can come out of the expectation in both the numerator and denominator of Equation 3.2 and then the terms cancel each other out.

Next, we consider the case when our trade size is insignificant, as compared to the total market volume. In this situation we have that Equation 2.2 changes to $V_t=\sum_{i=1}^{t}\nu_i$ where our personal volume has not affected the market volume. This also means that Equation 2.4 changes to

$$M_t=\sum_{i=1}^{t}\frac{\nu_i}{V_t}S_i,$$

meaning our personal trades have also not affected the market VWAP. Here, this can be taken further into Equation A.1 so that

$$(P_T-M_T)^2=\left(P_t\kappa_t-M_t\frac{V_t}{V_T}-S_t\sum_{i=1}^{T-t}\frac{\nu_{t+i}}{V_T}\prod_{j=1}^{i}\varepsilon_{t+j}+S_t\sum_{i=1}^{T-t}c_{t+i-1}\prod_{j=1}^{i}\varepsilon_{t+j}\right)^2,$$

which in turn means that $\phi=1$. Taking this all the way through the proof of Theorem 3.1

we find that

$$
\begin{aligned}
c_t &= 1 - \kappa_t - E\left[\sum_{i=2}^{T-t} \frac{\nu_{t+i}}{V_T} \middle| \mathcal{F}_t\right] \\
&= 1 - \kappa_t - E\left[\frac{V_T - V_{t+1}}{V_T} \middle| \mathcal{F}_t\right] \\
&= E\left[\frac{V_{t+1}}{V_T} \middle| \mathcal{F}_t\right] - \kappa_t.
\end{aligned}
$$

B Calibration of the Price Model

To fit the relationship between the volume and price we must calibrate the α and β parameters from Eq. 4.2.

Figure 7 presents plots of the logarithm of the daily volume vs. the logarithm of the daily standard deviation of stock returns for each of the stocks in the Dow Jones Industrial average. The plots indicate that the relationship between volume and returns can be reasonably approximated from the model in Eq. 4.2. Rather than using the plots to determine the values of the parameters α and β, we instead compute their values by maximizing the log likelihood function.

We have assumed that we have D days of data split into B buckets and we are trying to calibrate the parameters of the price model $S_{t+1} = S_t \cdot \varepsilon_{t+1}$. This means that our data is the set of points S_{t+1}/S_t. The data we have from TAQ counts each trade in each bucket and calculates the VWAP over that bucket. We use this calculated VWAP as the price at the end of the bucket. This however does not allow us to calculate the return on the first bucket on each day since we do not consider overnight returns because of splits, dividends and buy backs. Therefore in order to calculate the return over the first bucket we rely on the opening price reported by the exchange on each day. We write the ratio of stock prices on each day and in each bucket as ε_b^d and then the log likelihood function becomes

$$\sum_{b=1}^{B}\sum_{d=1}^{D}\left(-\log(\varepsilon_b^d)-\frac{1}{2}\log\left(2\pi\alpha(\nu_b^d)^\beta\right)-\frac{1}{2\alpha(\nu_b^d)^\beta}\left(\log(\varepsilon_b^d)+\tfrac{1}{2}\alpha(\nu_b^d)^\beta\right)^2\right)$$

which we can maximize, over α and β, using standard numerical packages for each stock.

C Calibration of Inverse Gaussian Volume Model

The density of an IG-distributed random variable, $X \sim IG(\mu, \lambda)$, is given by

$$p(x|\mu,\lambda)=\sqrt{\frac{\lambda}{2\pi x^3}}\exp\left(-\frac{\lambda(x-\mu)^2}{2\mu^2 x}\right)$$

and it's expectation and variance are

$$\mathbb{E}[X]=\mu,$$
$$\mathrm{Var}[X]=\mu^3/\lambda.$$

The parameters $\{f_b\}, \psi, \psi_0, \lambda$ can be found by maximizing the log-likelihood function for each stock. Given B buckets, and D days, the log-likelihood function is given by

$$\begin{aligned}\mathcal{L}(\{f_b\},\lambda,\psi,\psi_0) &= \sum_{d=1}^{D}\left(\tfrac{1}{2}\log\left(\frac{\lambda}{2\pi}\right)-\tfrac{3}{2}\log(\nu_1^d)-\lambda\sum_{d=1}^{D}\frac{(\nu_1^d-f_1-\psi_0 V_B^{d-1})^2}{2\nu_1^d(f_1+\psi_0 V_B^{d-1})^2}\right)\\ &+ \sum_{b=2}^{B}\sum_{d=1}^{D}\left(\tfrac{1}{2}\log\left(\frac{\lambda}{2\pi}\right)-\tfrac{3}{2}\log(\nu_b^d)-\lambda\sum_{d=1}^{D}\frac{(\nu_b^d-f_b-\psi\nu_{b-1}^d)^2}{2\nu_b^d(f_b+\psi\nu_{b-1}^d)^2}\right).\end{aligned}$$

In order to maximize this log-likelihood function we first find the first order conditions for the f_b's and ψ's by differentiating the log-likelihood function with respect to these variables and setting the derivatives equal to zero. The first order conditions for the f's are given by

$$f_1 = \left(\sum_{d=1}^{D} \frac{\nu_1^d - \psi_0 V_B^{d-1}}{(f_1 + \psi_0 V_B^{d-1})^3} \right) \Bigg/ \left(\sum_{d=1}^{D} \frac{1}{(f_1 + \psi_0 V_B^{d-1})^3} \right) \tag{C.1}$$

$$f_b = \left(\sum_{d=1}^{D} \frac{\nu_b^d - \psi \nu_{b-1}^d}{(f_b + \psi \nu_{b-1}^d)^3} \right) \Bigg/ \left(\sum_{d=1}^{D} \frac{1}{(f_b + \psi \nu_{b-1}^d)^3} \right), \quad b = 2, ..., B. \tag{C.2}$$

We can see here that the first order condition for each f_b only depends of that f_b and the ψ's; they do not depend on λ, and thus for any values of ψ and ψ_0 we can solve the first order conditions using any numerical root finder.

Once we have each f we find that the first order condition for λ is

$$\frac{1}{\lambda} = \frac{1}{B \cdot D} \left(\sum_{d=1}^{D} \frac{(\nu_1^d - f_1 - \psi_0 V_B^{d-1})^2}{\nu_1^d (f_1 + \psi_0 V_B^{d-1})^2} + \sum_{b=2}^{B} \sum_{d=1}^{D} \frac{(\nu_b^d - f_b - \psi \nu_{b-1}^d)^2}{\nu_b^d (f_b + \psi \nu_{b-1}^d)^2} \right)$$

which is a closed form, given ψ, ψ_0 and all the f_b's. Therefore given any values of ψ and ψ_0 we can calculate the optimal f_b's and λ, so all that is left is to find the optimal ψ and ψ_0. To do this we must use numerical optimization over these two variables to find the maximal value of the log-likelihood function given the values of the f_b's and λ calculated above. We can do this numerical optimization with any standard non-linear optimization routine.

D Calibration of Market Impact Parameters

The permanent price impact and execution slippage functions were chosen to correspond with the ones in Almgren et al. (2005) and here we go through the steps to estimate γ and η. Almgren et al. (2005) finds that the parameters, γ and η, are functions of a few characteristics of the stock and volume data, namely the volatility of daily returns, the expected number of shares traded by the market and the number of shares outstanding.

We call $\bar{\sigma}$ the volatility of daily returns and Almgren et al. (2005) finds that

$$\eta = 0.142 \cdot \bar{\sigma}.$$

We can easily estimate $\bar{\sigma}$ from the logarithmic returns data. We note however that $\bar{\sigma}$ is not the annualized volatility, but rather the daily volatility. In order to calculate γ we also need to incorporate the expected number of shares traded by the market and the number of shares outstanding. For this we call Θ our expectation of the daily turnover. Here turnover means number of shares traded in a day divided by shares outstanding. Additionally we call $\widetilde{V}$ our expectation of market volume in a day. With these Almgren et al. (2005) finds that

$$\gamma = 0.314 \cdot \frac{\bar{\sigma}}{\widetilde{V}\Theta^{1/4}}.$$

Both Θ and $\widetilde{V}$ can be calculated given our volume data and with this we have found all we need to incorporate market impact into our pricing model.

E Additional Statistics

Figures 8, 9 provide time-series plots of the traded volume and the market VWAPs for each of the stocks in the data over the period of our study. Figure 10 provides additional information on the empirical distribution of the difference between the market and personal VWAP achieved following the optimal strategy, for the case without trading costs, when the personal volume traded equals 1% of the expected market volume. The figure illustrates the existence of significant outliers in the distribution.

References

Admati, A. R., Pfleiderer, P., 1988. A theory of intraday patterns: volume and price variability. Review of Financial Studies 1, 3–40.

Almgren, R., Thum, C., Hauptmann, E., Li, H., 2005. Direct estimation of equity market impact. Risk 62, 18–57.

Anand, A., Irvine, P., Puckett, A., Venkataraman, K., 2012. Performance of institutional trading desks: An analysis of persistence in trading costs. Review of Financial Studies 25, 557–598.

Andersen, T., 1996. Return volatility and trading volume: An information flow interpretation of stochastic volatility. Journal of Finance 51, 169–204.

Australian Securities Exchange, 2010. Algorithmic trading and market access arrangements. ASX Review .

Bank of America, 2007. Marching up the learning curve: The second buy-side algorithmic trading survey.

Berkowitz, S. A., Logue, D. E., Noser Jr., E. A., 1988. The total cost of transactions on the NYSE. The Journal of Finance 43, 97–112.

Białkowski, J., Darolles, S., Le Fol, G., 2008. Improving VWAP strategies: A dynamic volume approach. Journal of Banking and Finance 32, 1709–1722.

Bouchaud, J.-P., Gefen, Y., Potters, M., Wyart, M., 2004. Fluctuations and response in financial markets: the subtle nature of random price changes. Quantitative Finance 4, 176–190.

Brownlees, C. T., Cipollini, F., Gallo, G. M., 2011. Intra-daily volume modeling and prediction for algorithmic trading. Journal of Financial Econometrics 9, 489–518.

Chan, K., Christie, W. G., Schultz, P. H., 1995. Market structure and the intraday pattern of bid-ask spreads for NASDAQ securities. Journal of Business 68, 35–60.

Chan, K., Fong, W., 2000. Trade size, order imbalance, and the volatilityvolume relation. Journal of Financial Economics 57, 247–273.

Chiyachantana, C., Jain, P., Jiang, C., Wood, R., 2004. International evidence on institutional trading behavior and price impact. The Journal of Finance 59, 869–898.

Clark, P., 1973. A subordinated stochastic process model with finite variance for speculative prices. Econometrica 41, 135–155.

Darat, A., Rahman, S., Zhong, M., 2003. Intraday trading volume and return volatility of the DJIA stocks: A note. Journal of Banking and Finance 27, 20352043.

Domowitz, I., Yegerman, H., 2011. Algorithmic trading usage patterns and their costs. Journal of Trading pp. 9–21.

Easley, D., O'Hara, M., 1987. Price, trade size, and information in securities markets. Journal of Financial Economics 19, 6990.

Engle, R. F., Furstenberg, R., Russell, J., 2008. Measuring and modeling execution cost and risk, technical report, University of Chicago Graduate School of Business.

Epps, T., Epps, M., 1976. The stochastic dependence of security price changes and transaction volumes: implications for the mixture-of-distributions hypothesis. Econometrica 44, 305–321.

Foster, D., Viswanathan, S., 1990. A theory of the interday variations in volume, variance, and trading costs in securities markets. Review of Financial Studies 3, 593–624.

Frei, C., Westray, N., 2012. Optimal execution of a VWAP order: a stochastic control approach, working paper.

Fuh, C.-D., Teng, H.-W., Wang, R.-H., 2010. On-line VWAP trading strategies. Sequential Analysis 29, 292–310.

Gallant, A. R., Rossi, P. E., Tauchen, G., 1992. Stock prices and volume. Review of Financial Studies 5, 199–242.

Gatheral, J., Schied, A., 2013. Dynamical models of market impact and algorithms for order execution. In: Fouque, J.-P., Langsam, J. (eds.), *Handbook on Systemic Risk*, Cambridge University Press, to appear.

Gatheral, J., Schied, A., Slynko, A., 2012. Transient linear price impact and Fredholm integral equations. Mathematical Finance 22, 445–474.

Giot, P., Laurent, S., Petitjean, M., 2010. Trading activity, realized volatility and jumps. Journal of Empirical Finance 17, 168–175.

Gouriéroux, C., Jasiak, J., Fol, G. L., 1999. Intra-day market activity. Journal of Financial Markets 2, 193–226.

Humphery-Jenner, M. L., 2011. Optimal VWAP trading under noisy conditions. Journal of Banking and Finance 35, 2319–2329.

Jain, P., Joh, G., 1988. The dependence between hourly prices and trading volume. Journal of Financial and Quantitative Analysis 23, 269–283.

Johnsey, G., 2006. Insights on volume weighted average price. Report available at http://www-ac.northerntrust.com/content/media/attachment/data/white_paper/0603/document/ntgil_whitepaper_transmgmt.pdf.

Kaastra, I., Boyd, M., 1995. Forecasting futures trading volume using neural networks. Journal of Futures Markets 15, 953–970.

Kakade, S., Kearns, M. J., Mansour, Y., Ortiz, L. E., 2004. Competitive algorithms for VWAP and limit order trading. ACM Conference on Electronic Commerce pp. 189–198.

Karpoff, J., Boyd, M., 1987. The relation between price changes and trading volume: a survey. Journal of Financial and Quantitative Analysis 22, 109–126.

Keim, D., Madhavan, A., 1997. Transactions costs and investment style; an inter-exchange analysis of institutional equity trades. Journal of Financial Economics 46, 265–292.

Konishi, H., 2002. Optimal slice of a VWAP trade. Journal of Financial Markets 5, 197–221.

Kyle, A., 1985. Continuous auctions and insider trading. Econometrica 53, 1315–1335.

Lo, A., Wang, J., 2000. Trading volume: Definition, data analysis, and implications for portfolio theory. Review of Financial Studies 13, 257–300.

Madhavan, A. N., 2002. VWAP Strategies Transaction Performance: The Changing Face of Trading, Institutional Investor Inc., pp. 32–38.

Malinova, K., Park, A., 2010. Trading volume in dealer markets. Journal of Financial and Quantitative Analysis 45, 1447–1484.

Manchaldore, J., Palit, I., Soloviev, O., 2010. Wavelet decomposition for intra-day volume dynamics. Quantitative Finance 10, 917–930.

Manganelli, S., 2005. Duration, volume, and volatility impact of trades. Journal of Financial Markets 8, 377–399.

McCulloch, J., Kazakov, V., 2012. Optimal VWAP trading strategy and optimal volume, working paper.

NASDAQ, 2012a. Comments on NASDAQ rulemaking: Order instituting proceedings to determine whether to approve or disapprove proposed rule change to establish "Benchmark Orders" under NASDAQ rule 4751(f), (release no. 34-67655; file no. sr-nasdaq-2012-059). Available at http://www.sec.gov/comments/sr-nasdaq-2012-059/nasdaq2012059.shtml.

NASDAQ, 2012b. Notice of filing of proposed rule change to establish "Benchmark Orders" under NASDAQ rule 4751(f). Available at http://www.sec.gov/rules/sro/nasdaq/2012/34-66972.pdf.

Obizhaeva, A., Wang, J., 2006. Optimal trading strategy and supply/demand dynamics, technical report, MIT Sloan School of Management.

SEC, 2013. Order disapproving proposed rule change to establish "Benchmark Orders" under NASDAQ rule 4751(f). Available at http://www.sec.gov/rules/sro/nasdaq/2013/34-68629.pdf.

Tauchen, G., Pitts, M., 1983. The price variability-volume relationship on speculative markets. Econometrica 51, 485–505.

Table 1: **Dow Jones Industrial Average**

This table specifies the stocks included in the Dow Jones Industrial Average during the period of our study, July 1, 2010 to December 30, 2011.

Ticker	Company Name
AA	Alcoa Inc.
AXP	American Express Company
BA	The Boeing Company
BAC	Bank of America Corporation
CAT	Caterpillar Inc.
CSCO	Cisco Systems, Inc.
CVX	Chevron Corporation
DD	E.I. Du Pont de Nemours and Company
DIS	Walt Disney Co.
GE	General Electric Company
HD	The Home Depot, Inc.
HPQ	Hewlett-Packard Company
IBM	International Business Machines Corporation
INTC	Intel Corporation
JNJ	Johnson & Johnson
JPM	JPMorgan and Chase & Co.
KFT	Kraft Foods Group Inc.
KO	The Coca-Cola Company
MCD	McDonalds Corporation
MMM	3M Co.
MRK	Merck and Co., Inc.
MSFT	Microsoft Corporation
PFE	Pfizer Inc.
PG	The Procter & Gamble Company
T	AT&T Inc.
TRV	The Travelers Companies, Inc.
UTX	United Technologies Corporation
VZ	Verizon Communications Inc.
WMT	Wal-Mart Stores, Inc.
XOM	Exxon Mobil Corporation

Table 2: **Volume model coefficients.**

This table reports the mean value as well as the 10% and 90% deciles of the coefficients for the model of volume. Ψ_1 is the coefficient of the previous bucket volume, Ψ_{on} the coefficient of the volume of the previous trading day when there the previous day was a trading day, and Ψ_{md} is the coefficient of the volume of the previous trading day, when the previous day was not a trading day either due to a weekend, a holiday, or both. The values are calculated for the 30 stocks in the Dow Jones Industrial Average over consecutive 60-day rolling windows during the period July 1, 2010 to December 30, 2011.

	10%	Average	90%
Ψ_1	0.56	0.62	0.71
Ψ_{on}	0.45	0.59	0.73
Ψ_{md}	0.42	0.57	0.71

Table 3: **Coefficients** β_s

This table reports the mean value, along with the values of the 10% and 90% deciles, of the coefficient β_s in the relationship between the variance of stocks returns and market volume, for each stock of the Dow Jones Industrial Average. The values are calculated over consecutive 60-day rolling windows during the period July 1, 2010 to December 30, 2011.

Ticker	10%	mean	90%
AA	1.20	1.32	1.53
AXP	0.95	1.13	1.36
BA	1.02	1.21	1.39
BAC	1.29	1.53	1.81
CAT	1.25	1.42	1.70
CSCO	0.89	1.11	1.36
CVX	1.26	1.42	1.67
DD	1.06	1.16	1.38
DIS	0.95	1.08	1.24
GE	1.27	1.37	1.48
HD	1.07	1.18	1.32
HPQ	1.18	1.32	1.57
IBM	1.22	1.33	1.48
INTC	1.09	1.29	1.47
JNJ	1.21	1.33	1.54
JPM	1.32	1.46	1.61
KFT	0.71	0.96	1.26
KO	1.03	1.27	1.50
MCD	1.03	1.19	1.38
MMM	1.07	1.19	1.39
MRK	0.95	1.11	1.46
MSFT	1.11	1.29	1.58
PFE	1.00	1.14	1.42
PG	1.11	1.23	1.50
T	1.11	1.23	1.44
TRV	0.91	1.07	1.33
UTX	1.07	1.16	1.38
VZ	1.16	1.33	1.45
WMT	1.29	1.41	1.52
XOM	1.24	1.36	1.58

Table 4: **No trading costs, trading 1% of expected volume**

This table displays characteristics of the empirical distribution of the difference between the market VWAP and the VWAP achieved following the strategy described in this paper for each of the stocks of the Dow Jones Industrial Average. The distribution was determined by implementing the VWAP strategy on 15 minute intervals on data from the TAQ database, where the volume model was calibrated in 60-day rolling windows. The table reports the ticker for the stock (Ticker), the mean (Mean), standard deviation (Stdev), skewness (Skew), and kurtosis (Kurt) of the difference of VWAPs, as well as the annualized Sharpe ratio (Sharpe) that would be achieved by the strategy assuming compensation of 10 basis points for implementing the strategy, and the daily empirical standard deviation (Vol) of returns for each stock. The amount traded is equal to 1% of the expected daily volume of each stock, and there are no trading costs. The mean and standard deviation are reported in basis points.

Ticker	**Mean**	**Stdev**	**Skew**	**Kurt**	**Sharpe**	**Vol**
AA	-2.6	15.6	-2.9	29.2	7.6	1.9%
AXP	-2.1	18.6	-8.2	97.9	6.7	1.5%
BA	-0.9	11.3	-1.6	17.2	12.8	1.4%
BAC	-2.2	27.1	-5.9	60.9	4.6	2.4%
CAT	-1.8	10.1	-2.1	14.8	12.8	1.6%
CSCO	-1.2	11.1	-1.1	16.7	12.6	1.3%
CVX	-1.5	11.7	-3.2	29.5	11.5	1.1%
DD	-1.6	11.2	-1.9	24.9	11.9	1.3%
DIS	-1.3	15.5	-2.9	79.4	8.9	1.2%
GE	-0.7	11.5	-0.2	16.3	12.8	1.4%
HD	-0.7	10.6	0.1	20.9	13.9	1.4%
HPQ	0.6	18.7	8.6	98.8	8.0	1.5%
IBM	-1.1	7.7	-1.7	18.0	18.5	0.9%
INTC	-0.6	8.7	0.7	19.6	17.2	1.2%
JNJ	-0.3	7.1	4.1	56.9	21.6	0.8%
JPM	-1.4	15.9	-1.2	41.1	8.7	1.8%
KFT	-0.7	9.1	-5.7	78.8	16.1	0.9%
KO	-0.3	7.1	0.3	26.4	21.7	0.8%

Continued on next page

Table 4 – **continued from previous page**

Ticker	Mean	Stdev	Skew	Kurt	Sharpe	Vol
MCD	-0.7	6.5	-2.8	31.2	22.6	0.8%
MMM	-1.6	9.2	-2.1	21.9	14.4	1.2%
MRK	-1.0	8.6	-1.9	32.7	16.7	0.9%
MSFT	-0.4	7.7	-0.5	16.7	19.8	1.1%
PFE	-1.1	10.6	-3.7	50.2	13.3	1.1%
PG	-0.6	5.9	-0.5	18.5	25.3	0.7%
T	-0.9	6.9	-3.4	27.7	20.9	0.8%
TRV	-0.7	10.6	-2.7	52.0	14.0	1.1%
UTX	-0.7	8.2	-1.7	20.8	18.0	1.1%
VZ	-0.6	6.7	-3.4	36.1	22.3	0.9%
WMT	-0.7	5.7	-2.1	26.1	26.2	0.8%
XOM	0.1	9.7	-0.7	23.4	16.1	0.9%

Table 5: **No trading costs, trading 5% of expected volume**

This table displays characteristics of the empirical distribution of the difference between the market VWAP and the VWAP achieved following the strategy described in this paper for each of the stocks of the Dow Jones Industrial Average. The distribution was determined by implementing the VWAP strategy on 15 minute intervals on data from the TAQ database, where the volume model was calibrated in 60-day rolling windows. The table reports the ticker for the stock (Ticker), the mean (Mean), standard deviation (Stdev), skewness (Skew), and kurtosis (Kurt) of the difference of VWAPs, as well as the annualized Sharpe ratio (Sharpe) that would be achieved by the strategy assuming compensation of 10 basis points for implementing the strategy, and the daily empirical standard deviation (Vol) of returns for each stock. The amount traded is equal to 5% of the expected daily volume of each stock, and there are no trading costs. The mean and standard deviation are reported in basis points.

Ticker	**Mean**	**Stdev**	**Skew**	**Kurt**	**Sharpe**	**Vol**
AA	-2.5	15.2	-3.0	30.1	7.9	1.9%
AXP	-2.1	18.3	-8.4	101.1	6.9	1.5%
BA	-0.9	11.0	-1.6	17.6	13.2	1.4%
BAC	-2.2	26.6	-6.0	62.5	4.7	2.4%
CAT	-1.8	9.9	-2.2	15.2	13.2	1.6%
CSCO	-1.1	10.9	-1.2	17.4	13.0	1.3%
CVX	-1.5	11.4	-3.3	30.3	11.9	1.1%
DD	-1.6	10.9	-2.0	25.8	12.3	1.3%
DIS	-1.3	15.2	-2.9	81.5	9.1	1.2%
GE	-0.7	11.2	-0.3	16.7	13.2	1.4%
HD	-0.7	10.3	0.1	21.4	14.3	1.4%
HPQ	0.6	18.4	8.7	100.5	8.1	1.5%
IBM	-1.0	7.5	-1.8	18.5	19.1	0.9%
INTC	-0.6	8.4	0.7	20.5	17.8	1.2%
JNJ	-0.3	6.9	4.2	58.4	22.2	0.8%
JPM	-1.3	15.5	-1.3	42.4	8.9	1.8%
KFT	-0.7	9.0	-5.9	82.0	16.4	0.9%
KO	-0.3	7.0	0.2	26.9	22.2	0.8%

Continued on next page

Table 5 – **continued from previous page**

Ticker	Mean	Stdev	Skew	Kurt	Sharpe	Vol
MCD	-0.7	6.4	-3.0	32.9	23.2	0.8%
MMM	-1.6	9.0	-2.2	23.1	14.9	1.2%
MRK	-0.9	8.4	-1.9	33.5	17.2	0.9%
MSFT	-0.4	7.4	-0.6	17.3	20.5	1.1%
PFE	-1.1	10.4	-3.8	51.3	13.6	1.1%
PG	-0.6	5.7	-0.5	19.1	26.1	0.7%
T	-0.9	6.7	-3.5	28.5	21.5	0.8%
TRV	-0.6	10.3	-2.8	53.8	14.4	1.1%
UTX	-0.6	8.0	-1.7	21.4	18.5	1.1%
VZ	-0.6	6.5	-3.5	37.6	22.9	0.9%
WMT	-0.7	5.5	-2.2	27.1	26.9	0.8%
XOM	0.1	9.5	-0.8	24.0	16.5	0.9%

Table 6: **No trading costs, trading at 5 minute intervals**

This table displays characteristics of the empirical distribution of the difference between the market VWAP and the VWAP achieved following the strategy described in this paper for each of the stocks of the Dow Jones Industrial Average. The distribution was determined by implementing the VWAP strategy on 5 minute intervals on data from the TAQ database, where the volume model was calibrated in 60-day rolling windows. The table reports the ticker for the stock (Ticker), the mean (Mean), standard deviation (Stdev), skewness (Skew), and kurtosis (Kurt) of the difference of VWAPs, as well as the annualized Sharpe ratio (Sharpe) that would be achieved by the strategy assuming compensation of 10 basis points for implementing the strategy, and the daily empirical standard deviation (Vol) of returns for each stock. The amount traded is equal to 1% of the expected daily volume of each stock, and there are no trading costs. The mean and standard deviation are reported in basis points.

Ticker	**Mean**	**Stdev**	**Skew**	**Kurt**	**Sharpe**	**Vol**
AA	-2.3	15.4	-2.7	28.3	7.9	1.9%
AXP	-2.2	18.6	-8.2	97.8	6.7	1.5%
BA	-0.8	11.3	-1.5	16.3	13.0	1.4%
BAC	-2.3	28.6	-6.0	59.3	4.3	2.4%
CAT	-1.8	10.5	-1.9	13.4	12.5	1.6%
CSCO	-1.0	11.7	-0.9	18.3	12.3	1.3%
CVX	-1.5	11.7	-3.1	29.0	11.4	1.1%
DD	-1.4	11.5	-1.6	23.1	11.9	1.3%
DIS	-1.1	15.7	-1.8	78.3	9.0	1.2%
GE	-0.9	11.6	-0.3	16.1	12.4	1.4%
HD	-0.5	10.4	0.4	22.1	14.4	1.4%
HPQ	-0.2	15.9	10.9	167.9	9.9	1.5%
IBM	-1.2	8.2	-1.7	18.9	17.1	0.9%
INTC	-0.6	8.9	0.8	23.4	16.8	1.2%
JNJ	-0.4	7.3	4.3	60.5	20.9	0.8%
JPM	-1.4	16.4	-1.8	39.8	8.4	1.8%
KFT	-0.8	9.7	-7.5	106.1	15.1	0.9%
KO	-0.3	7.3	0.4	26.6	21.0	0.8%

Continued on next page

Table 6 – **continued from previous page**

Ticker	**Mean**	**Stdev**	**Skew**	**Kurt**	**Sharpe**	**Vol**
MCD	-0.8	6.7	-3.0	32.4	22.0	0.8%
MMM	-1.6	9.8	-3.2	34.2	13.6	1.2%
MRK	-1.1	8.9	-2.1	30.3	15.8	0.9%
MSFT	-0.5	8.1	0.0	20.6	18.6	1.1%
PFE	-1.2	10.2	-3.7	51.9	13.7	1.1%
PG	-0.6	5.9	-0.3	20.2	25.4	0.7%
T	-1.2	7.8	-4.9	48.1	17.9	0.8%
TRV	-0.6	10.6	-2.1	47.1	14.1	1.1%
UTX	-0.8	8.6	-1.7	19.8	17.0	1.1%
VZ	-0.8	6.8	-3.6	38.5	21.5	0.9%
WMT	-0.7	5.7	-2.3	26.7	25.7	0.8%
XOM	0.1	9.6	-0.4	23.4	16.3	0.9%

Table 7: **No trading costs, IG volume model**

This table displays characteristics of the empirical distribution of the difference between the market VWAP and the VWAP achieved following the strategy described in this paper for each of the stocks of the Dow Jones Industrial Average. The distribution was determined by implementing the VWAP strategy on 15 minute intervals on data from the TAQ database, where the volume model was calibrated in 60-day rolling windows. The volume model is described in Section 4.4.4. The table reports the ticker for the stock (Ticker), the mean (Mean), standard deviation (Stdev), skewness (Skew), and kurtosis (Kurt) of the difference of VWAPs, as well as the annualized Sharpe ratio (Sharpe) that would be achieved by the strategy assuming compensation of 10 basis points for implementing the strategy, and the daily empirical standard deviation (Vol) of returns for each stock. The amount traded is equal to 1% of the expected daily volume of each stock, and there are no trading costs. The mean and standard deviation are reported in basis points.

Ticker	**Mean**	**Stdev**	**Skew**	**Kurt**	**Sharpe**	**Vol**
AA	-2.7	18.6	-2.1	24.0	6.2	1.9%
AXP	-2.0	18.3	-7.1	81.1	6.9	1.5%
BA	-0.8	11.9	-1.4	17.1	12.2	1.4%
BAC	-3.1	30.5	-4.5	44.8	3.6	2.4%
CAT	-2.0	11.9	-2.1	14.4	10.7	1.6%
CSCO	-1.3	11.8	-1.2	16.1	11.7	1.3%
CVX	-1.4	12.3	-2.8	28.2	11.1	1.1%
DD	-1.6	11.9	-1.3	23.3	11.3	1.3%
DIS	-1.4	15.4	-5.2	80.9	8.9	1.2%
GE	-0.6	12.6	0.2	15.6	11.8	1.4%
HD	-0.7	11.1	0.5	21.2	13.4	1.4%
HPQ	0.5	21.8	11.2	164.4	7.0	1.5%
IBM	-0.9	7.9	-1.6	16.9	18.5	0.9%
INTC	-0.3	9.7	0.3	15.7	15.9	1.2%
JNJ	-0.2	7.3	4.7	62.3	21.4	0.8%
JPM	-1.5	17.6	-0.3	33.1	7.7	1.8%
KFT	-0.6	9.6	-4.9	74.4	15.5	0.9%
KO	-0.2	7.5	0.5	29.9	20.7	0.8%

Continued on next page

Table 7 – **continued from previous page**

Ticker	**Mean**	**Stdev**	**Skew**	**Kurt**	**Sharpe**	**Vol**
MCD	-0.7	6.8	-2.7	29.3	21.8	0.8%
MMM	-1.6	9.8	-1.4	18.2	13.7	1.2%
MRK	-0.9	8.8	-1.0	30.9	16.4	0.9%
MSFT	-0.2	8.0	-0.2	14.8	19.4	1.1%
PFE	-1.0	11.2	-3.2	47.3	12.7	1.1%
PG	-0.6	5.7	0.2	17.4	26.0	0.7%
T	-0.9	7.1	-3.0	25.7	20.5	0.8%
TRV	-0.6	10.0	-1.9	49.2	14.8	1.1%
UTX	-0.6	8.6	-1.8	21.9	17.4	1.1%
VZ	-0.5	6.8	-2.7	30.2	22.1	0.9%
WMT	-0.6	5.8	-1.9	26.5	25.9	0.8%
XOM	0.0	9.2	-0.5	25.5	17.3	0.9%

Table 8: **Trading costs, trading 1% of expected market volume**

This table displays characteristics of the empirical distribution of the difference between the market VWAP and the VWAP achieved following the strategy described in this paper for each of the stocks of the Dow Jones Industrial Average. The distribution was determined by implementing the VWAP strategy on 15 minute intervals on data from the TAQ database, where the volume model was calibrated in 60-day rolling windows. The table reports the ticker for the stock (Ticker), the mean (Mean), standard deviation (Stdev), skewness (Skew), and kurtosis (Kurt) of the difference of VWAPs, as well as the annualized Sharpe ratio (Sharpe) that would be achieved by the strategy assuming compensation of 10 basis points for implementing the strategy, and the daily empirical standard deviation (Vol) of returns for each stock. The amount traded is equal to 1% of the expected daily volume of each stock. Each trade has a permanent effect on market price and is subject to execution slippage, according to the model in Section 5. The mean and standard deviation are reported in basis points.

Ticker	**Mean**	**Stdev**	**Skew**	**Kurt**	**Sharpe**	**Vol**
AA	-4.0	15.6	-2.9	29.3	6.1	1.9%
AXP	-3.3	18.6	-8.2	97.7	5.7	1.5%
BA	-2.0	11.3	-1.6	17.2	11.3	1.4%
BAC	-3.7	27.1	-5.9	61.8	3.7	2.4%
CAT	-3.1	10.2	-2.2	15.0	10.8	1.6%
CSCO	-2.3	11.1	-1.1	16.5	11.0	1.3%
CVX	-2.5	11.7	-3.3	29.5	10.2	1.1%
DD	-2.7	11.2	-2.0	24.8	10.4	1.3%
DIS	-2.4	15.5	-2.9	79.6	7.8	1.2%
GE	-1.8	11.5	-0.2	16.3	11.3	1.4%
HD	-1.7	10.6	0.1	20.7	12.4	1.4%
HPQ	-0.6	18.5	8.5	98.1	8.1	1.5%
IBM	-1.8	7.7	-1.7	17.8	16.9	0.9%
INTC	-1.6	8.7	0.7	19.4	15.3	1.2%
JNJ	-1.0	7.1	4.1	56.7	20.2	0.8%
JPM	-2.6	15.9	-1.3	41.6	7.4	1.8%

Continued on next page

Table 8 – **continued from previous page**

Ticker	Mean	Stdev	Skew	Kurt	Sharpe	Vol
KFT	-1.5	9.1	-5.7	78.4	14.8	0.9%
KO	-0.9	7.2	0.2	25.9	20.2	0.8%
MCD	-1.4	6.5	-2.8	30.9	21.0	0.8%
MMM	-2.5	9.2	-2.1	21.5	12.9	1.2%
MRK	-1.7	8.6	-1.9	32.7	15.3	0.9%
MSFT	-1.4	7.7	-0.6	16.8	17.7	1.1%
PFE	-2.1	10.6	-3.7	50.3	11.8	1.1%
PG	-1.2	5.9	-0.5	18.4	23.6	0.7%
T	-1.7	6.9	-3.4	28.0	19.2	0.8%
TRV	-1.5	10.6	-2.8	52.6	12.7	1.1%
UTX	-1.6	8.2	-1.7	20.5	16.3	1.1%
VZ	-1.4	6.7	-3.4	35.5	20.4	0.9%
WMT	-1.3	5.6	-2.1	26.1	24.3	0.8%
XOM	-0.7	9.7	-0.7	23.2	15.2	0.9%

Table 9: **Trading costs, trading 5% of expected market volume**

This table displays characteristics of the empirical distribution of the difference between the market VWAP and the VWAP achieved following the strategy described in this paper for each of the stocks of the Dow Jones Industrial Average. The distribution was determined by implementing the VWAP strategy on 15 minute intervals on data from the TAQ database, where the volume model was calibrated in 60-day rolling windows. The table reports the ticker for the stock (Ticker), the mean (Mean), standard deviation (Stdev), skewness (Skew), and kurtosis (Kurt) of the difference of VWAPs, as well as the annualized Sharpe ratio (Sharpe) that would be achieved by the strategy assuming compensation of 10 basis points for implementing the strategy, and the daily empirical standard deviation (Vol) of returns for each stock. The amount traded is equal to 5% of the expected daily volume of each stock. Each trade has a permanent effect on market price and is subject to execution slippage, according to the model in Section 5. The mean and standard deviation are reported in basis points.

Ticker	**Mean**	**Stdev**	**Skew**	**Kurt**	**Sharpe**	**Vol**
AA	-6.1	15.2	-2.9	29.8	4.0	1.9%
AXP	-4.9	18.2	-8.3	99.7	4.4	1.5%
BA	-3.5	11.0	-1.5	17.2	9.4	1.4%
BAC	-6.0	26.7	-6.1	63.9	2.4	2.4%
CAT	-5.0	9.9	-2.2	15.2	8.1	1.6%
CSCO	-3.9	10.9	-1.1	16.6	8.9	1.3%
CVX	-3.9	11.4	-3.3	29.6	8.6	1.1%
DD	-4.2	10.8	-2.0	25.1	8.5	1.3%
DIS	-3.8	15.2	-2.8	80.6	6.4	1.2%
GE	-3.4	11.2	-0.1	16.7	9.3	1.4%
HD	-3.2	10.4	0.2	20.8	10.4	1.4%
HPQ	-2.4	18.2	8.6	100.2	6.6	1.5%
IBM	-2.9	7.5	-1.7	17.4	14.9	0.9%
INTC	-3.2	8.5	0.8	20.2	12.7	1.2%
JNJ	-1.9	7.0	4.2	58.9	18.5	0.8%
JPM	-4.4	15.5	-1.3	42.9	5.7	1.8%

Continued on next page

Table 9 – **continued from previous page**

Ticker	Mean	Stdev	Skew	Kurt	Sharpe	Vol
KFT	-2.5	9.0	-5.8	80.3	13.3	0.9%
KO	-1.8	7.0	0.3	25.7	18.5	0.8%
MCD	-2.4	6.3	-2.8	31.2	19.1	0.8%
MMM	-3.8	9.0	-2.1	21.8	10.9	1.2%
MRK	-2.9	8.3	-1.9	33.4	13.6	0.9%
MSFT	-2.8	7.6	-0.5	17.0	15.0	1.1%
PFE	-3.4	10.4	-3.7	50.9	10.0	1.1%
PG	-2.1	5.7	-0.5	18.4	21.9	0.7%
T	-2.7	6.7	-3.4	28.4	17.2	0.8%
TRV	-2.7	10.3	-2.9	54.2	11.2	1.1%
UTX	-2.9	8.0	-1.6	20.2	14.1	1.1%
VZ	-2.5	6.6	-3.3	34.6	18.0	0.9%
WMT	-2.3	5.5	-2.1	25.8	22.3	0.8%
XOM	-1.8	9.5	-0.7	23.3	13.6	0.9%

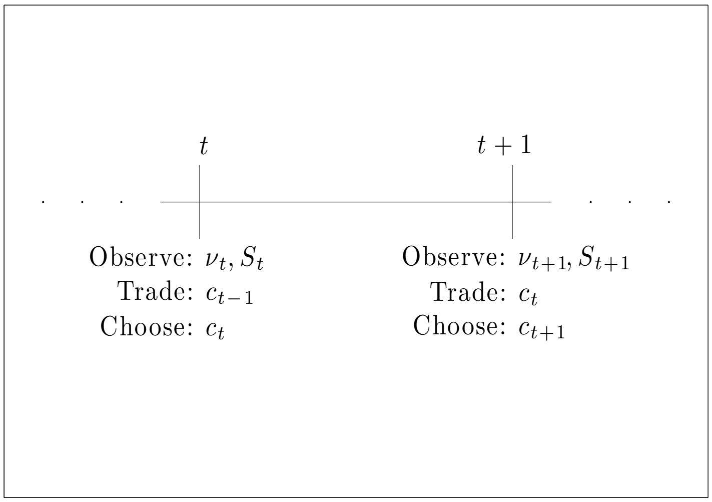

Figure 1: This figure provides a timeline for the events in our model. At time t, the market price S_t, and the market trading volume ν_t, are observed, and the personal trade of volume c_{t-1}, is executed at the market price, S_t. The market and personal VWAPs, M_t, P_t, and the running market and personal volumes V_t, κ_t, are updated, and the personal volume to be traded at time $t+1$, c_t, is chosen. The same events occur at time $t+1$, until the terminal time $t = T$.

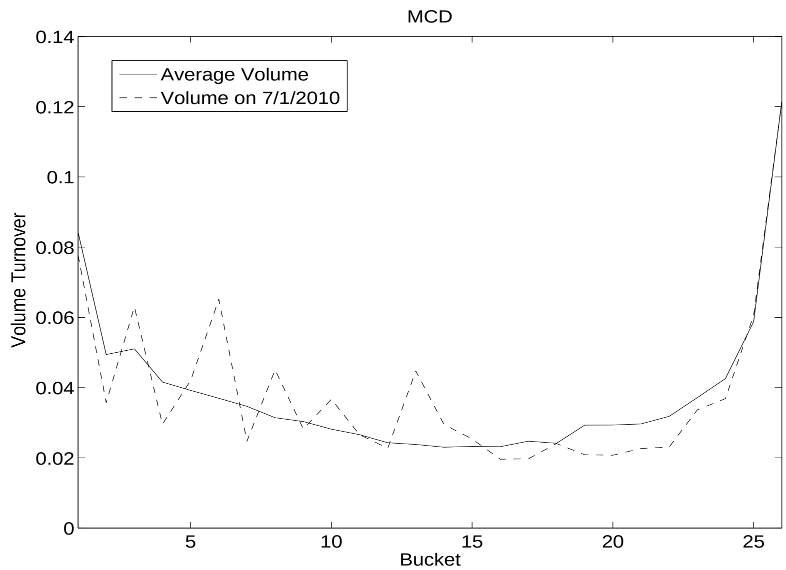

Figure 2: McDonald's volume. The solid line presents the average volume in each time bucket, averaged over all the 18 months in the data, July 1, 2010, to December 30, 2011, as a percentage of total daily volume. The dotted line presents the bucket volume for McDonald's on the day of July 1st, 2010.

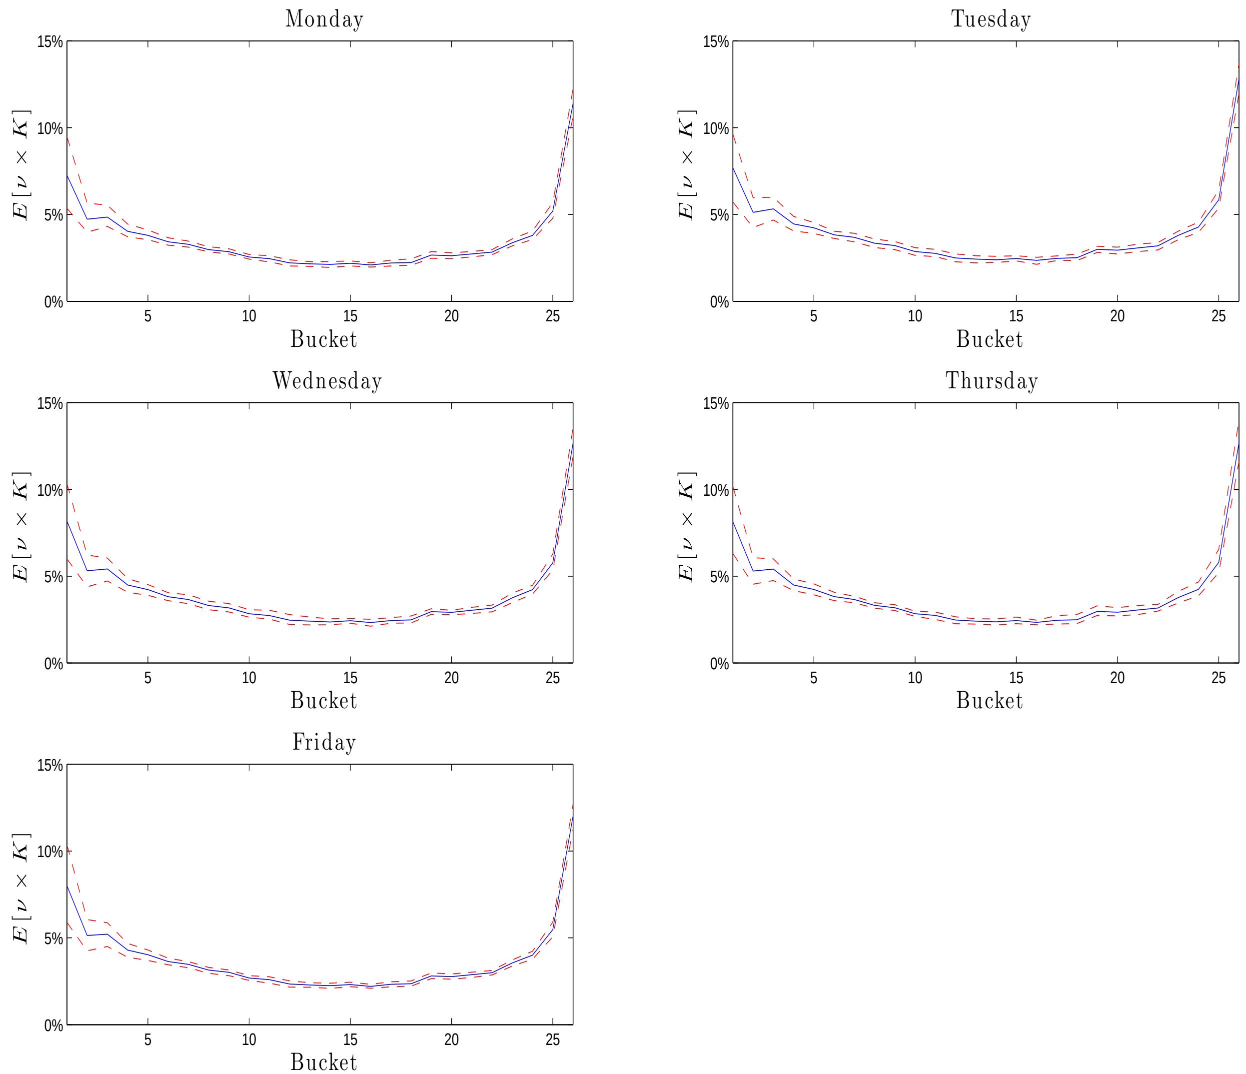

Figure 3: Model calculated, average value of the expected normalized volume per time bucket on each day of the week, along with 10% and 90% deciles of the distribution of trading volume during the 18 months in the data, July 1, 2010, to December 30, 2011.

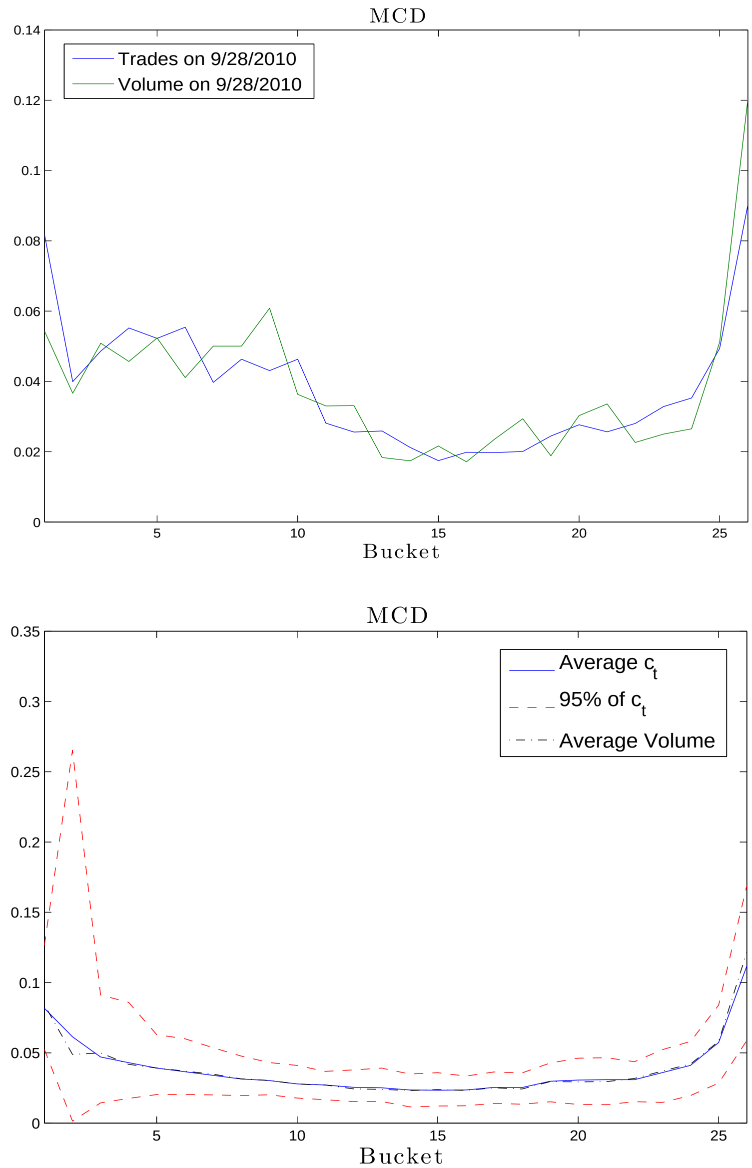

Figure 4: Bucket volume and optimal trading strategy for McDonalds. The top panel presents the market bucket volume (dotted line) and the personal traded volume according to the optimal VWAP tracking strategy (solid line) for McDonalds stock on September 28, 2010. The bottom panel presents the average bucket volume, and the average bucket optimal personal volume and 95% interval, averaged over the 18 months in the data. The total personal traded volume is equal to 1% of the expected market volume for the day.

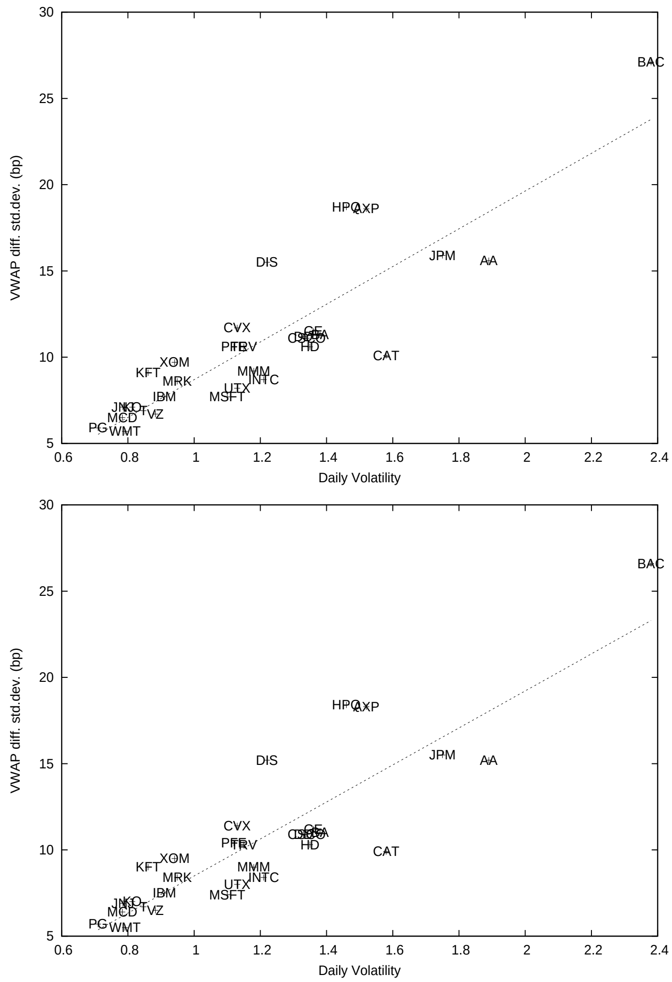

Figure 5: Standard deviation of the difference between market and personal VWAP vs. daily standard deviation (volatility) of logarithmic returns for the thirty stocks of the Dow Jones Industrial Average. The daily personal volume is equal to 1% of the expected daily market volume for each stock in the top panel, and 5% of the expected daily market volume in the bottom panel.

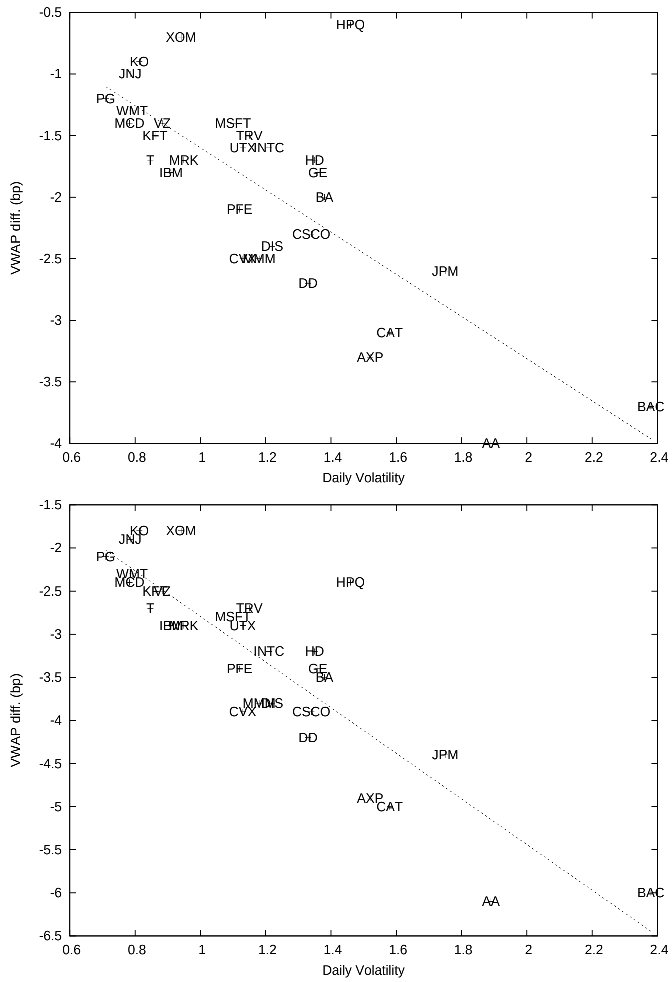

Figure 6: Average difference between market and personal VWAP vs. daily standard deviation (volatility) of logarithmic returns for the thirty stocks of the Dow Jones Industrial Average when both permanent market impact and slippage are included. The daily personal volume is equal to 1% of the expected daily market volume for each stock in the top panel, and 5% of the expected daily market volume in the bottom panel.